PULSED FLYBACK GENERATOR

Assembly & Guidance Manual

by

Julian Perry MSc, PGCE

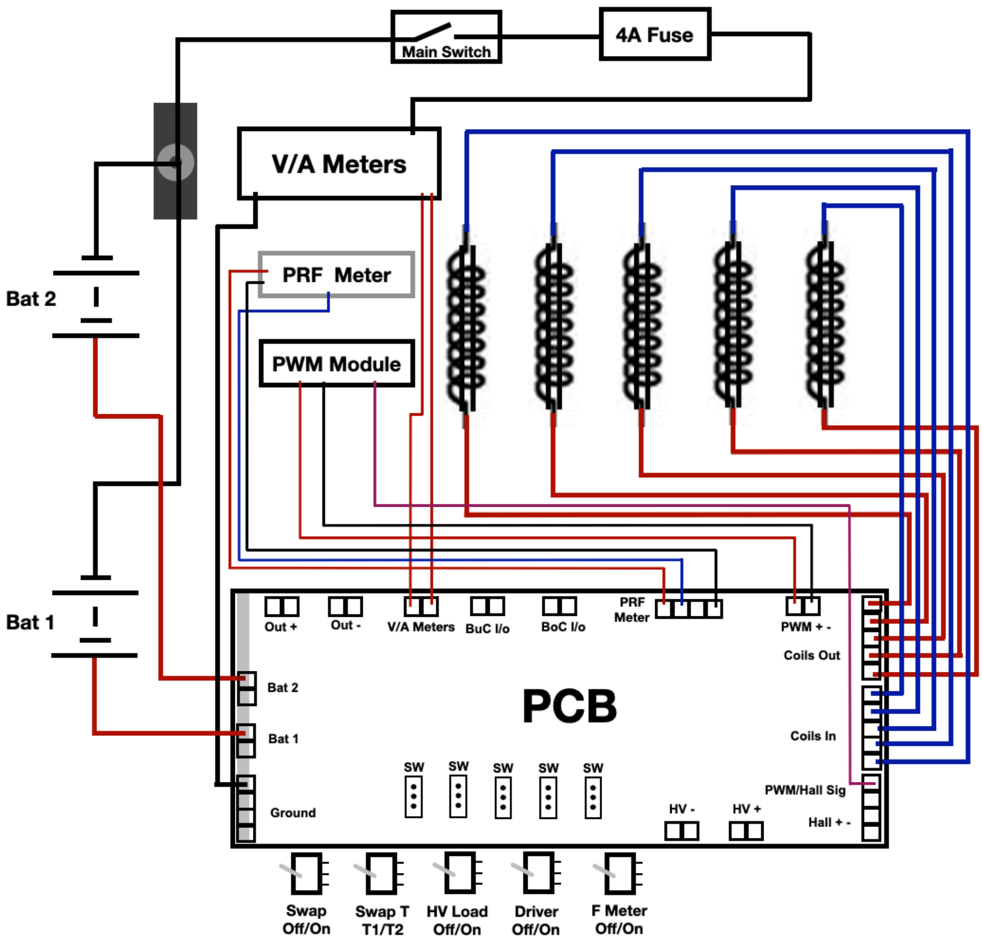

This work is dedicated to those who dare to question that which others assume.

Foreword

What type of situation does it take to cause one to question the status quo, to encourage one to probe and poke the hornets' nest again and again to see how it responds to different stimuli and under different conditions? Such an approach is befitting scientific exploration and discovery since, unless one pushes against the boundaries and questions everything, one can never know the limits of knowledge and within that, your own.

The scientific method has evolved over centuries to encourage the formulation of questions, the testing of them in the crucible of the laboratory and the development of new questions and further hypotheses. Despite this noble venture, whether driven first by the discovery of a phenomenon or the formulation of a hypothesis and the uncovering of experimental answers, it is conducted by human beings with innate perceptual biases and within a wide variety of cultural settings. In many ways it is also a conservative venture in that established and repeated observations, that are enshrined, codified and referred to as 'Laws', posses an inertia that serves to protect them from being too easily overthrown. Like a large cargo ship, it is resistant to minor storms that might otherwise throw it off course and upset its contents. This inertia then serves to give a necessary and important stability to Science but which also renders it somewhat impervious to paradigm shifts and new ideas that threaten to shift the centre of gravity of our lumbering ship.

This is especially true with regard to the topic of novel energy systems. The foundation principles of electromagnetic induction, for example, were laid down nearly two hundred years ago. In some cases the inception of Laws included a set of assumptions and decisions about how and where they were to be applied and, for many generations, they have served our society, our innovators and engineers well. However, the context and the larger paradigm in which we live has developed and moved on; it has expanded beyond recognition in some cases. We have more accurate and precise measurement systems and are more acutely aware of the deeper and finer structure of the world around us such that some of those early assumptions and approximations are no longer valid or have a shorter reach.

With regard to its participants, Science posses a broad brush. They hail from all sorts of backgrounds and innate persuasions, a fact that the scientific method itself must try and normalise and compensate for through the peer review process. Like a growing number of people, I come from a background that takes it as axiomatic that the Universe is much grander than we have yet discovered and that there are future revolutions yet to unfold and be told. That we are immersed in energy of one kind and another is one particular derivative of that view, but that, for good reasons, this hoped for energy is not freely accessible until such time as we demonstrate an ethical mode of conduct that recognises its use and benefit for all and not just the few - to demonstrate power with and not power over.

With current work on fusion power and the urgent drive to combat climate change and reach net zero, it is perhaps inevitable that in the near future Pandora's box will be cracked open and that, besides hope, there will be a surge of new ideas and possibilities, the glinting of a larger reality in the light of our new awareness. The interests of those I benevolently refer to as ‘garage enthusiasts’, speak to that desire for a larger reality to make itself known. Working largely on the ‘fringes’ of mainstream Science, perhaps there is also the belief that the immersive energy spoken of is actually potentially available to some of the more modest of technical systems and individual abilities. Such motivations drive a significant number of ‘seekers’ after their own specific value driven goals. In many respects, I share those goals, albeit with the proviso that, if the data should reveal that my proposed working hypothesis is wrong, then I would accept that and move on to ask other questions of the world.

It is to these aspiring individuals, probing and searching with often very limited resources, as well as to large academic establishments that are often driven by what attracts funding, that I seek to address through this manual. My hope is that, by the end of this document, anyone with some key skills, namely students, post-grads, researchers and their assistants, technicians or otherwise, can demonstrate that the current framework of Physics, embracing principles in the domains of electromagnetism, energy and quantum theory for example, can show hints of evidence for a few gaps in the fabric of our understanding and some unanswered questions. That is not to say that these theoretical pillars of Science are wrong, for that would undermine the work of countless individuals who have tested many of these theories to the highest level possible and they have repeatedly come out triumphant. Rather, by asking slightly different questions, by probing from a different angle and in a different context, some features of Nature may have been overlooked or ignored for perfectly justifiable reasons. We may be surprised at what has been there all along, quietly waiting in the wings for the asking of the appropriate questions and with the right tools.

It goes without saying then that those who will endeavour to replicate the findings that I have presented, both through interim reports and several forthcoming papers for publication, will need a variety of skills. In pitching this document at the right level, I have chosen to assume only a moderate degree of physics and technical background. In fact I have found it very helpful to consider myself back in the classroom with my eighteen year old ‘A’ level Physics students and where I am describing to them how to construct and test such a device. If you did not reach that particular educational level, then my aim is that this project is still open to you and, if you moved beyond that, either formally or through your own efforts and studies in later life, then I simply ask for your understanding of where I am coming from.

It is appropriate to mention here that there is no commercial motivation behind this manual or the research project. It has been entirely self-funded and, over the last three years, I could easily have spent as much time in a quaint Cornish pub looking at the plimsole line on a pint glass while thinking about my ideas - it’s just a matter of choice. As such all the information I have assembled here is offered freely on the basis of open knowledge and scientific enquiry. It is for all those who chose to work with it and find some value therein. Equally, if anyone believes they can produce something practical and useful based on these findings, as indeed I will likely try myself at some point, then that is all well and good and I wish them every success.

The ideas presented here will, in due course, be superseded by others who will bring new ideas to the fore. They will break through the ‘glass ceiling’ of awareness prescribed by this particular work and discover how to take it forward in some new and exciting directions. This is the way of discovery and of evolution - the bringing of fresh questions and insights into the arena so that the ‘game’ is raised to the next level. It is right and natural that the baton be passed on and that each step on the journey only goes so far.

As a physicist, rather than someone trained in electronics or electrical engineering, I seek then to act as a bridge between the scientific establishment, in which I did my training and had my career, and the ‘open hearted’ seeker after a larger point of view - a more embracing reality. The skills I have developed on this ‘curiosity driven’ project have been on a ‘need to have’ basis, and I am learning new things every day. But even without those particular skills, so long as the reader can apply diligence and focus, pay attention to the small details, so that the larger ones can unfold naturally, then this build and its testing should be straightforward and rewarding.

I would even go so far as to say that, if you have a good degree of manual dexterity, a healthy measure of dogged determination, can think reasonably logically and yet are open to the unexpected and your own deeper instincts, then you should do very well. No more is required of you than that, besides being open to the possibilities of what Nature chooses to reveal to you, based on how and in what context you ask your questions to tease out Her truths.

I was reminded recently of a quote by John Archibald Wheeler, the American physicist well known for his work on Quantum theory, Relativity and Unified Field Theory amongst others. He said: "In any field, find the strange thing and explore it." Well, this is indeed a strange thing and it most definitely needs exploring.

Enjoy the journey of discovery!

Kerrow Energetics,
Nov 2022

ACKNOWLEDGMENTS

Over the years working on this project, various people have inspired me, helped to steer my efforts and compensate for my lack of knowledge and experience in many areas of electronics. In particular I am very grateful for the developer who, for reasons explained in the Introduction, I will simply refer to as Ray. Also for the suggestions and contributions from the late Patrick Kelly and the participants of various electronics forums, including Eric Gibbs, Dick Cappels, Albert Hall, 'Alec-t', 'HarryA' and Brian Kelly.

I would also like to acknowledge the published work of Dr Peter Lindemann and Aaron Murakami from the Energy Science forum, who laid out the key principles behind the work of the late John Bedini. Also to Gary Hammond for his suggestions with regard to some of the circuit developments and for his continued interest and enthusiasm in the work. Also, a thank you to Lee Ferrand for his suggestions after reading through a first draft of this manual.

Lastly to my wife Rhona for her tolerance at my spending endless hours focused on this project and for the supportive ambience.

CONTENTS

INTRODUCTION

This manual has been written for the express purpose of supporting anyone with moderate technical skills in replicating the results of experiments that were conducted during 2022 on a device referred to as a ‘Pulsed Flyback Generator’. These were undertaken to test a hypothesis for the purpose of confirming or refuting prior claims of energy gains by applying high voltage inductively generated transients to chemical batteries.

Such claims have been made since the 60s by independent researchers, in particular the late John Bedini, and replicated many times by followers and enthusiasts, but with varying degrees of successes and clarity. Despite this, very little data and coordinated results have been forthcoming to the scientific community for evaluation, perhaps for fear of ridicule, with the result that the phenomenon has languished on the fringes of study and not endeared itself to widespread investigation and replication.

As a retired medical physicist (radiation) and science teacher, but with an embracing cosmological paradigm, I straddle both viewpoints and was fortunate to be in a position to be able to invest time and effort into addressing this question in a manner that would stand up to scientific scrutiny. As such, I am as curious as the next man to find out if the long standing claims have any merit or if the results were such as to be buried well within the uncertainty range so as to render them of little value and import or in the end to arise from some more mundane mechanism.

An initial build in 2018 was in part inspired by the work of a contact and developer abroad who, for good reasons, chooses to remain anonymous due to the nature of the political regime where he lives. With his national energy infrastructure being very poor and fragmented, he is subjected to daily power cuts of varying duration requiring some form of backup system. He developed his own take on a ‘Bedini’ and ‘Adams’ style generator in order to convince himself that not only was it possible to extract useful energy from the environment, but that the current theories of electromagnetism and thermodynamics are not fully complete.

While coming from a more technical than scientific background, he succeeded in producing several small rotor based generators that, as far as he is concerned, demonstrated an energy harvesting phenomenon. In more recent times he has further developed a pair of solid state ‘outage’ lights that he uses during his protracted power failures and which have run for over two years without the batteries ever running down. This is sufficient for his own needs.

After long discussions with him, and also the late Patrick Kelly, who in this area of study was considered by many to act as a ‘clearing house’ for this type of technology, I decided to throw my hat into the ring and try and replicate his findings. In so doing, I would apply sufficient scientific rigour and repeatability for any results to stand up to analysis and replication, which ever way they fell.

With the first build in 2018, I developed a device with an operational rotor switching system and that produced pulses with a modest voltage of around 500V. I could see that while one battery was supplying the circuit, the other was receiving the pulses causing its voltage to rise via an only moderately accurate panel meter. However, my understanding of the system was undeveloped at that time and I was not aware of the many factors that contributed to the overall behaviour and performance of the whole system.

In time, I came to the view that the fixed frequency produced by the rotary switching system was a limitation and that to see more positive results I needed to produce higher and adjustable frequencies. So I went on to construct a purely solid state device which gave me free reign on the pulse repetition frequency (PRF) with the expected potential for a noticeable power gain. This unit offered me PRFs up to 15kHz and I applied them to a set of seven solenoids in the hope that this would result in some clear evidence of energy gain. My naive optimism was again uninformed about the many factors and the optimisation of the parameters that are of importance and based on the assumption that pulse frequency was the most important factor for the phenomenon.

Seeing no good evidence of such, I put the work aside and started to investigate electrolysis efficiency using resonant circuits. During this research I developed the skills for designing and building PCBs and which would find suitable application later.

The electrolysis research, with its explosive results, sometimes of the unhealthy kind, did not produce results better than was typical for various commonly used types of electrolysis systems. In the end I decided to return to my first rotor-based generator and see if I could look into it in more depth and also step outside the various parameters based on the commercially available documents that I had acquired for the project. For this I would develop the circuits, using PCBs instead of the old veroboard and connector block method, and with many more features that would allow me to investigate the hypothesis of energy harvesting in greater detail and with more control over the various parameters.

Starting in late 2021, I rebuilt the generator with a PCB and then further developed a ‘low sided’ capacitive discharge circuit design as a ‘high sided’ version to suit my own already existing generator design. I also incorporated features to allow for both a rotor based switching and ‘555’ chip trigger input, as well as using a pulse width modulation (PWM) unit for even more precise control.

Over time improvements were made to the PCB design resulting in a version 2 and which has been used for the bulk of the testing to date. A v3 board was built and made ready for installation but this has been superseded by a v4 board due to requests to make available a circuit that allowed others to replicate the findings with minimal effort.

Such materials and accompanying guidance are a required part of the planned scientific paper, that will be written during 2023, and which would provide enough information to enable others to replicate the findings. To this end I decided to design a new updated PCB together with associated information, so that anyone new to this topic, and without specific training in electrical engineering, could construct the device. This PCB would operate with only those elements of the device that have been found to work or be useful in a net energy gain. Other elements, such as the rotor system and the capacitive discharge circuit, were excluded to limit complexity and cost.

In moving away to some degree from the prescriptions of some of the commercially available material, I endeavoured to start with a clean slate on what might work rather than be led by prior expectations. This approach has paid off in that some ideas that were considered important for this phenomenon to be observed have turned out to produce either very modest or minimal results, whereas other factors that have never been publicly shared, have proved to be crucial for optimum results. So being guided more by my own logic and instincts, has paved the way for a project that has not only confirmed the presence of an energy gain phenomenon but has far exceeded my original expectations.

It is important to say that I have not invented anything new here. Any attributions of invention must be credited to the likes of Nikola Tesla and John Bedini amongst others. If I have done anything 'novel' it is to circumvent certain assumptions and boundaries regarding the operational parameters, as set out by others in the field. For example, it was a given that the limiting factor for the peak voltage of the pulses was the coils such that bigger and beefier coils would give better results. In fact, the limiting factor is not the coils but the 'avalanche rating', or breakdown value, of the active device used. With even quite modest sized coils, the way to improve the peak voltage is to use active devices with higher reverse polarity ratings. Similarly, the position on the charging profile where the pulses are applied, has a very significant effect on how the battery responds, just as optimising the pulse frequency to a specific battery configuration is crucial.

It is equally relevant to mention here that there is nothing special or unusual about the workings of the circuit described in this manual. Its behaviour is as defined by electronics and electromagnetic theory with regard to the production of flyback pulses from the solenoids. The energetic processes involved in the fast switching of the currents in the coils, with the consequent rise and fall of the associated magnetic fields, is all well understood. This fact alone should make diagnostics relatively straightforward should any faults develop. However, the device has been developed in such a way so as to make it straightforward to optimise the pulses for the particular battery configuration being used. It is perhaps the customisability of the stimuli to the specifics of the batteries and the electrochemistry, of the whole system, that is its greatest strength.

Any 'mystery' involved is focused squarely on where the pulses meet the electrochemistry at the battery's positive electrode. This is where something strange and unexpected is going on and which might be an example of an energetic process that baulks at conventional wisdom and understanding, especially with regard to the Second Law of Thermodynamics. As has been experimentally shown and reported on many occasions over the last decade in particular, there are examples of both organic processes and electrical phenomena that, while strictly adhering to energy conservation, nevertheless break free from the confines of some versions of the 2nd Law with the prospect of the emergence of new technologies as we collectively face new challenges. Discovering the likely source of the energy will take further detailed and elaborate tests.

At the time of writing, the highest value of Coefficient of Performance (CoP) measured is 37.8 ± 1.33 and with a theoretical power output of 219 W. However, until the live power tests are done, there is no evidence regarding how much of that power is actually usefully available and how much is due to one or more unrecognised artefacts. Equally, there may be hidden factors, physical or computational, that are not yet taken account of and that are contributing to the results. Only when the power tests are done, and when the power that can be extracted over a number of cycles and without the batteries running down over an extended period of time, is measured, will there be enough concrete data to confirm an effect. This is the inductive scientific method at work.

Even if the results were to show a viable output of only a few tens of watts, while that is not very useful as a power source, scientifically speaking it could be highly significant since there should not be any at all. Indeed, there should be a loss of energy due to the normal thermal and other losses appropriate to a regular switched inductive circuit. That fact that this appears not to be the case implies that an undefined process is at work, one that will require further experimental design and testing to clarify. This will include seeing if some process within the electrolyte itself is causing the results.

This particular research project then is still not complete in that the CoP results stage is to be followed by power tests during early 2023 and which will serve to integrate various uncertainties and confirm the theoretical power outputs derived from the CoP tests. Only then will it be clear to what extent there is a viable phenomenon of energy gain worthy of wider reporting.

This manual has been prepared in a logical sequence and does not need to be read from cover to cover before construction can begin. However, for some, reading the whole manual before starting is the preferred choice so as to get the bigger picture and to feel fully prepared for a build. Whatever your preference, if you undertake all the stages in the sequence they are presented and discussed, then you should ‘arrive’ at a working unit. Additionally, there will be some who already have a well developed device in operation and who, instead of starting again, can adjust or revise certain aspects of their system to achieve better results. For them that would be a constructive use of their time and energy.

It is my hope then that, as and when experimentalists, investigators and explorers find confirmatory evidence, or develop new insights and experiments, then this is shared with the wider community for the benefit of all using recognised routes. Equally, if there are aspects of this manual that require further elaboration or improvement, or you just wish to comment upon or share your results, then I welcome sensible feedback and there is an address at the end where I can be reached.

All the relevant files and appendices are available on either of the following links:

https://mega.nz/folder/YUM0nLoT#bYpLIazqMM5K2IrEQjghDQ

https://www.dropbox.com/sh/td55b8675vvqtbg/AADzPSKMOI8q_YM1cFUT2T07a?dl=0

PRINCIPLES OF OPERATION

The following description explains the operation of the device and which is based upon standard electrical theory and principles.

With reference to Fig 1, a supply battery, or for testing purposes a power supply, provides power to the circuit and the five coils (solenoids). These are energised when the rising edge of the square wave from the pulse width modulation (PWM) unit oscillator arrives at the Gate of the main active device which is either a MOSFET (Metal Oxide Semiconductor Field-effect Transistor) or an IGBT (Insulated Gate Bipolar Transistor).

While the square wave input is high, the magnetic fields in the coils build up to a maximum and, in keeping with Lenz's Law, result in a high voltage (HV) pulse being generated and which is then earthed through a diode to ground.

On the falling edge of the square wave, the coils switch off and again generate a reverse polarity 'flyback' pulse as they oppose the collapse of the magnetic fields. This pulse, in the range of 0.5 - 1.7kV depending on the active device used, and with a pulse width of 40-50μs, appears at the Drain of the active device and is directed to the receiving battery through a set of diodes. The rate of change of the voltage dV/dt approximately equals 1.5E+08 V/s (1.5 x 10^8 V/s) and which may be a relevant factor for the degree of the energy gain observed.

Alternatively, as done in earlier experiments, the pulses can be directed to a bank of storage capacitors whereupon, at a set voltage, the capacitors will discharge into the receiving battery as a high current, high intensity pulse. However, experiments using this approach showed much less effect than using the HV (high voltage) pulses directly and so this is the approach used in this system for replication and which also reduces the system complexity.

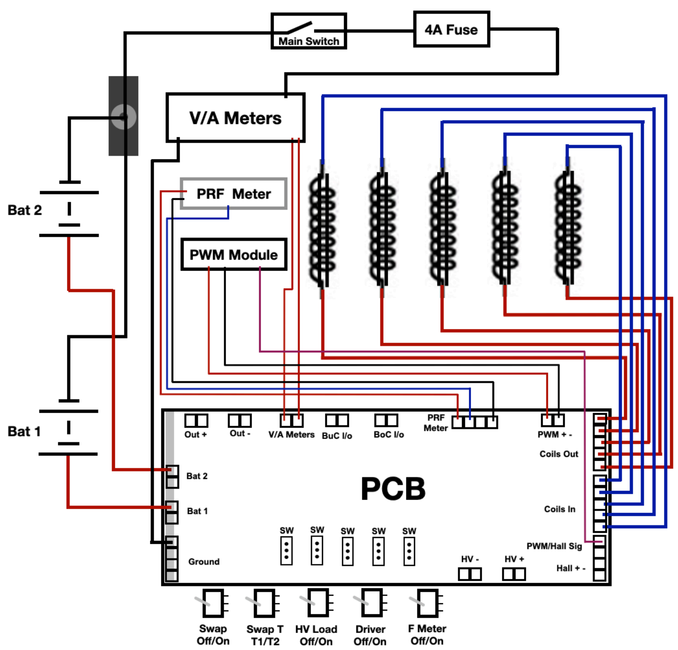

Fig 1: Generator Circuit

The peak flyback voltage has been measured using a custom built 10:1 potential divider and calibrated using transients generated by a signal generator. The peak pulse voltage has also been found to be limited by the 'Avalanche breakdown' rating of the active device and so the indigenous peak voltage of the coil is even higher, an issue discussed later at length. Changing the active device for one with a higher ‘avalanche’ rating is the means by which higher peak voltages are obtained. The effect of this change has been found to depend on the chemistry type of battery being charged and is more pronounced with certain parameters than others. Further clarity will come from the forthcoming power tests due to start in early 2023.

As an alternative to using the PWM (Pulse Width Modulation) module for the trigger, switching the active device can be done using a rotor driven by the energised coils in conjunction with a Hall sensor. This results in a fixed pulse repetition frequency, at max RPM (revolutions per minute) as each of the five sets of rotor magnets switches the solenoids on and off every rotation of the rotor. As such, at 3,000 RPM, the PRF will be 250Hz (3,000/60 x 5). Again, since the CoP values measured using the rotor and Hall sensor switching were significantly lower than with the PWM, this less flexible arrangement is not a part of the replication build and again considerably reduces the complexity and cost.

Using the battery swapper circuit, with a predetermined swap interval, the batteries are automatically switched over so that the energy that has been expended to the circuit by the ‘run’ (supply) battery can be replenished when it becomes the ‘receiving’ battery. If a CoP> 1 is measured then the ‘run’ battery can also deliver some useful power to an external load while the device is operating with an internal efficiency that has been measured to be in the range of 40 - 50%.

For the purposes of testing, the swapper is switched off so that the effect of variations in individual battery properties is negated and a stable power supply is used in place of the ‘supply’ battery. This also has the advantage that only one battery needs to be fully recharged for the start of another test run. The use of a power supply also facilitates the adjustment of the coil voltage in test runs, since the voltage would remain stable over the test period and not drop under continuing load as would occur when using a regular battery. Various methods by which the coil voltage can be adjusted, separately from the supply to the circuit, are explained in its own section entitled ‘Coil Voltage’.

When it comes to live power tests, then the swapper needs to be enabled so that the energy expended in one swap interval can be replenished in the next, as described in later sections. The use of the battery swapping system then is the basis of the normal operation of the system.

THE PCB

The revised PCB v4 is available for anyone interested in replication using the Gerber zip file loaded into the 'PCB Files' folder on the file storage links at the end of the Introduction.

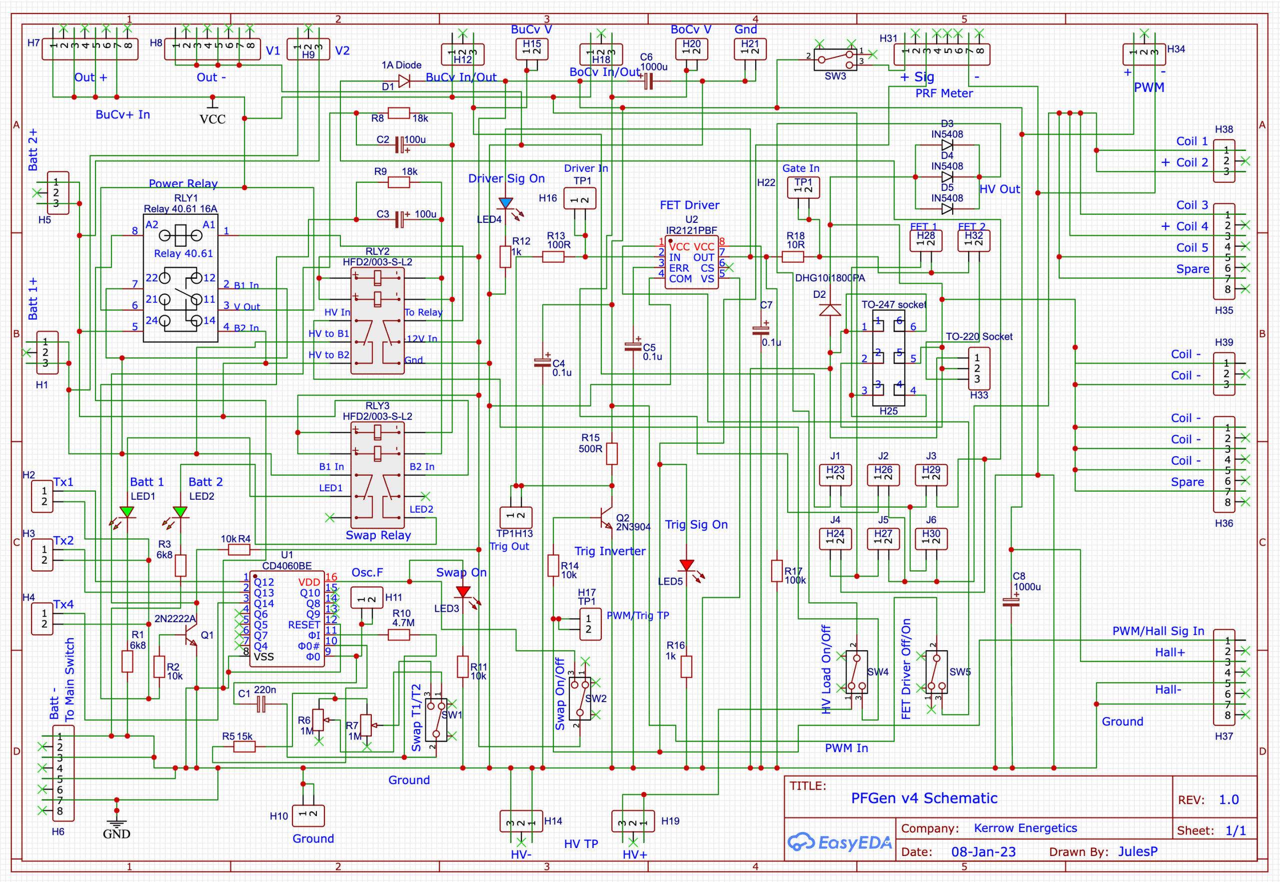

Fig 2: PCB Schematic

Following on from extensive research and testing of the pulsed flyback generator, the findings gave rise to the need to update the circuit PCB in order to reflect them. To this end the PCB that I have used for the majority of my testing (v2) has been updated to v4 for use by anyone wanting to replicate the findings.

The 'PFGen v4' PCB will not have connections or components related to the use of a rotor based switching system or a capacitive discharge unit (Cap Dump Circuit) that delivered high current pulses to the batteries. Instead it focuses on what has been shown to give the best results and will provide the following features:

- Battery swapper circuit
- Trigger circuit
- Drive circuit (FET/IGBT based)
- Coil (solenoid) & circuit supply options
- Coil connections

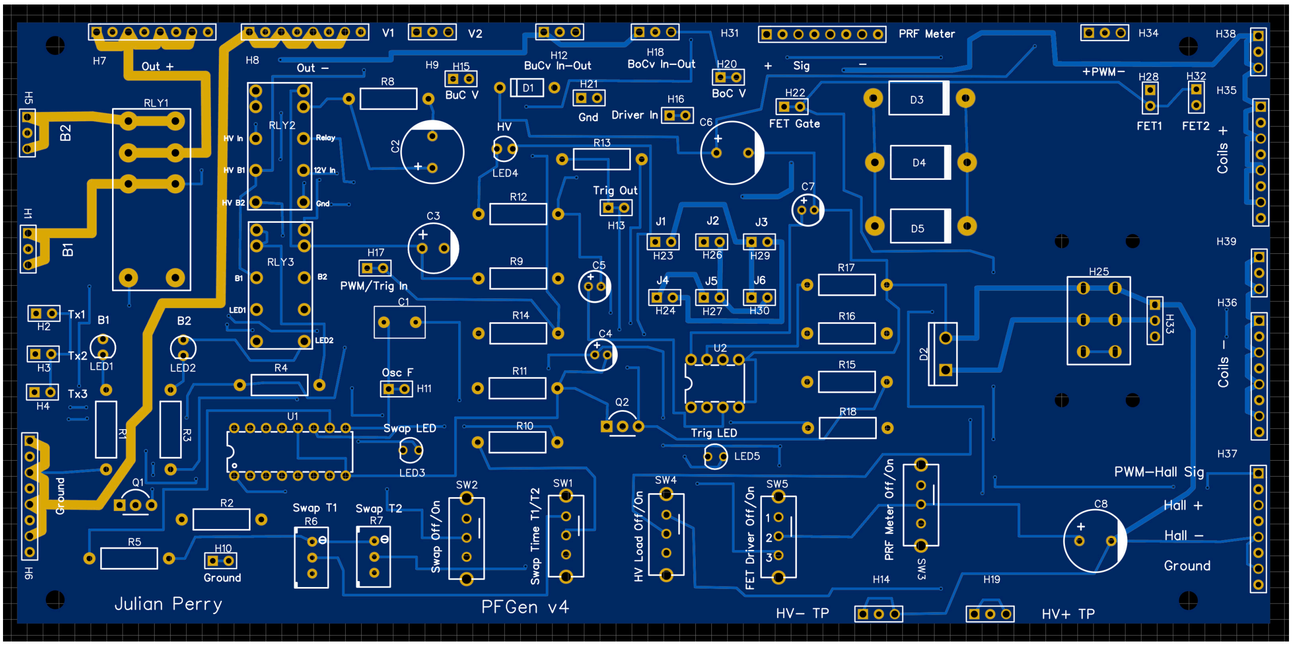

Fig 3: PCB - 2D view

- External power outputs
- Connections to Buck/Boost converters (if required)
- Connection to PWM module (or Hall sensor if required)
- Connection to external PRF meter
- Connections for measuring the HV pulses (via a potential divider)

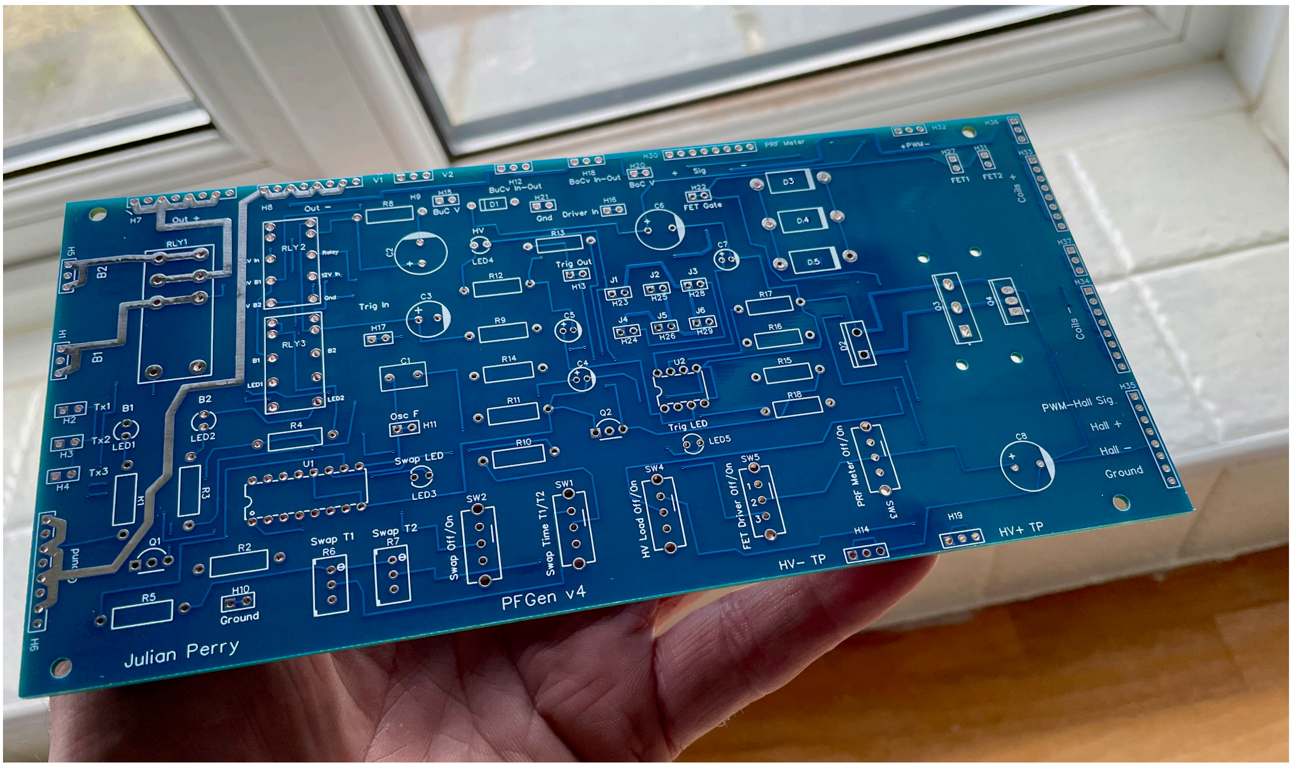

Fig 4: v4 PCB (first design)

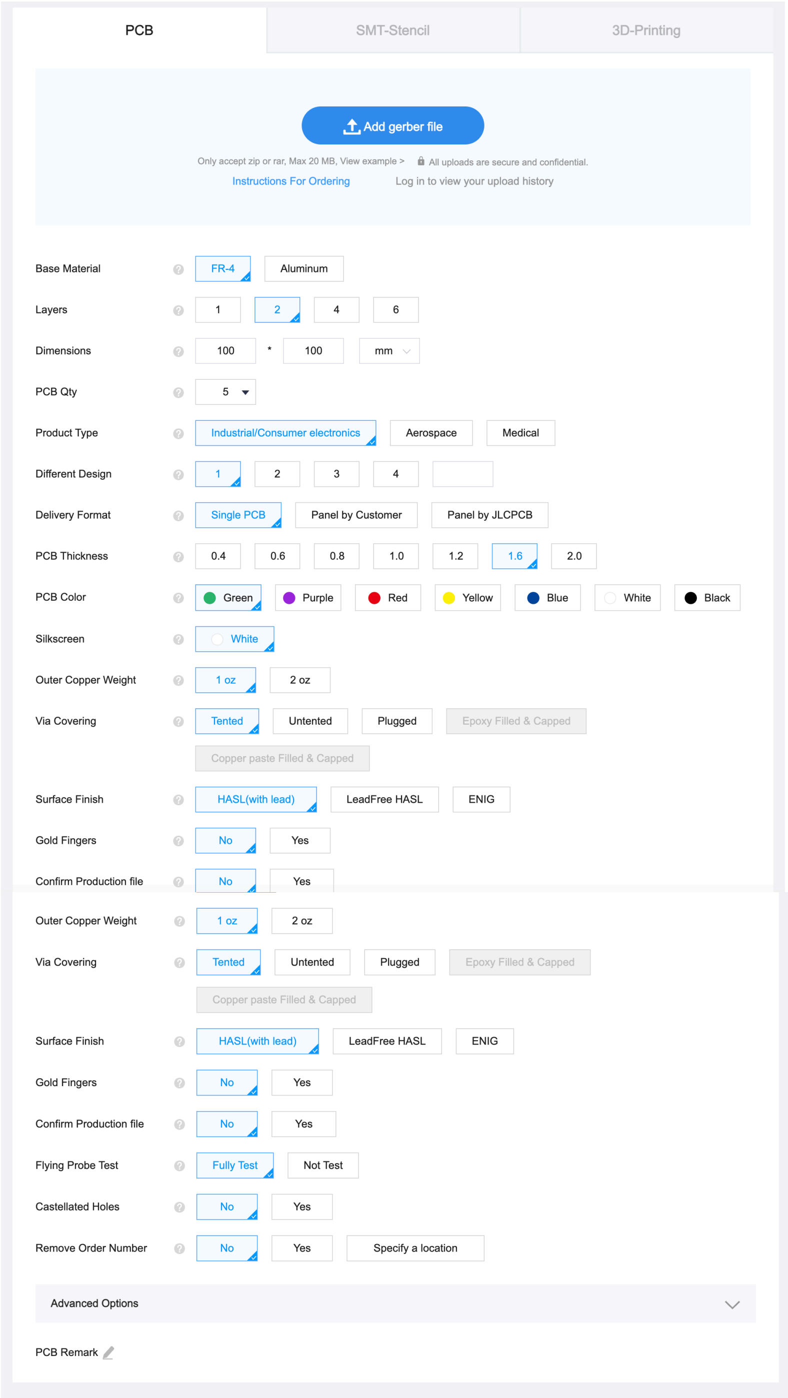

Fig 5: PCB Ordering

Figs 2 and 3 show the circuit schematic and the 2D view of the PCB and Fig 4 the actual printed board.

In the 'PCB files' folder on the above links is the Gerber zip file as well as the above mentioned files. I have not included any editable files in this folder as the focus is on replication rather than going off at a tangent. However, it is recognised that if somebody so chose they could copy the schematic, make modifications and try other things. Just so long as such departures are mentioned along with any results then there should be no confusion regarding the replication process.

The Gerber files are not specific to any production site but can be used with many. I used EasyEDA to design the board and JLCPCB (https://jlcpcb.com) to print it so I will describe this route.

The Gerber file should be downloaded to a suitable personal location and then imported to the 'instant quote' option on the JLCPCB website at https://jlcpcb.com. Click on the 'Instant Quote' button and upload the Gerber file as a zip file - do not unpack it first.

Once imported you can select a range of variables, although the defaults are usually acceptable. The minimum number of boards is five and the few extras are useful to have around to share, or practice certain tasks such as track thickening with solder. I personally prefer the blue coloured boards rather than the default green.

The cost of producing five boards is relatively cheap, at less than £7 depending on offers they may have on and other discounts, but this is more than matched by the shipping costs shown during the checkout stage. The cheapest method, Parcel International Direct, can still arrive within two weeks.

What you should see when your boards finally arrive is shown in Fig 4 (this one is the result before the sockets were added for Q3 & Q4) and where the thicker tracks on the left hand side are those for which the solder mask has been removed so that you can add solder to thicken the tracks, as explained later. This will allow a higher current rating than the track width would normally accommodate.

COMPONENTS

The components that make up this board are shown in Table 1 and several links are included to help source some items.

Component ID	Details	Component ID	Details
R1	6k8	D3	UF5408
R2	10k	D4	UF5408
R3	6k8	D5	UF5408
R4	10k	RLY1	Relay
R5	15k	RLY2	Relay
R6	1M	RLY3	Relay
R7	1M	U1	CD4060BE
R8	18k	U2	IR2121PBF
R9	18k	Q1	P2N2222A
R10	4.7M	Q2	2N3904
R11	10k	TO-247 Mount (2x5P)	https://www.aliexpress.com/item/1005001884522039.html?
R12	1k	TO-220 Mount (1x3)	https://www.ebay.co.uk/itm/121267837242
R13	100R	Active Device 1	IRF840
R14	10k	Active Device 2	STP20N95K5
R15	510R	Active Device 3	STP12N120K5
R16	1k	Active Device 4	STW12N150K5
R17	100k	Active Device 5	STW12N170K5
R18	10R	Heat sink	SK 129 38,1 STS
C1	220n	LED 1 - G	
C2	100u+	LED 2 - G	
C3	100u+	LED 3 - R	
C4	100n+	LED 4 - B	
C5	100n+	LED 5 - R	
C6	1000u+	SW1-SW5	SPDT
C7	100n+	H (1x4) 7	https://www.ebay.co.uk/itm/234615666086?
C8	1000u+	H (1x2) 10	https://www.ebay.co.uk/itm/234615666086?
D1	IN4007	TPs (1x2) 20	https://www.ebay.co.uk/itm/265539912049?
D2	DHG10i1800PA	TP Jumpers (x5)	To go over TPs
		PCB Feet (x4)	
Notes: *Resistors are all 0.25W. For capacitors, + means electrolytic.*			

Table 1: PCB Components

Terminal Ref	Pins	Function
Going clockwise around the board from the bottom left corner		
H6	**H6-1 to 4**	**Ground connection wire (to meters, fuse and main switch)**
H1	**H1-1/2**	**To battery 1**
H5	**H5-1/2**	**To battery 2**
H7	**H7-1/2/3/4**	**External load +**
H8	**H8-1/2/3/4**	**External load -**
H9	**H9-1**	**Batt1 V to small panel meter**
H9	**H9-2**	**Batt2 V to small panel meter**
H12	**H12-1**	**Out to Buck converter (optional)**
H12	**H12-2**	**Return from Buck converter (optional)**
H18	**H18-1**	**Out to Boost converter (optional)**
H18	**H18-2**	**Return from Boost converter (optional)**
H31	**H31-1**	**Supply + to Frequency meter**
H31	**H31-2**	**Signal to Frequency meter**
H31	**H31-4**	**Supply - Frequency meter**
H34	**H34-1**	**Supply + to PWM module**
H34	**H34-2**	**Supply - to PWM module**
H38	**H38-1 to 2**	**To coils + out**
H35	**H35-1 to 3**	**To coils + out (H35-4 spare)**
H39	**H39-1 to 2**	**Return from coils**
H36	**H36-1 to 3**	**Return from coils (H36-4 spare)**
H37	**H37-1**	**PWM (or Hall signal in if using rotor)**
H37	**H37-2**	**Hall sensor + supply**
H37	**H37-3**	**Hall sensor - supply**
H37	**H37-4**	**Ground connection**

Table 2: Terminal Connectors & Function

For many components, such as resistors and capacitors, then any reasonable source is acceptable. When it comes to active devices (MOSFETs and IGBTs), in particular the harder to come by ones since the pandemic, then some caution is advised since there is a lucrative market for suboptimal components that some would reasonably refer to them as ‘dodgy'. These may still work but will often fall outside the specifications limits or be unstable with respect to temperature. Most of those who build circuits regularly will have stories about components available cheaply but which did not live up to their specification. Use you common sense but eBay or AliExpress has proved to be a good source for many of the components and so far I have only encountered one ‘dodgy’ comparator used in the cap dump circuit build, but which is not a part of the revised circuit build.

Other good sources are Farnell, RS online, Digikey, Mouser etc and, depending on your location, some will offer more reasonable delivery costs than others. For the IGBTs, I used a company called Jinftry (jinftry.com) to put in a request for sources of those items that were hard to find. They act as a distribution and advice centre for components and will seek out sources for you from multiple suppliers and manufacturers. They usually quote FedEx delivery charges that will cost more than the components so ask about low cost postal delivery. This can be as low as £5 and will still arrive within 2-3 weeks; if the wind is in your favour.

Via them I was contacted by someone whose details are below who stated that they had the STW150k5 and 170k5 devices that have since proved effective. With the provision for device sockets, you can easily plug in and out any of these devices and select which one is active via jumpers. This is explained in depth in the section on ‘Active Device Selection’.

If you are struggling to obtain some of the more ‘avalanche rugged’ devices and have to use just one MOSFET, then use the STP20N95k5 which was used for 95% of the testing and yielded very good results. It is simply that the higher voltages available with the other devices can push those CoP values even higher but, so long as you are above about 800-900V, the phenomenon should be readily apparent and measurable.

The component list in Table 1 is further broken down into terminals and connectors in an Table 2 and includes the function of the connection.

When referring to a connector, whether it be one of the type positioned around the edge of the board or the two pin connectors serving as jumpers or test points (TPs), I have used the following labelling convention as indicated in the Fig 6 which shows the schematic symbol, the board net and 2D views, the actual component and it inserted onto the PCB.

For example, with connector H35 shown in Fig 6, and arranged vertically on the top right hand edge of the board (Coils +), the schematic symbol shows 8 holes and are referred to as H35-h1 etc. to H35-h8 from top to bottom and where hole 1 is surrounded by a square area.

The actual component consists of 2 x 2 pin terminal blocks dovetailed together, and therefore has 4 pins referred to as H35-1, H35-2, H35-3 and H35-4. The pins go into alternative holes on the PCB starting with H35-h1 then H35-h3, h5 and h7. In this case the fourth terminal, H35-4, is a spare connection should anyone decide they want to to use 6 coils rather than the 5 that I have used.

This naming convention applies equally to terminals arranged horizontally along the sides of the board. For example, with H7 (Out +), the 4pin terminal block will be placed in holes H7-h1, h3, h5 & h7 and where h-1 is on the left indicated why a square area around the hole.

If you look closely at the Net or 2D view in Fig 6, you might be able to make out that the holes H7-h1, h3, h5 & h7 are of slightly large diameter at 1.1mm compared to the other holes at 0.89mm. As such the terminal block will easily go into the correct holes but be a tight fit, or be hard to insert at all if shifted over one hole to h2, h4, h6 & h8.

In the tables the stated function will aid in understanding the device better when it comes

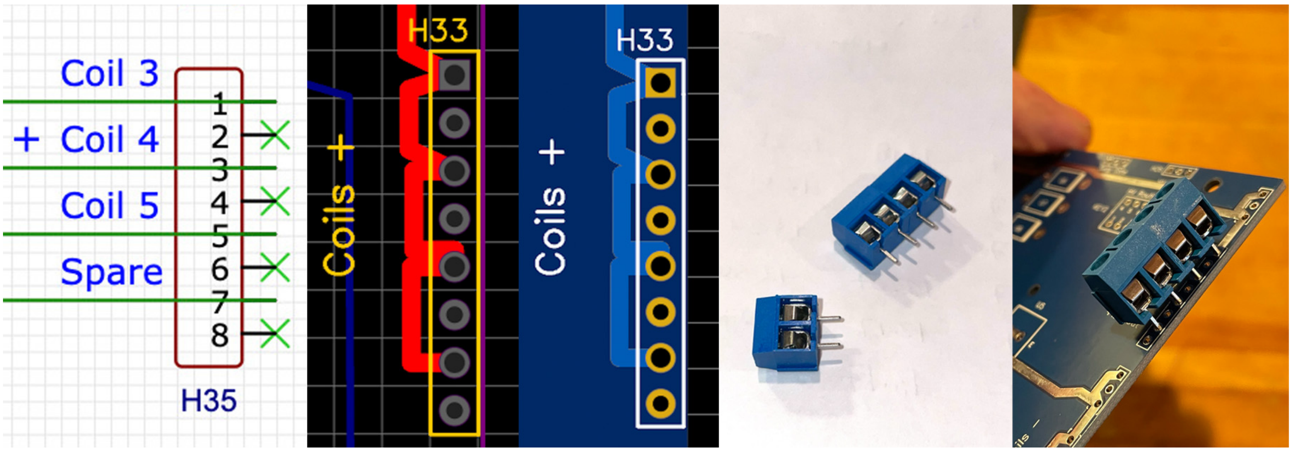

Fig 6: Terminal formats

to diagnostics, and process that inevitably arises when, for example at switch on, nothing happens and you are left wondering where to start with tracing down the fault.

In my experience, one of the most valuable tools, besides heat sink clips, is the continuity meter allowing you to find breaks in a circuit where there should not be any. This might be due to a dry joint or a poor quality component. As has happened to me on several occasions, a wire was inserted into the wrong terminal which is easily done when a terminal blocks carries both the + and - feed to a device, such as with H34 for the PWM module.

Devan/ Sales Manager
Company name: JINGHONGYANG ELECTRONICS TECHNOLOGY LIMITED
Contact Person:Devan Chen
Email:Devan@jhykjic.com
Phone:+86 13430828789
Instant Messaging: (Skype) 134308282789
WhatAPP: +86134308282789
website: www.jhykjic.com

I addition to the 'Terminals List' in Table 2, there are two more lists consisting of the 2 pin terminals that are used as on board switches and those used as test points (TPs) for the purpose of circuit diagnostics. These are shown in Tables 3 and 4. These will make it easier to check that you have wires going to the right places and to provide a rationale for the use of the terminal, on board switches and test points (TPs).

Terminal Ref	Function
Going from left to right across the board	
H2	**To set the 4060 chip output to a value of 4096 x the Oscillator frequency**
H3	**To set the 4060 chip output to a value of 8,192 x the Oscillator frequency**
H4	**To set the 4060 chip output to a value of 16,384 x the Oscillator frequency**
H23	**Used in conjunction with five other jumpers to set the coil supply**
H24	**Used in conjunction with five other jumpers to set the coil supply**
H26	**Used in conjunction with five other jumpers to set the coil supply**
H27	**Used in conjunction with five other jumpers to set the coil supply**
H29	**Used in conjunction with five other jumpers to set the coil supply**
H30	**Used in conjunction with five other jumpers to set the coil supply**
H28	**To enable FET1 to be active with a signal to its Gate**
H32	**To enable FET2 to be active with a signal to its Gate**

Table 3: List of 2 pin terminals used as switches

TP Ref	Function
Going clockwise around the board from left to right.	
H10	**A ground connection, useful for scope measurements**
H11	**For measuring the 4060 chip oscillator frequency, the basis for setting the swap timer**
H13	**To check the output of the output transistor in response to the trigger square wave input**
H14	**The negative/ground connection for observing the HV pulses using a potential divider**
H15	**To enable measurement of the Buck converter output, if used to modify the coil voltage**
H16	**To check the input to the FET driver chip**
H17	**To check the input to the transistor in response to the trigger square wave input**
H19	**The positive connection for observing the HV pulses using a potential divider**
H20	**To enable measurement of the Boost converter output, if used to modify the coil voltage**
H21	**A ground connection, useful for measurements of H15 and H20**
H22	**The check point for the output of the driver and input to the active device Gate (selectable via H28/H32)**

Table 4: List of 2 Pin Test points

ASSEMBLY SEQUENCE

The sequence in building a PCB is very much a matter of personal preference but for what it's worth I will elaborate on the sequence I used.

After securing the four feet to the PCB, to keep the underside from abrading on a work surface, the first step is to thicken some of the solder tracks. These show as yellow on the 2D view and silver with no resin coating on the PCB. The aim here is to allow up to 16A to flow from the battery without any damage to the PCB track. As this is only likely to happen when a substantial external load is attached to the circuit, this thickening is only required on the area around the swapper, in particular the track from the batteries to the larger relay and then on to the outputs at H7 and H8.

The maximum current demand for the coils themselves is going to be less than 3A and so can be accommodated by the 0.8mm tracks to the coils and voltage selection jumpers. Creating resin covered tracks wide enough to handle 16A would have impacted adjacent components and so track thickening with solder is the best alternative and easy to do.

Move the soldering tip slowly backwards along the track while holding the solder wire next to the tip and a thickening of between 0.5 - 0.8mm will naturally occur. It can be a good idea to use one of the five PCBs received to practice this before adding other components. Don't add too much thickness on to the terminal block areas but enough to allow it to reflow when the terminals are added and so bond with the terminal pins.

After track thickening, start by adding the terminal blocks around all the edges consisting of 10 x 2pin terminal connectors and 7 x 4pin connectors and where the 4pin blocks are made from sliding together two of the 2pin ones using the fine grooves on their sides.

This is followed by installing all the resistors and the capacitors. Most of the resistors are inline TTH (through the hole) components but the two 1M trimmers, used for the swapper, are three pin.

The installation of the right components in the right place is facilitated by printing off the components sheet (see the file in the Appendices folder) and then ticking off each one as it is installed. A mistake in putting the wrong value into place could mean a whole section of the device will not work, and most likely the whole unit, and so any method to ensure that the right component value goes in the right place is not to be underestimated.

After the capacitors, add the dip sockets to take the 4040 decade counter (U1-16pin), the FET driver (U2-8 pin) and the two smaller relays (RLY2 and RLY3). In the case of the two smaller relays, some of the pins will need to be trimmed from the underside of the dip socket so that the socket can be fitted to the board which has only 10 holes available for each socket.

To do this the following legs on each of the two 16pin dip sockets need to be cut where they meet the plastic body before laying them on the board and soldered as shown in Fig 8. The pins to be removed are 3, 5, 7, 10,12 & 14 where pin 1 is closest to the ‘U’ shaped notch on the upper short side of the socket.

It doesn’t matter if you take pin 1 to be to the left or right of the notch since it will still result in the correct pins being removed.

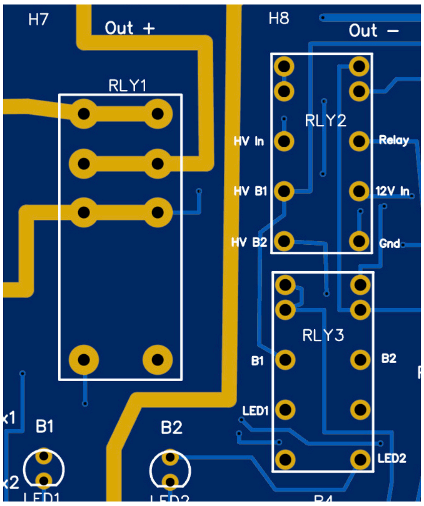

Fig 7: Relay location

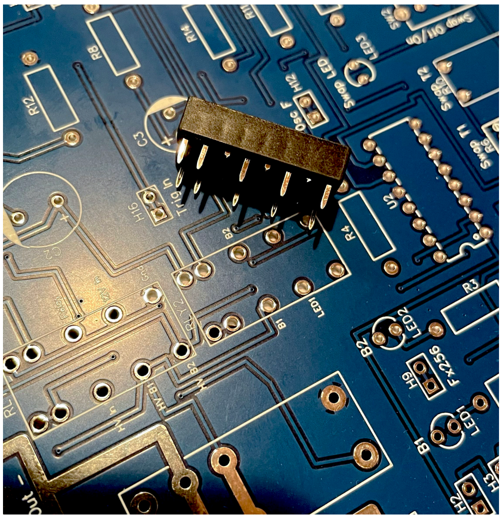

Fig 8: Trimming a Dip socket

After the sockets are added then mount the various test points (TPs) which are either for test measurements at various points of the circuit or for use with headers to act as switches. Such switches, incorporated into the board, rather than being wired from the board to an external SPDT switch, are used with the decade counter (integral to the battery swapper system), the coil voltage selection (in conjunction with an optional Buck and a Boost converter) and for the active device selection. Details of the combinations of jumpers to be used for specific supply options for the circuit and coils are given in the section on ‘Coil Voltage’.

The 2 pin connectors are also known as single line male pin headers and should be inserted with the shorter length of leg into the board. The longer length projecting above the board will be needed for the switch jumpers to fit over to create sufficient contact area to be secure and allow adequate amperage to flow. In the case of the test points, having the length helps attaching probes and clips.

After this, solder in the large 16A relay and plug in the two IC chips (U1 and U2). The two smaller signal relays are push fit into the dip sockets and so can be quickly inserted now or later. Fig 9 shows various stages of assembly and the use of a heat sink clip to hold one leg of a component in place on the underside while soldering the other.

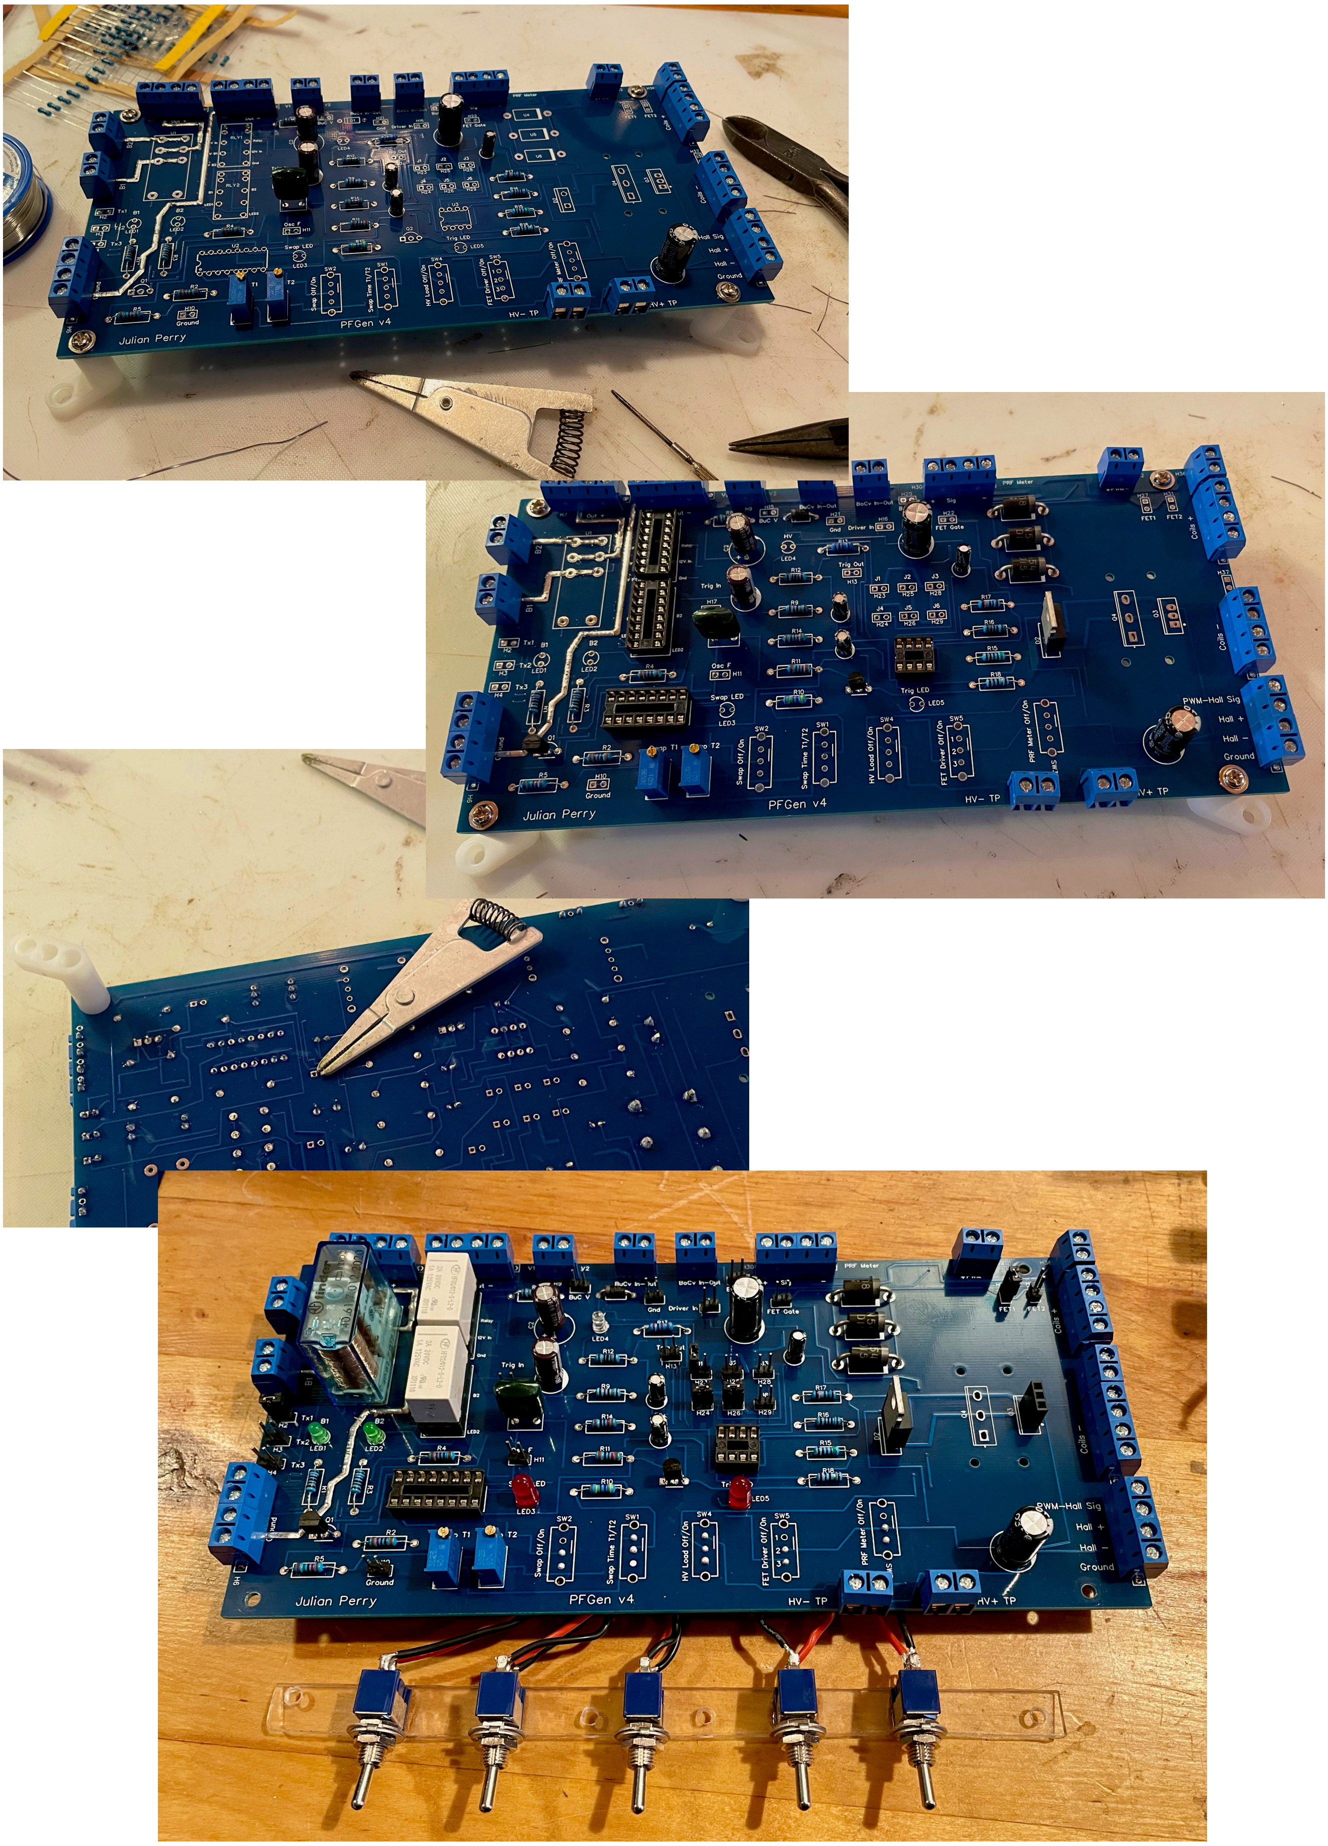

Fig 9: Pics of assembly stages

THE BATTERY SWAPPER

The battery swapper is a central and important part of the circuit comprising an accurate timing mechanism and a set of three relays.

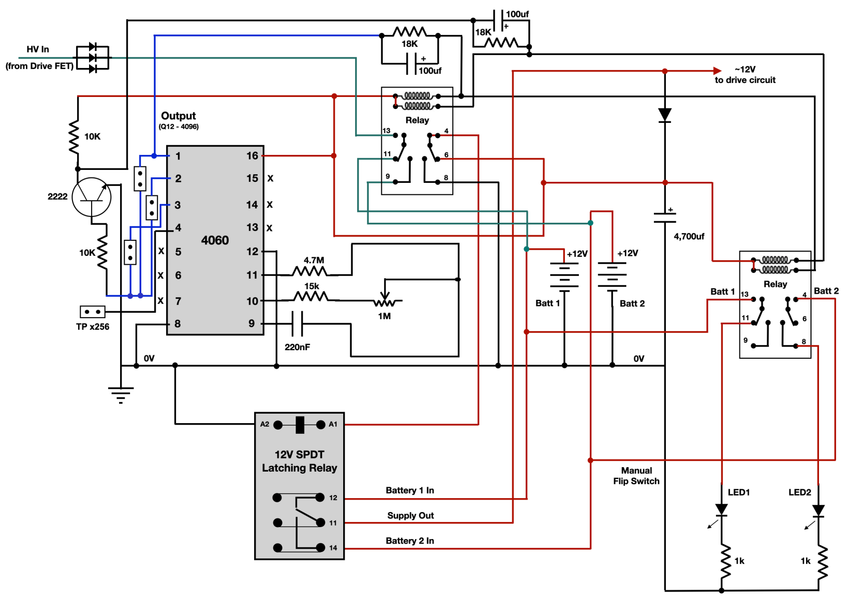

Fig 10: Battery Swapper Circuit

Its role is to allow one of the two batteries to be pulse charged while the other supplies the energy to the circuit and any external load. Then, after a preset time, the batteries swap over their respective roles and the now depleted battery is pulse charged while the battery that is now more fully charged becomes the supply for the circuit and the external load.

The interval between the swap events is an important factor in the CoP and power results since it determines the whereabouts on the charging profile that the pulses act on the battery. The position on the charging profile will be influenced by the load used and the rate at which the battery responds to the HV pulses at a given PRF.

Looking at the design used for the v4 board in Fig 10, the swapper is located on the left hand side the PCB immediately after the battery inputs as shown in the 2D view in Fig 11. This location, delineated by the red box, is so that the main relay can route the current from the supply battery to both the circuit and the external load, while one of the smaller relays triggers the larger one and the other routes the HV pulses to the receiving battery. At swap over, the large contacts within the main relay move to pick up the other battery input and similarly, the smaller relays are enabled to route the HV pulses to the appropriate battery and trigger the correct LED (green LED 1 or 2) to indicate which battery is now providing the power.

The relays used must be able to handle the maximum expected current and most of that demand will result from external load. As such, and particularly importantly for the PCB tracks from the main relay to the output terminals of the PCB (H7, H8) and shown as the wide yellow tracks where they are exposed Copper, the tracks should be thickened with solder as explained earlier. Without the addition of the solder, the tracks would need to be impractically wide to accommodate up to 16A.

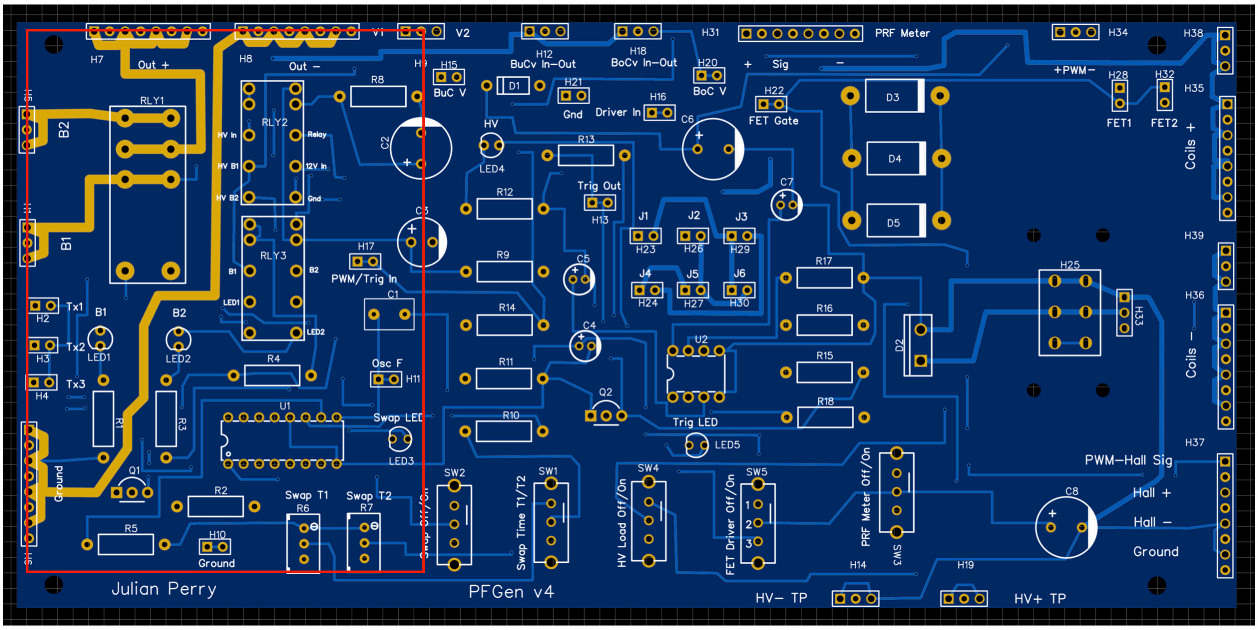

Fig 11: Battery swapper on PCB

The need to have an accurate timing mechanism allows for the optimum positioning of charging on the charging profile by determining how much energy from the supply battery is discharged before the swap interval is reached. Fig 12 shows the sequence of the energy flow.

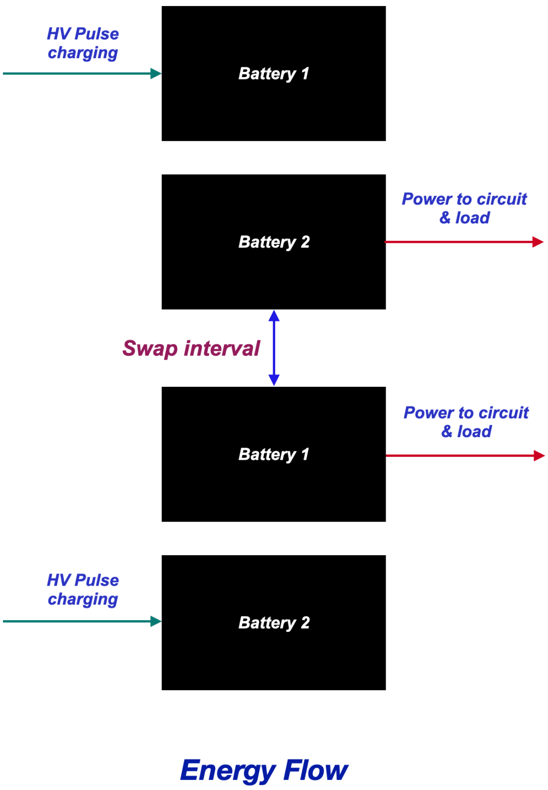

Fig 12: Energy flow using swapper

The basis of the timing mechanism is the CD4060BE decade counter, which has a built in oscillator and is fairly independent of temperature drift. The onboard oscillator frequency is set by several external resistor and capacitor components, in the same way that a 555 timer will use them to set up an RC circuit.

Once the oscillator is running, then the square wave pulses it produces are counted up by the various registers in the chip and result in the various pins going from low to high after a fixed number of cycles and then back to low again after the same number of cycles. The process is repeated until power to the decade counter is turned off.

In operation, the output of the appropriate 4060 pin is routed via a transistor to one of the coil inputs in one of the small relays, the other coil input arriving directly from the 4060 output pin. It is the fact that one of the relay coil inputs is high and the other low, due to the action of the transistor, that results in the coil being activated and operating the relay and which then sends its 12V output to the larger relay to enable it.

Looking at chip's outputs in more detail, with reference to the Fig 13, we can see that the pins are annotated with the letter Q and a number, for example, pin 5 is Q4 representing 2^4 (16). This means that after every 16 cycles of the clock/oscillator, pin 5 will go high and low again after another 16 cycles. So in effect it is dividing the clock frequency by 16 or 2^4.

Similarly pin Q12, will undergo a change from low to high every 4096 (2^{12}) cycles of the oscillator and is therefore dividing the oscillator frequency by 4096.

The chips can vary as to the location of the various Q values, as shown in my circuit and

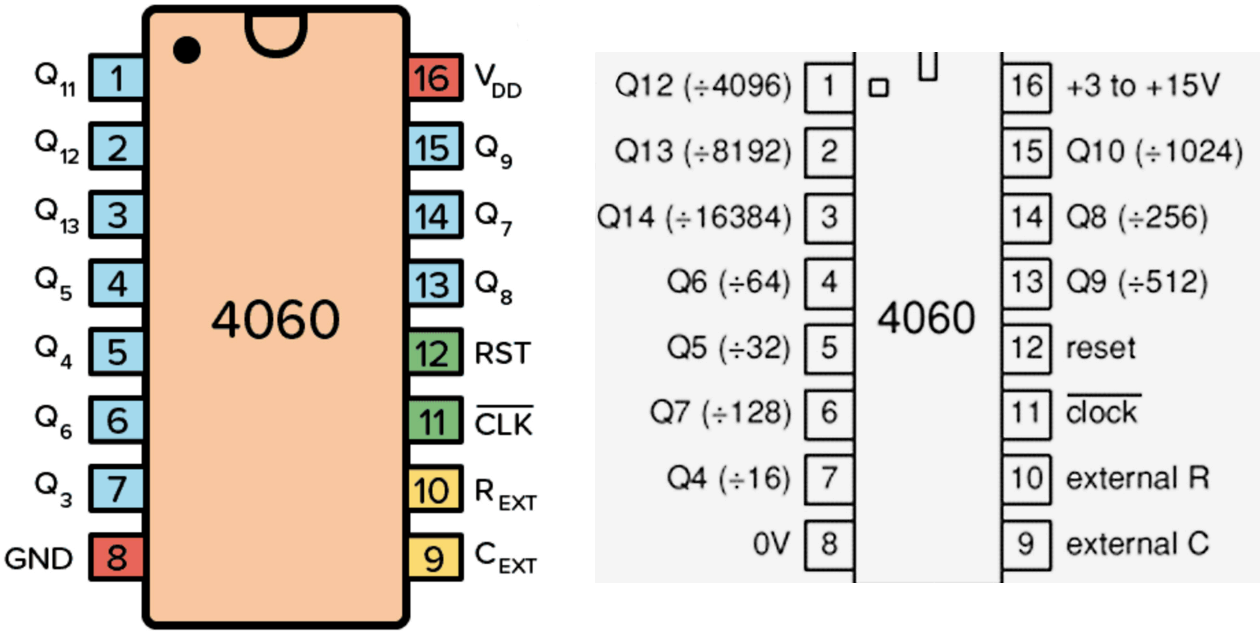

Fig 13: 4060 pin configurations

the two figures, but it will be accurately described on the spec sheet for a particular make of component.

The value of this dividing function is that we can set the oscillator frequency such that the output of a pin will change at any desired time based on the clock frequency and the output pin we choose to use. The pin output, in going from low to high, can be used to trigger a relay coil and flip the relay or, in this case result in the sequential triggering of another relay.

In Fig 14 we can see that for every complete cycle of the clock on the top row (from any point on the square wave back to the same position), the pin below on the second row only does half a complete cycle and the one below that a quarter of a cycle.

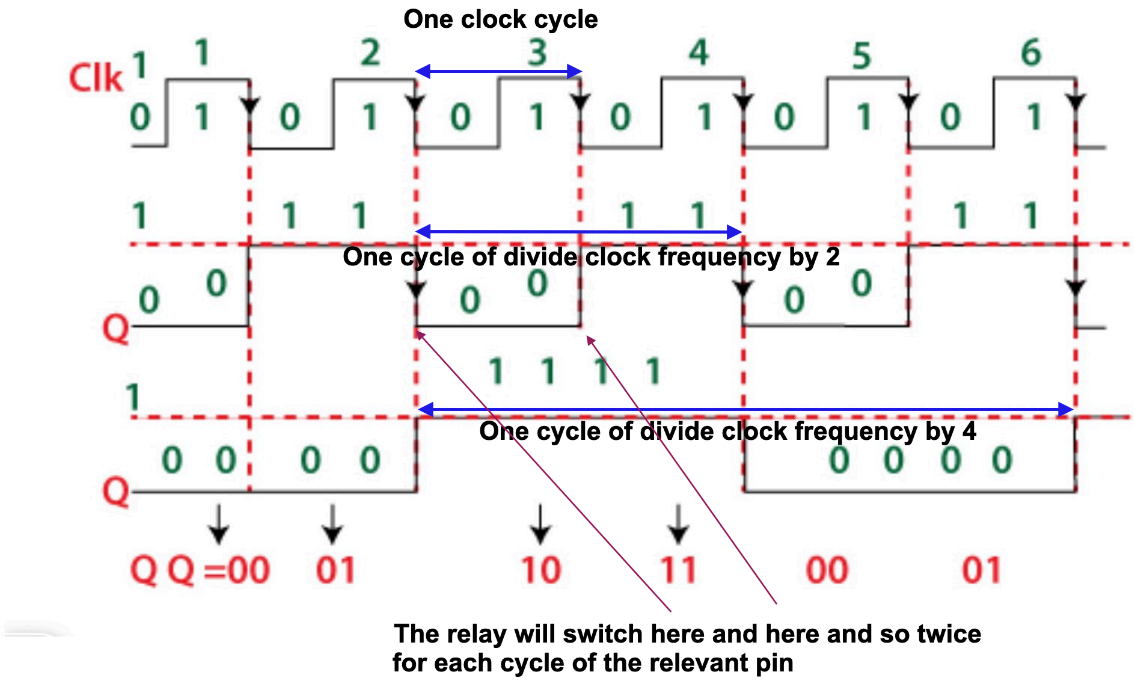

Fig 14: Various cycle frequencies

It is important to note that, having explained how the chip can divide frequencies, that the interval between each swap event is half the time for a complete cycle. This is simply because the relay will switch when the pin goes both high and low and so, with reference to the middle row in Fig 14, for each cycle (or period if you want to think in terms of time), then the relay will switch over twice, once when the pulse goes high and again when it goes low. If we were connected to the pin showing as the bottom line in the figure, the clock frequency is divided by 4, and the swapper would be switched at twice that frequency, i.e. half the frequency of the clock pulses. This is why I refer to a ‘swap interval’ and not a ‘swap period’ since correctly speaking the period is the time taken for a complete wavelength (cycle) to occur and using the same term will cause unnecessary confusion.

So, without going into details of the internal cascading process, the clock frequency and the different pins provide us with a means to set the interval for the battery swapping. Now we can move on see how, numerically, we can adjust the swap interval.

Table 5 displays the simple numerical equations you will need to calculate the oscillator frequency for any swap interval you want, within the limits set by the installed oscillator components (resistor and capacitor values). Conversely, you can calculate the swap time from a particular frequency. Here are some practical examples so you can see how they work in practice.

To Calculate:	From a desired swap time T in minutes	Comments
Osc. Frequency (Hz) using J1	34.43/T	From RC components: 1/(2.2xC1xR1)
Osc. Frequency (Hz) using J2	68.86/T	
Osc. Frequency (Hz) using J3	137.72/T	
Q4 Frequency (Hz)	2.15/T	'X256' test point
To Calculate:	**From an oscillator frequency**	
Swap time (min) using Q12	34.43/F	Using Jumper 1
Swap time (min) using Q13	68.86/F	Using Jumper 2
Swap time (min) using Q14	137.72/F	Using Jumper 3
JI, J2 & J3 refer to the jumper positions, to the left of the 4060 chip, to enable other outputs to be used for swap times larger than 15min (with the present RC values)		

Table 5: Equations for use with the timer circuit

Oscillator Frequency:

Suppose you have installed circuit values for R1 and C1 of 1MΩ and 220nF respectively, and where R1 is in fact a pot trimmer that you can adjust from 15k to 1MΩ (the 15k is provided by a fixed resistor to stop the effect of having the trimmer set to 0Ω).

If you apply the equation at the end of line one of the table, 'From RC components' then this is: $1 / (2.2 \times 1\text{x}10^6 \times 2.2\text{x}10^{-7}) = 2.07$Hz and with the minimum values possible with these components: $1 / (2.2 \times 1.5\text{x}10^4 \times 2.2\text{x}10^{-7}) = 137.7$Hz

These two values of 2 and 138Hz, to the nearest whole number, are the minimum and maximum values of the inbuilt oscillator frequency that you can have based on your installed components and are measured with the scope using H11 (Osc. F) and the ground clip on H10. Using these two values we can calculate the longest and shortest swap intervals that you can set depending on which output pins you choose to use. Clearly the higher oscillator/clock frequency of 138Hz will do its thing inside the chip, and cause all the pins to change their state, much faster than the lower value of 2Hz.

Using the equation on the 5th line for 'Swap time (min) using Q12' which is 34.43/F, with the maximum and minimum frequencies we have determined from the components, you can see that we will have a minimum swap interval of 34.43/F = 34.43/138 = 0.25min = 15s, and a maximum interval of 34.43/2 = 17.22min = 17min 13s. This means that if you are connected to the Q12 output pin, with the jumper on H2, then with the minimum and maximum clock frequencies that you can generate, you will be able to set a swap interval between 15s and 17min 13s, give or take a small percentage.

If we would like longer intervals then we can employ one of the other output pins, such as Q13 and Q14, using the jumpers H3 and H4 also positioned to the left of the 4060. These connect the relevant output to the base of the transistor so that just switching the jumper position will bring a different frequency divider into play and therefore a different multiplication of the swap interval.

Normally we have a swap interval in mind that we want to set up and we would like to know the oscillator frequency so we can quickly set that using a scope using the 'Osc. F' test point (with the scope earth lead attached to a nearby Ground point) and we will then have the desired swap time or interval. This is very much faster than making a clock frequency adjustment and then waiting for the relay to swap using a stopwatch.

For example, if we want a swap time of 10 mins then we must again use the equation at the top of the table which is F_{Osc} = 34.43/T this gives us 34.43/10 = 3.44Hz. To set this we adjust one of the two trimmers, one for each of the two times we can select with switch SW1, so that the clock frequency is 3.44Hz. To achieve this turn the small trimmer screw clockwise to lower the frequency, and so increase the swap interval, or anticlockwise to increase the frequency and reduce the interval, as shown in Fig 15. After setting this frequency, pin Q12 will go high after 10min (600s) and low again after another 10min and keep doing that until you turn off the power to pin 16 of the 4060 chip using switch SW2.

Using other output pins:

The above calculations were done based on using the Q12 output, usually pin 1 or 2 on the actual component. However, you can get longer times than these by using the Q13 or Q14 outputs that divide by a further 2 and 4 respectively. This option is available by using the aforementioned jumper pins to the left of the 4060 chip. If the top jumper in in place then Q12 is being used for the output, the next down H3 is using Q13 and H4 is using Q14. This gives a greater range of options when perhaps wanting to use longer swap intervals and without having to build in an excessive resistance value to R1.

To give some examples, if you want to use a swap time of 20mins, since this is more than the 15min max using the Q12 pin, then you can set the oscillator frequency for 10mins at 3.44Hz (34.43/10) and use the Q13 output, which takes twice as long to change state than the Q12 pin. If you wish to use the Q13 output then use the equation in line 6, which is 68.86/20 = 3.44Hz. Again just use the desired swap time and the right equation for the pin you are using as set by the jumper header position Tx1 (H2), Tx2 (H3) or Tx4 (H4), and measure the frequency at TP-H11 (Osc. F).

The circuit has two timer options set by a switch to T1 or T2. My own preference is to set T1 at the minimum of about 15s so I can use that to check the operation of the swapper and to make a 'manual' change to the battery that is providing the power if I need to. T2 is then set to my preferred swap interval depending on what operational parameter I need using the above process.

Here are some other examples:

To set a swap interval of 8min use either jumper 1 and adjust F_{Osc} to 4.3Hz (34.43/8) or use jumper 2 and set F_{Osc} to 8.6Hz (68.86/8) and which, if you were using Q12 and jumper 1, would result in 4min.

To set a swap interval of 25min, which can't be done using Q12 as it's longer than the maximum available) use jumper 2 and F_{Osc} to 2.75Hz (68.86/25) or use jumper 3 and adjust F_{Osc} to 5.51Hz (137.72/25).

To set a swap interval of 42min use jumper 3 and F_{Osc} to 3.28Hz (137.72/42) since 42 min is longer than either J1 or J2 can generate using the RC values in place.

To set a swap interval of 16min 20sec, convert to decimal minutes as 16.33min then, using jumper 2 (Q13), adjust F_{Osc} to 4.22Hz (68.86/16.33)

These techniques will allow you to easily setup your timer to swap the batteries over to within an accuracy of a few seconds.

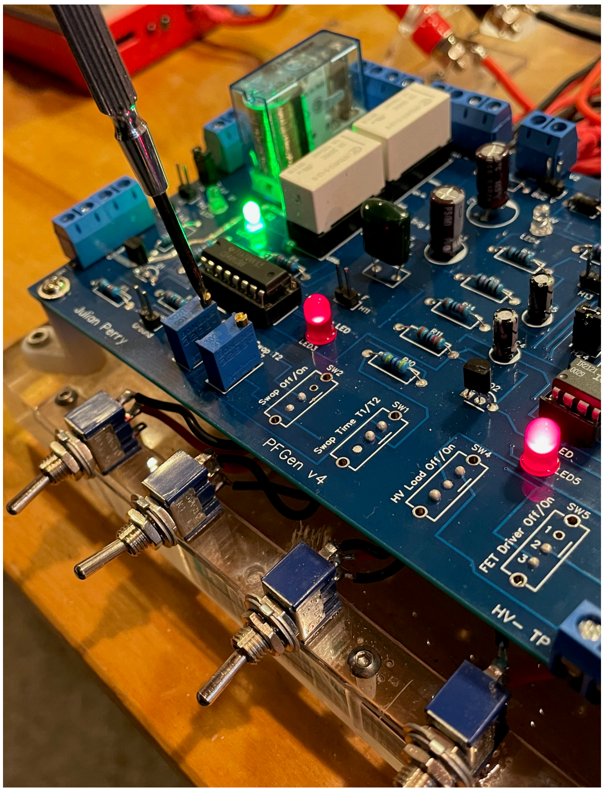

Fig 15: Adjusting the T1 swap interval

TRIGGER CIRCUIT

The function of the trigger circuit is to produces a square wave input to the Gate of the MOSFET or IGBT to initiate the flow of current in the solenoids and back through the Drain and Source of the active device to Ground. This can be achieved in various ways including using a rotor and Hall sensor assembly, a 555 timing chip with associated RC components or a PWM module.

Early tests using the rotor, while clearly demonstrating that the generator was using energy by virtue of the approx 3,000 rpm that the rotor reached, showed that the CoP results were not as good as those achieved using a flexible PRF system. With the rotor's 5 permanent magnets, which triggered the Hall sensor 5 times with each revolution, the PRF was essentially fixed by the rpm it achieved. In this case 250Hz (3,000/60 x 5).

The 555 chip, while easy to use and design, was subject to some temperature drift over the course of experiments as components heated up. This required constant checking of the frequency generated using the installed frequency meter. Additionally, the simple RC component design tended towards changing the duty cycle (the proportion of the period where the signal is high) as the frequency was adjusted. Since duty cycle transpired to also be an important factor in minimising the energy supplied to the generator, it was required to be stable and easily selectable.

Fig 16: PWM - Signal generator unit

These issues were successfully resolved by using a PWM module of the type that I had used in 2018 with my fully solid state generator build. They are inexpensive, easily available and allow for precise setting of both the frequency and duty cycle with no obvious drift. The type recommend is shown in Fig 16 and exists with various types of control formats, such as with a dial knob version and 'soft press' controls as shown here.

They typically take a 10-20V supply, provided via H34-1(+) and H34-2(-) on the board and the PWM signal output being connected to H37-1. A viable trigger input to the PCB will show as LED5 being lit with each incoming square wave trigger pulse and therefore, above about 25Hz, showing as continuously on.

WIRING AND SWITCHES

Most of the wiring used in the generator is silicon coated AWG18 that is sufficient for modest amperage of less than 10A and is very flexible. This would be used for all the coil connections and to and from the Buck and Boost converters, if you are choosing to use them, as well as to and from the two batteries. For the external loads then AWG16, or potentially AWG14, should be used. Where one can use the smaller diameter AWG20 wire is for the connections to the two small (and not very accurate panel meters), to and from the Hall sensor, if you are using a rotor, the supply and signal connections to the PWM module and for the SPDT switches. These applications are listed in Table 6.

AWG	Connecting
All wiring is of the silicone coated type for maximum flexibility	
20	**To PRF meter via H31, PWM module via H34, for both supply and signal, panels meters via H9 and wiring to SPDT switches (SW1-5)**
18	**From batteries 1 and 2 to PCB, ground wire from PCB H6 to panel meters, fuse and main switch and back to battery common negative terminal. Also to and from the coils (H35, H38, H36 & H39), Buck and Boost converters (H12 & H18) and the optional potential divider, for HV peak pulse measurements, via H14 & H19**
16	**To external loads via H7 & H8 (suggested up to 15A). For higher power use 24/36V multiple batteries**
8	**Used for capacitive discharge circuit delivering ~100A pulses**

Table 6: Wiring gauges used

While it is possible to just insert the wire, minus a few millimetres of its silicone coating, into the terminal blocks, it is preferable to add some solder to the end of the wire. However, inserting a soldered wire into the terminal blocks can be rather tight so it is suggested to crimp down on the soldered end to make a slightly wider and flatter shape to insert into the terminal holes that have been unscrewed as far as possible.

An alternative to soldering the ends, or just leaving them unsoldered, is to use ferrules over the wire end and to crimp that. This may be seen as overkill by some as the applications is not subject to vibration or other factors that would require the use of ferrules but for some it is a tidy option even if not really necessary.

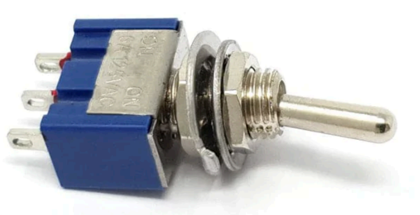

Fig 17: Mini SPDT Switch

Besides the switches on the board that use the small 2 pin jumpers, there are a selection that exist as mini SPDT switches, of the type shown in Fig 17. It is strongly recommended to mount these switches on a detachable plate so that when you need to lift the PCB to get to the underside, you do not need to unsolder all the connections, a fiddly and time consuming task as you will no doubt discover when wiring them up.

Using some form of mounting plate, such as a piece of perspex, you can lift all the switches away together with the board. Given the revised generator design, certain switches previously used are no longer required and the new circuit requires just five as shown in Fig 20. This diagram contains details of how to actually wire the switches up and therefore to ensure the correct operation when a switch is turned on. A detailed description will be given regarding how to wire up a switch with the correct connections from the PCB to the terminals at the back of the SPDT switch.

Looking a the schematic representation in Fig 18, we can see that the default position for the switches is where terminal 2 inside the symbol is connected to terminal 1, and this is the default ‘off’ position (the two green Xs just indicate that those parts of the symbol are unused). When the switch is turned ‘on’, then terminal 2 will be connected to terminal 3 and, with this being so, we want the feed to terminal 2 to be what we are going to switch by the action of the switch.

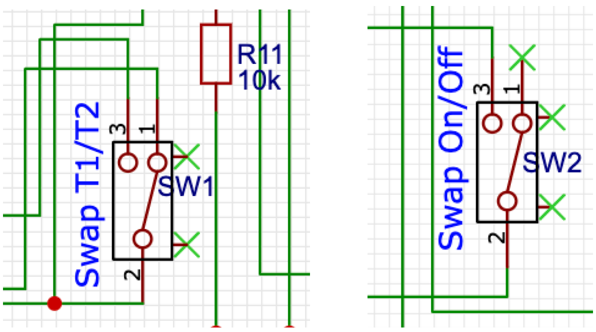

Fig 18: Switch Schematics

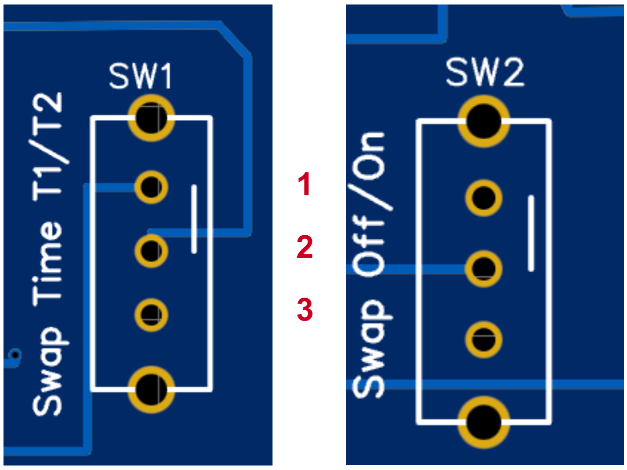

Fig 19: PCB switches

In the example of SW1 in Fig 18, which selects which of two resistor trimmers are engaged to determine the oscillator frequency of the 4060 timer chip, then terminal 2 will be connected to the port on the chip that requires the RC components to set its frequency.

If you were to look at the whole schematic you will see that terminal 2 connects to the capacitor C1 and which is part of what determines the frequency. So in this switch, terminal 1 will be to one of the 1M trimmers and terminal 3 to the other trimmer resulting in the change in frequency. Looking at the PCB arrangement in Fig 19, the white line to the right of two of the connecting pads represents the default switch position such that when the switch is operated the line would theoretically slide downwards to connect the middle contact with the bottom one instead of the top contact. So here the middle pad 2 is the ‘common’ one, the top pad 1 is the default connection and the bottom pad 3 is the ‘on’ position.

Now relaying that to an actual single pole double throw switch (SPDT) as shown in Fig 17, if you hold the switch in the orientation that you want to secure it to the mounting plate, then with the switch lever in the upward position, the two terminals that are electrically connected are the middle one and the lower one. Similarly, when you flip the switch down

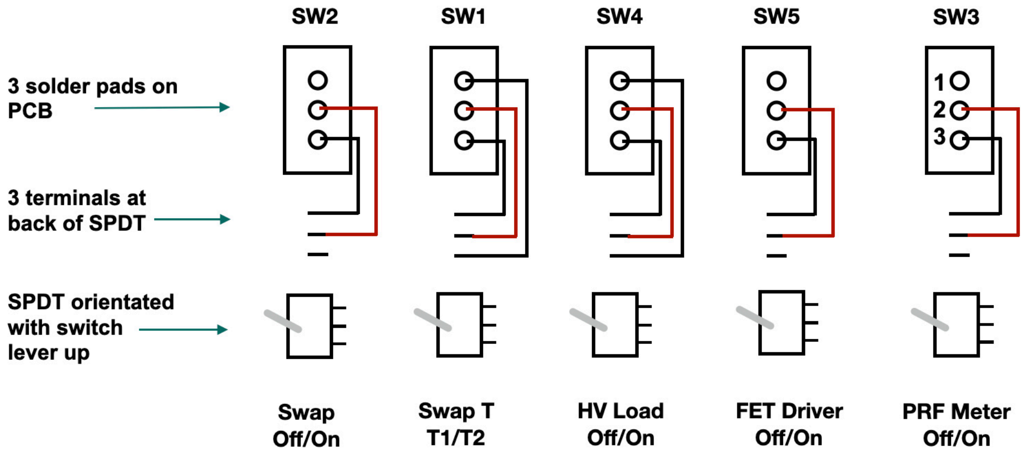

Fig 20: Switch connections

the middle terminal connects to the upper terminal - perhaps the opposite of what you might expect.

In translating these details to how to connect the solder pads on the PCB with the physical switch, we need to refer to Fig 20 where the relevant connections are clearly shown (this image is also in the Appendices). With reference to SW1 in Fig 19, the top solder pad on the underside of the board is wired using AWG 20 wire to the bottom terminal of the switch and the middle pad to the middle terminal. This will connect the correct track on the PCB with the switch in the default or 'off' position. Connecting pad 3 to the upper terminal on the back of the actual switch will then connect pads 2 and 3 together when the switch lever is down and 'on', since doing so connects the upper and middle switch terminals.

The same principles are applied to the other four switches, although in SW2, SW3 and SW5 there is no connection in the default position but only when the switch is 'on', thereby connecting solder pads 2 and 3 via the middle and upper terminals of the SPDT switch. Only SW1 and SW4 have connections in both switch positions, the former selecting the resistor trimmer for the swapper timer, and the latter the routing of the HV pulses to the receiving battery, as directed by the relay RLY1 positions, or to the test terminals H14 and H19, where a potential divider can be connected to use with a scope. This will be discussed in more depth in the relevant section.

Since wiring up the switches can be a fiddly task, here is a breakdown of the steps, with accompanying pictures, that will help make it a smooth task that can be completed in less than an hour. Fig 21 shows these steps in a range of images.

1. Place the strip of switches next to the PCB to approximately match where it will sit on the base.

2. Pre solder all the PCB switch holes that will be used (some switches will use just two holes and others all three).

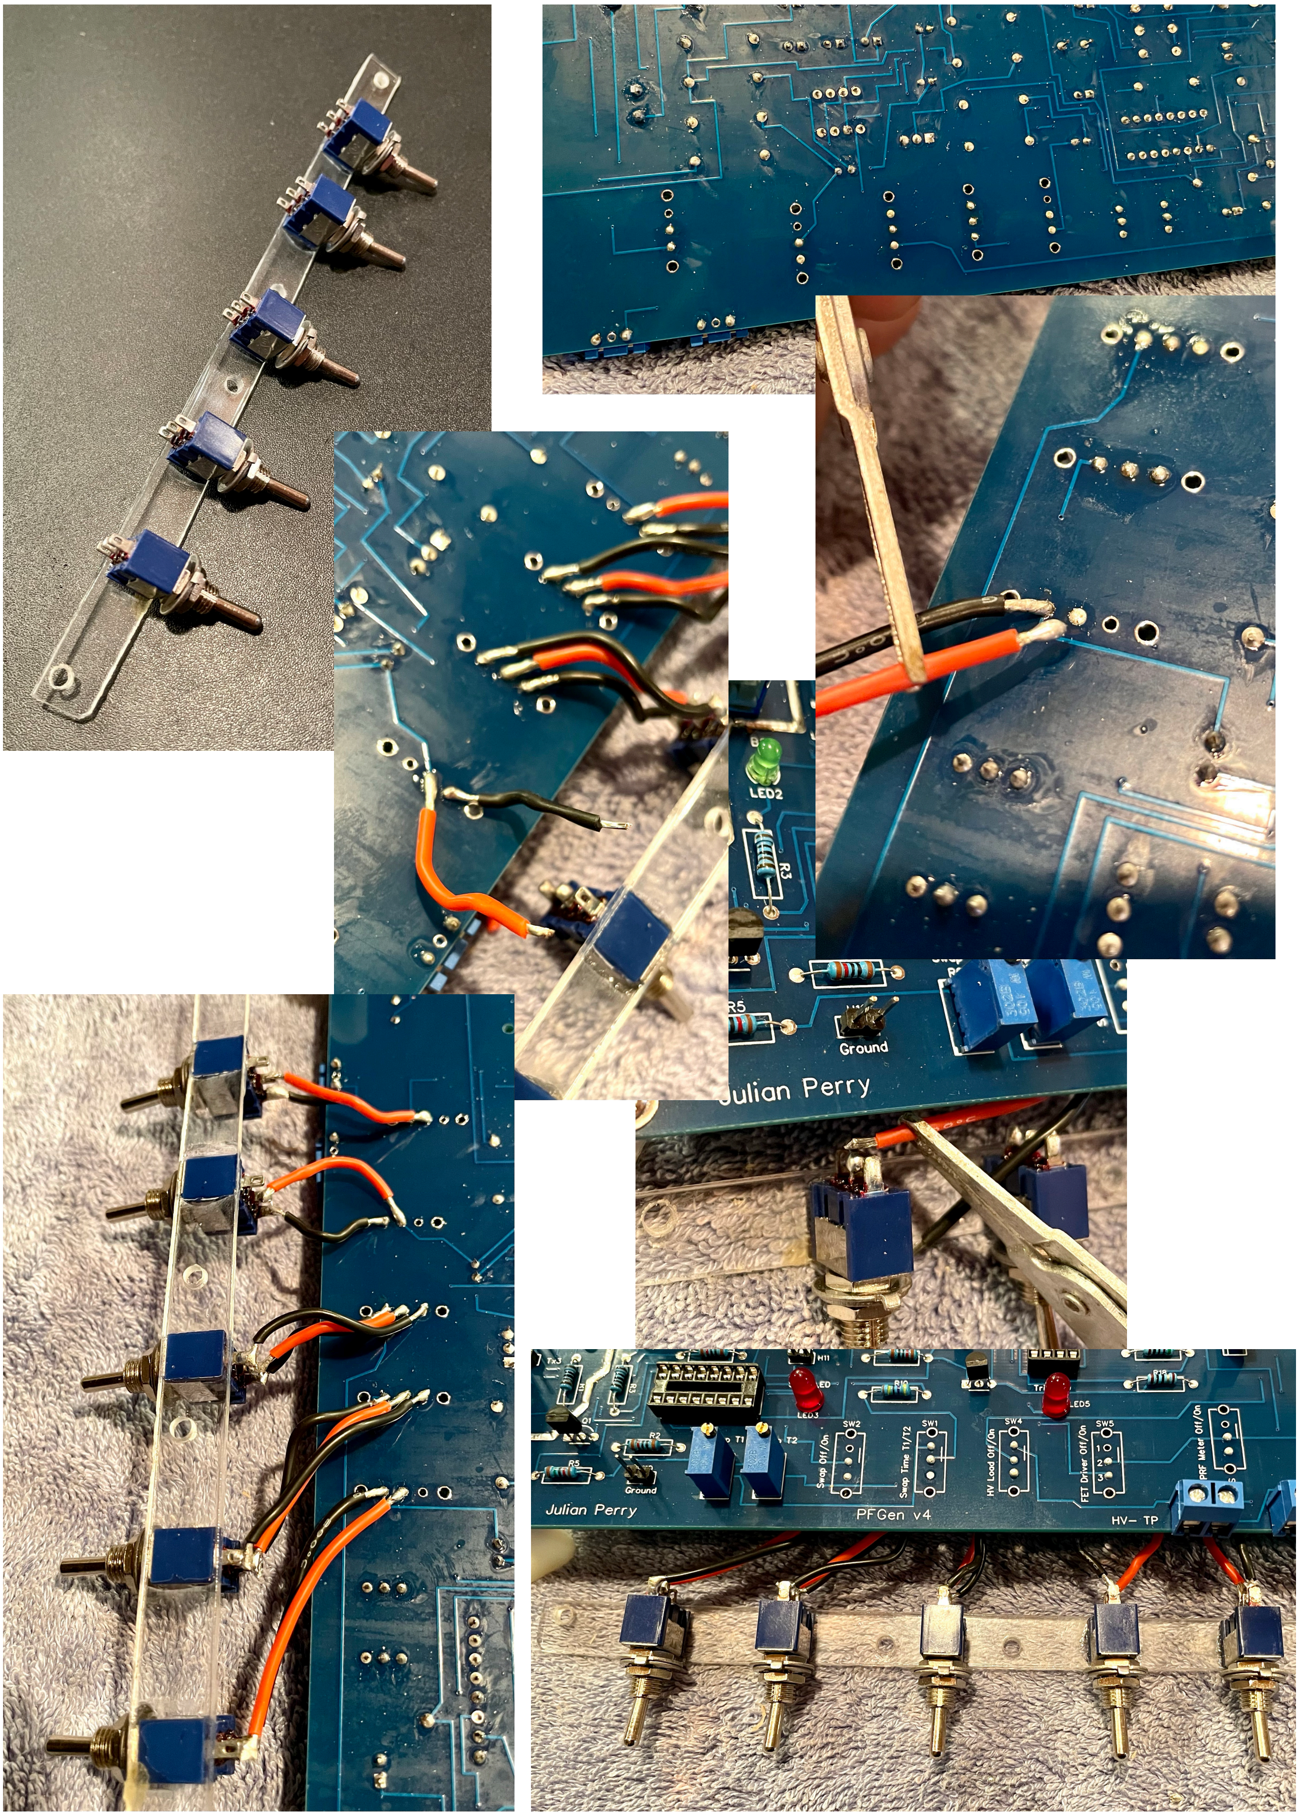

Fig 21: Switch installation

3. Cut lengths of AWG 20 wire a bit longer than needed, strip one end 2-3mm and flood end with solder. Using red wire for the common connection - PCB pad 2 - to the centre terminal of the switch, will help identify which wire goes where using the ‘Switch connections’ schematic sheet.

4. Turn the PCB over and connect the soldered wire end to the underside of the relevant PCB switch contacts using the ‘Switch connections’ sheet.

5. Turn the PCB over again and trim the wire lengths to meet the switch terminals with a bit of flex.

6. Carefully strip the wire ends (be aware of the force that may be applied to the wire-PCB switch tab junctions) and add solder to them

7. Turn the PCB back the right way up and using a heat sink clip or pliers hold the wire end to the correct switch terminal and solder it. Start with the lowest terminal and move towards the top.

8. Repeat for the next switch.

Once completed the PCB can be installed and removed from the generator as a whole assembly with ease, complete with the switches.

Perhaps the most frequently used of these switches is the FET Driver. The advantage of this switch is that you can have the main power switch on but no pulses being produced since the supply to the driver circuit (U2) is turned off with this switch. This aids precise timing of events when one needs the main circuit to be on, in order to set the PWM module going and which, in my particular model, requires a soft press button on the module to be closed after the power supply to it is on to start generating pulses.

While the MOSFET/IGBT is firing you may hear a quiet buzzing sound, depending on the PRF, but the blue HV LED will be brightly lit to indicate the FET driver output is active and therefore that HV pulses are being generated.

ACTIVE DEVICE SELECTION

The choice of active device used, in this case either a MOSFET or an IGBT, is more important that has previously been conveyed in the work of others. This is in part due to it imposing a limiting factor on the maximum flyback voltage due to the device's 'avalanche rating'. When pulses arrive at the Drain with a higher voltage than the avalanche rating value, a breakdown occurs within the device that effectively grounds the pulses above that value. This then clamps or clips the peak voltage as shown in Fig 22. A document explaining 'avalanche breakdown' in greater detail is in the Appendices and Fig 23 shows some relevant data sheet specifications.

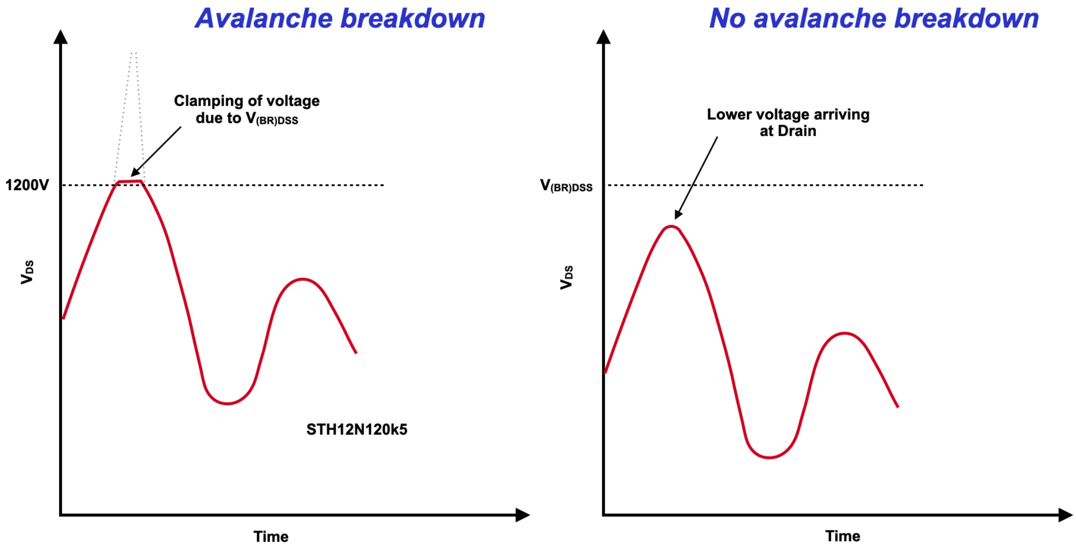

Fig 22: $V_{(BR)DSS}$ clamping due to internal breakdown

Electrical characteristics

(T_{CASE} = 25 °C unless otherwise specified)

Table 4: On/off states

Symbol	Parameter	Test conditions	Min.	Typ.	Max.	Unit
$V_{(BR)DSS}$	Drain-source breakdown voltage	V_{GS} = 0 V, I_D = 1 mA	1200			V
I_{DSS}	Zero gate voltage drain current	V_{GS} = 0 V, V_{DS} = 1200 V			1	µA
		V_{GS} = 0, V_{DS} = 1200 V, Tc = 125 °C			50	µA
I_{GSS}	Gate body leakage current	V_{DS} = 0 V, V_{GS} = ± 20 V			±10	µA
$V_{GS(th)}$	Gate threshold voltage	V_{DS} = V_{GS}, I_D = 100 µA	3	4	5	V
$R_{DS(on)}$	Static drain-source on-resistance	V_{GS} = 10 V, I_D= 6 A		0.62	0.69	Ω

Fig 23: $V_{(BR)DSS}$ and other specifications

A quick calculation of the flyback voltage using the coil inductance, the current and the switch off speed of the active device, indicates that the flyback voltage can easily be several tens of thousands of volts. However, that voltage is never seen at the device Drain due to this clamping effect and is orders of magnitude smaller. Using commonly available devices, such as the ubiquitous IRF840 that is cheap and freely available, will clamp the voltage at 500-600V whatever the theoretical voltage arriving from the coils.

The observed energy gain seems to be a function of dV/dt and the electric field stress resulting from it, and so the simplistic view is that the higher the spike voltage seen at the Drain the higher the resulting CoP value will be. In fact the issue is more nuanced in that the Lead Acid and Lithium batteries respond differently to the increased spike voltages. For example, the 18Ah $LiFePO_4$ battery responds roughly the same at 1.0kV as at 1.5kV but drops significantly at 1.1kV. The 7Ah Lead Acid battery peaks at 1.7kV, dips at 1.0kV and rises substantially at 0.5kV.

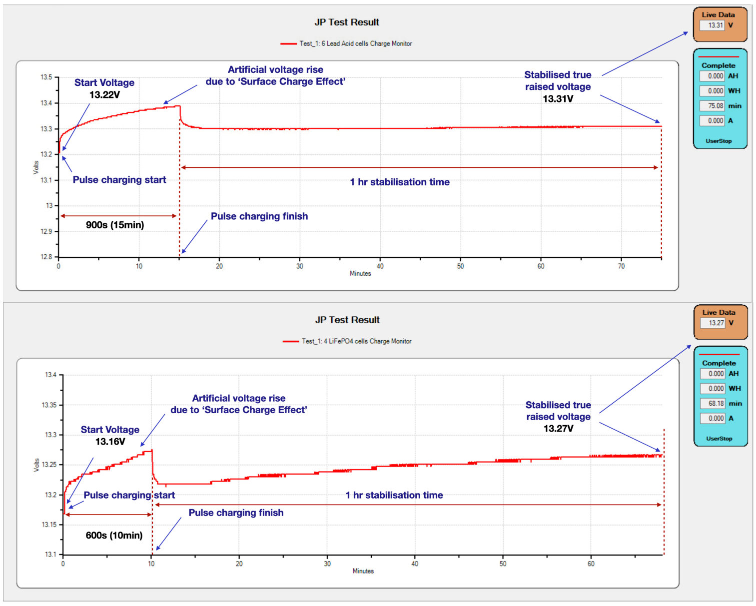

Fig 24: Recovery and battery type

For a 7Ah $LiFePO_4$ battery the optimum HV was 1.1kV. As with every other variable, the optimum kV will depend on your configuration and experimentation is required. However, the results I have provided may give you a head start in what to try, along with the suggested starter settings provided in the Appendices.

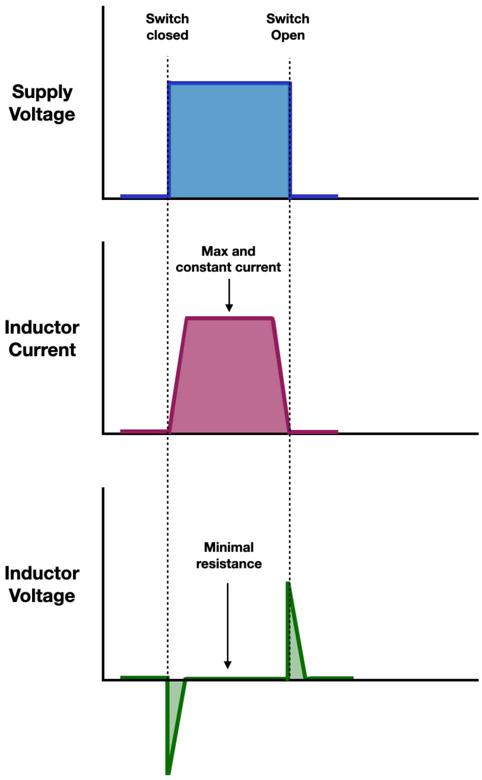

Fig 25: Behaviour of an inductor

It was also noted that during the stabilisation period after pulse charging, the SLA battery voltage remained essentially flat after the initial fall off, whereas the Lithium battery showed a drop, as expected and then a slow recovery, as depicted in Fig 24. Again, only the power tests will clarify the overall response when the batteries are swapped continuously instead of being left to stabilise after pulse charging.

For much of the testing dV/dt was around $1.5x10^8$ V/s and was sufficient to realise the CoP>1 results. The limiting factor is the 'avalanche rating' of the active device and not the coils, which will normally be creating a flyback voltage (as indicated by Faraday's and Lenz's Law) much higher than what you see at the Drain.

Making larger coils is not required to obtain sufficient voltages to demonstrate the energy gain process. Equally, using for the highest voltage possible will not necessarily achieve the best results. As with all the other parameters, there is an optimum value for the particular battery and setup so one needs to be able to change over the active device to suit the situation and provide the optimum HV for the pulses.

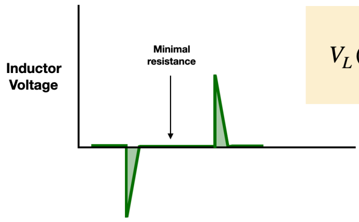

$$V_L(t) = -L\,\frac{\delta i}{\delta t}$$

$V_{L\ (t)}$ = - 0.35 x 0.25/1E-6 = 87.5 kV!

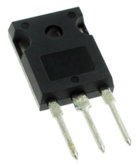

We don't see this high voltage at the MOSFET drain, and transferred to the receiving battery, since any voltage above the 'avalanche breakdown' value is taken to ground.

Fig 26: Calculating flyback voltage

The v4 PCB was initially designed to have two sets of holes for two selected TTH devices to be soldered in place, and then the active device selected using one of two headers (H28 & H32). This was also partly due to it proving hard to find a suitable socket that was easily available to plug in and unplug devices quickly and which would allow the use of a heatsink. Having now found sockets that are relatively easy to obtain, I have returned to my original plan to allow for the mounting of two sockets, one for TO-247 and one for TO-220 type devices in the same holes that can also take single regular TTH devices if desired. As such the PCB component list has been updated to include these two sockets, one available from eBay and one from AliExpress.

When fitting the larger socket, cut and remove the pins in positions 2 and 4 so that pins 1, 3 and 5 will fit into the three holes. If using just one TO-247 format device, then solder into the left hand set of holes nearest the heat sink base pin holes.

Once you have a working system and are starting to see results then you can switch the active switching device in operation very easily. The use of the above mentioned sockets means you can quickly unplug and insert a different MOSFET or IGBT device and then select which one for operation by putting the jumper header on either H28 or H32.

Device	***kV***
IRF840	0.6
STP20N95K5	1.0
STP12N120K5	1.1
STW12N150K5	1.6
STW12N170K5	1.7

Fig 27: HV outputs from active devices

The formats for the STW12N150K5/170K5 (Q3) is TO-247 and for the STW12N120K5, STP20N95K5 and IRF840 (Q4) is TO-220, all used in conjunction with the DHG10i1800PA diode which will cope with the maximum HV at the front end of the pulse. Fig 27 shows measured HV values for these devices and indicates that the spec sheet values are only typical.

The use of a heatsink is advised although during operation for periods of up to an hour there is very little heat generated.

However, when it comes to running for many hours, days or even weeks without a break, then heat may become an issue and the heatsink is best fitted.

To install them, I suggest fitting each in turn to the heat sink then placing the assembled item onto the board and soldering the leg that projects through. Doing the reverse will result in the hole in the device body not aligning with the heat sink hole. Once soldered and the legs trimmed, unscrew from the heat sink and repeat with the other one.

At this juncture it is a value to discuss the rationale behind the use of the diode shunting the active device, in this case the DHG10i1800PA.

A MOSFET inherently contains what is known as a body or parasitic diode on account of its intrinsic structure and which is formed by the PN junctions that make up its structure. The role of the body diode is to allow a current path to ground during the period between switching on and off in order to release the charge that collects at the junction during operation and is particularly relevant when switching inductive loads. However, they naturally have a slow response time that can result in a significant build up of charge at the junction and with associated problems depending on the circuit and the application. Adding an external fast recovery diode, thereby shunting the body diode in parallel, will release that charge more effectively than the indigenous body diode and usually prevent the body diode from ever having to switch on.

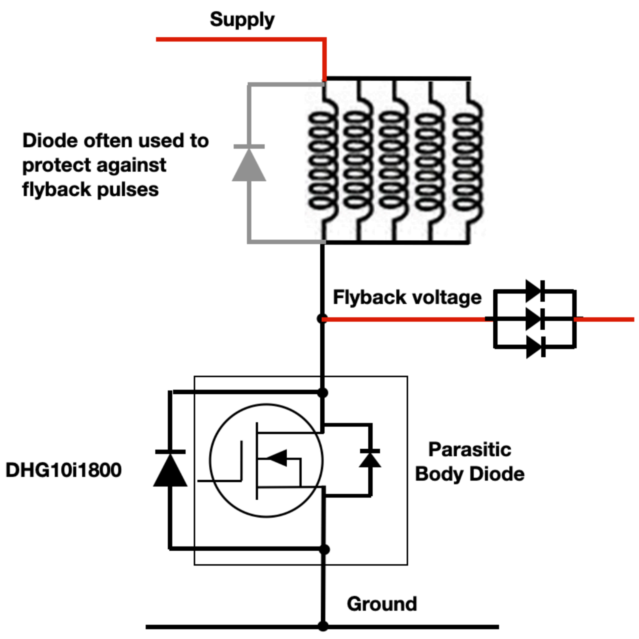

Fig 28: Active device diodes

In typical, regular power switching, where you don't want the flyback pulses appearing at the Drain at all and causing problems of electromagnetic interference (EMI), then another diode is often used to discharge the 'damaging' pulse and which the body diode is unable to achieve since it is in a reverse bias state. This is shown in grey in Fig 28 and would usually discharge back to the supply.

Alternatively, many of these active devices are designed to be 'avalanche rugged' and have an 'avalanche rating' whereby, over a certain voltage, they break down to discharge the pulse before recovering. This latter feature is the limiting factor in the peak flyback pulse voltage as mentioned above. Of course, we want this pulse so we do not add an additional diode across the inductor to discharge it to the battery.

You might reasonably ask why we don't do this anyway since the pulse is directed from the Drain to the battery via the 3 output diodes, D3, D4 and D5. The answer is that shunting the coils themselves with a diode would prevent the spike voltage from building up to the 1-2kV that we do see and need and instead, be shorted to a low value by the diode, along with a small amount of current. This is a very different result from delivering a pulse of high potential directly to the battery terminal which is what we need to elicit the 'phenomenon' at the interface between the pulse and the electrochemistry.

So instead we want and let the voltage spike exist as a pure voltage and electric field delivered directly to the battery terminal, instead of using an additional diode across the coil which would inevitably change the type, quality and magnitude of the pulse. However, we are still limited by the FETs avalanche rating and, if the FET we use didn't have such a rating, then it would be easily damaged by the very pulses we seek to use in the generator.

The external FRED (Fast Recovery Epitaxial Diode), Schottky or another type of diode, is added to support the discharge of the pulse energy generated during the coil switch on stage and which would otherwise build up as charge at the Drain-Source PN junction since the body diode is not there as a design feature but as a natural consequence of the device's structure.

Recognising that we need an appropriate and suitably rugged active device to switch the coils on and off, these higher voltage devices are sometimes hard to come by with the general shortage of silicon devices after COVID, however, the site I mentioned earlier will offer you a good chance to source them, especially if you arrange for postal rather than courier delivery.

So choosing which active device to use depends on the battery you plan to work with and of course you will need to experiment with different values. The ones I have used for all the testing so far are shown in Table 7 which presents CoP data for all the batteries using five different HV values. For example, the optimum response with the 7Ah SLA battery was the IRF840 at about 0.5kV, whereas with the 18Ah $LiFePO_4$ the optimum HV was 1.04kV using the appropriate device. Exploring the most suitable device is a lot easier when using the plug-in mount.

As is apparent, there is no simple relationship between HV and CoP which does have the advantage that even the easily obtainable IRF840, that produces an HV in the 500 - 600V range, is capable of producing some good results. It is the bigger picture that is important with all the optimised factors working together. The advantage of using the higher 'avalanche' rated devices is simply to allow you to optimise the HV to your specific battery and device setup and the devices that you can 'plug & run' will allow you to experiment.

Having installed your MOSFETs or IGBTs, it is very useful, but not essential, to be able to measure the peak voltage and see it on a scope so that any unwanted artefacts can be noted. This is described in the 'Potential Divider' section.

	Battery (Ah)	Chem.	kV	Coil V	PRF (Hz)	Start %Ah	%Ah Disch.	Charge Time (m)	CoP [1]	Ext Power (W) [2]	Active Device
1	7	SLA	0.50	12.50	108	100	20	15	4.82	29.2	IRF840
2	7	SLA	1.04	12.50	108	100	20	15	3.11	19.4	STP20N95K5
3	7	SLA	1.10	12.50	108	100	20	15	4.89	28.6	STP12N120K5
4	7	SLA	1.60	12.50	108	100	20	15	4.90	28.6	STW12N150K5
5	7	SLA	1.70	12.50	108	100	20	15	5.02	34.4	STW12N170K5
6	17	SLA	0.50	13.00	180	100	40	10	5.38	54.6	IRF840
7	17	SLA	1.04	13.00	180	100	40	10	5.37	44.8	STP20N95K5
8	17	SLA	1.10	13.00	180	100	40	10	4.93	33.3	STP12N120K5
9	17	SLA	1.60	13.00	180	100	40	10	4.90	28.6	STW12N150K5
10	17	SLA	1.70	13.00	180	100	40	10	4.24	42.4	STW12N170K5
11	7	$LiFePO_4$	0.50	12.50	108	95	20	10	6.97	56.3	IRF840
12	7	$LiFePO_4$	1.04	12.50	108	95	20	10	5.25	31.9	STP20N95K5
13	7	$LiFePO_4$	1.10	12.50	108	95	20	10	5.72	38.5	STP12N120K5
14	7	$LiFePO_4$	1.60	12.50	108	95	20	10	3.13	16.0	STW12N150K5
15	7	$LiFePO_4$	1.70	12.50	108	95	20	10	7.15	45.3	STW12N170K5
16	18	LiFePO4	0.50	12.00	155	85	10	6	13.14	83.1	IRF840
17	18	$LiFePO_4$	1.04	12.00	155	85	10	6	37.84	219.0	STP20N95K5
18	18	$LiFePO_4$	1.10	12.00	155	85	10	6	19.93	142.3	STP12N120K5
19	18	$LiFePO_4$	1.60	12.00	155	85	10	6	30.63	217.5	STW12N150K5
20	18	$LiFePO_4$	1.70	12.00	155	85	10	6	19.89	142.4	STW12N170K5
Notes	[1] *Uncertainty not shown but calculated in main spreadsheet.*										
Notes	[2] *Theoretical predicted available output power derived from the CoP value and the time taken to charge to $V_{(pk)}$. Used as a guide for the live power tests.*										

Table 7: CoP values vs HV using various battery configurations

COIL VOLTAGE

There has been much talk of increasing the power output of this type of generator by increasing the number of batteries used in series, thereby increasing the voltage supply to the coils to 24 or 36V. This has the effect of also increasing the voltage to the whole circuit and therefore the power to both the coils and the circuit components.

The latter can be a problem since most of the components used tend to have operating ranges in the 8 - 15V range and so anything above that will either require a higher rated set of components or the use of a Buck converter to run that part of the circuit. Using a converter is far simpler and I used one as part of the testing regime to see how adjustments to the coil voltage affected the results.

Of course, when you are using the normal two battery system, then the run battery will slowly drop in voltage, perhaps below the ideal value for the coils depending on the load you have attached. In that case you may also use a Boost Converter to be able to set and stabilise the coil voltage. The PCB is designed to be able to utilise both depending upon your requirements.

However, CoP tests I have done on two batteries so far (using three has yet to be completed) showed no real advantage to CoP over and above other factors that can influence the results. The only clear advantage I can state so far is that when you are powering an external device with 24V instead of 12V, then the current demand is halved with the power remaining the same. This is kinder on your swapper relay, so long as the total power is still within the rating. In an earlier build there was only a small signal relay, a Hongfa HFD2 capable of only 3A. I have since revised the board using the 16A capacity relay (Finder 40.61.7.012.2020) operated in conjunction with two of the smaller ones, one to flip the larger relay and handle the routing of the HV pulses and the other to operate the LEDs indicating which battery is providing power.

Unless and until you expect to be achieving 100W+ outputs, a fact that is as yet unconfirmed, then I suggest sticking with one pair of 12V batteries (see Capacity below) for the whole system or use suitably large capacity ones that can deliver 10 -15A continuously for 15 - 30mins, such as the 18Ah Lithium ones and as discussed it the ‘Batteries’ section. For testing and set up purposes you can use just one receiving battery together with a power supply for convenience.

Although I have no evidence to support it, there is the notion that the HV spikes ride on top of the DC voltage supplied to the coils and so, in effect, provide a baseline voltage level. This is shown in Fig 29 and, if the battery to which the pulses are being directed is at a higher voltage than this baseline levels then the pulses have to overcome a small energy hurdle before they can fully enter the battery and a small amount of its total energy is lost in the process. This theory might explain why using a coil voltage 1 - 2 volts below the current voltage state of the receiving battery has yielded reduced CoP results. On the other hand, raising the voltage above the receiving battery voltage will waste energy

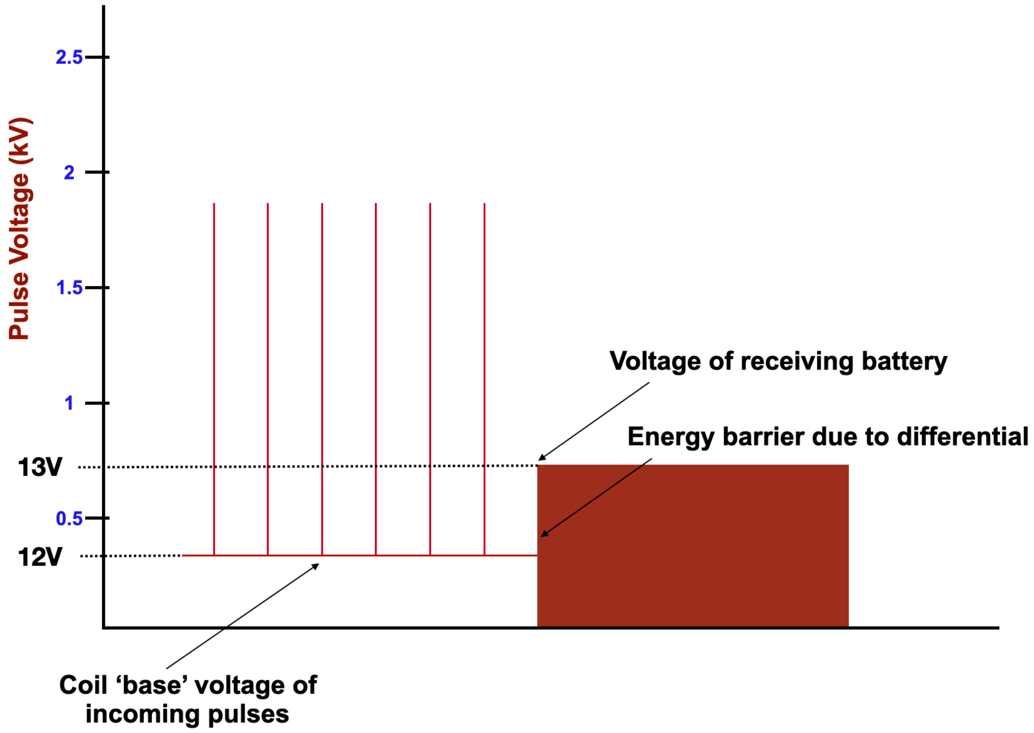

Fig 29: Spike and base coil voltage

unnecessarily and so increase the operator energy input and this will reduce the CoP value accordingly.

As with other factors, there is a 'sweet spot' where the best battery response is balanced with the minimum energy being delivered to the generator and which will show as a peak in a graph of the CoP, the dependent Y variable, against the independent X variable.

For the purposes of testing, I have used a power supply instead of a 'supply or run' battery. This allowed me to apply a consistent voltage to the coils to observe the effects. Of course in practice, when using the normal two battery system with the swapper, then the supply battery voltage may drop significantly during the swap interval depending on the battery type and condition. In order to enable adjustment of the coil voltage, to compensate either way, the option of using a Buck and a Boost converter is built into the PCB. The connectors for these are indicated in the terminals list on page 12. In this way you can choose to use the same supply for both the coils and the circuit or independent supplies for both, with one being higher than the other or vice versa.

Whether you use these options will depend on the voltage drop due to the total load applied to the supply battery and if you are inclined towards tinkering to find the optimum value for the battery type and configuration you are using, then it is a useful addition. The optimum coil voltage I found to vary between 12 and 13V depending on the battery type and capacity but trying to investigate the exact reasons why would not necessarily be a good use of time.

On the revised PCB there various supply options to the coils and circuit components has been designed so that you can choose the following arrangements shown in Table 8.

For example, if you want the same battery supply to be used for both the circuit components and the coils then you would place jumper headers on J4(H24) and J5(H27). If you decided instead to use two batteries in series, at around 24V, but still wanted the circuit to use 12V, to avoid having to utilise higher rated components, then use J1(H23), J3(H29) and J5(H27).

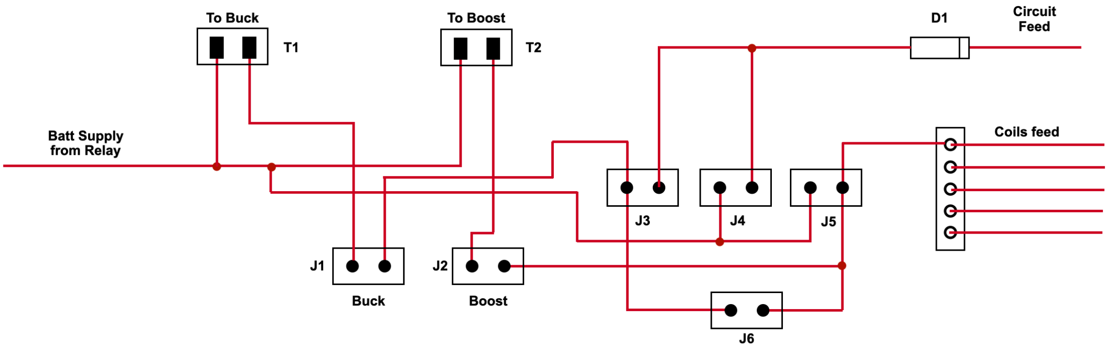

Fig 30: Circuit and coil supply network

This network is laid out in Fig 30 and the current pathways involved in each of the four options is shown in Fig 31. Option 1 is the only one that does not involve using either a Buck or a Boost converter. If you have those installed then their outputs arrive on the board at H12-2 and H18-2 respectively. In the pathway diagrams these are shown as terminals T1 and T2.

Option	***Supply wanted***	***Jumpers On***	***Comments***
1	Same battery supply to both circuit and coils	4 & 5	Coil voltage from supply battery adequate, even with load changes, and supply within circuit component tolerance (typically 6-18V)
2	Battery supply higher for coils and reduced supply for circuit	1, 3 & 5	For example, if using several batteries in series, therefore 24/36V to coils, but need 12V for standard circuit components
3	Battery supply adjusted down for coils with circuit voltage as battery	1, 4 & 6	For example, the single battery voltage is too high for optimum coil supply and needs reducing to an appropriate level
4	Battery supply adjusted up for coils with circuit voltage as battery	2 & 4	For example, the supply battery has too low a voltage, due to discharge and load, and needs to be raised for an optimum coil voltage.

Table 8: Circuit and coil supply options

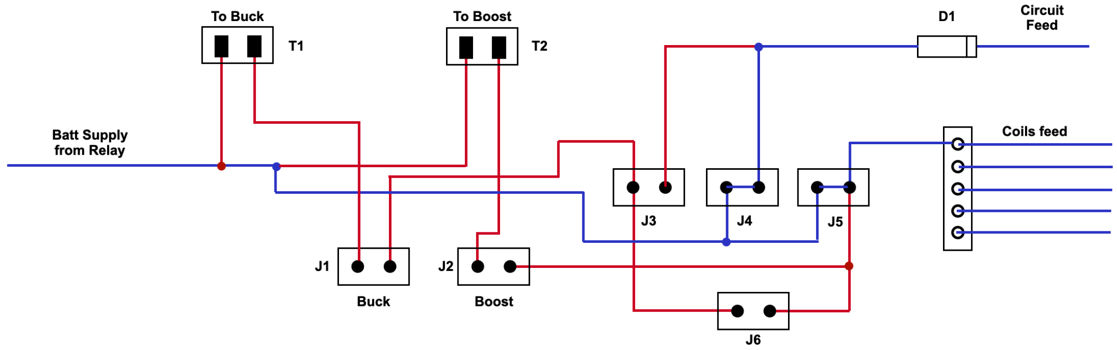

Routing for option 1

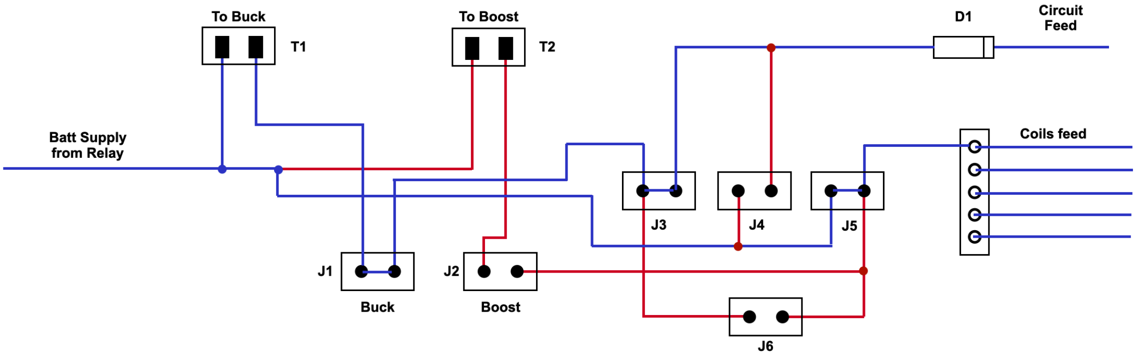

Routing for option 2

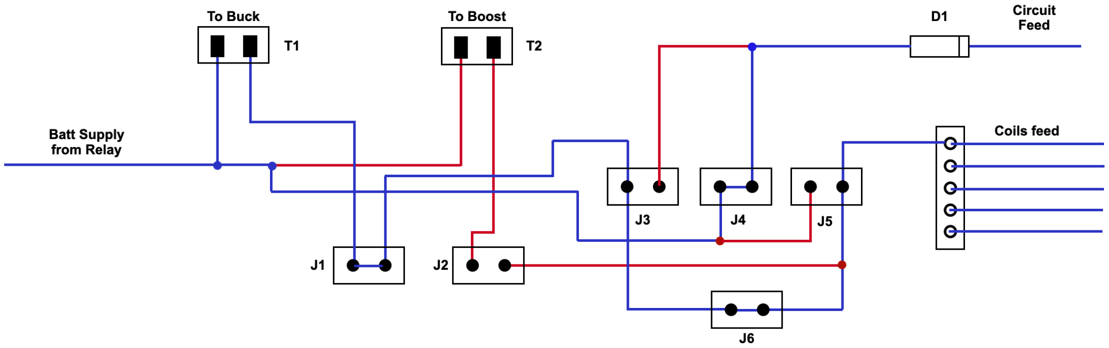

Routing for option 3

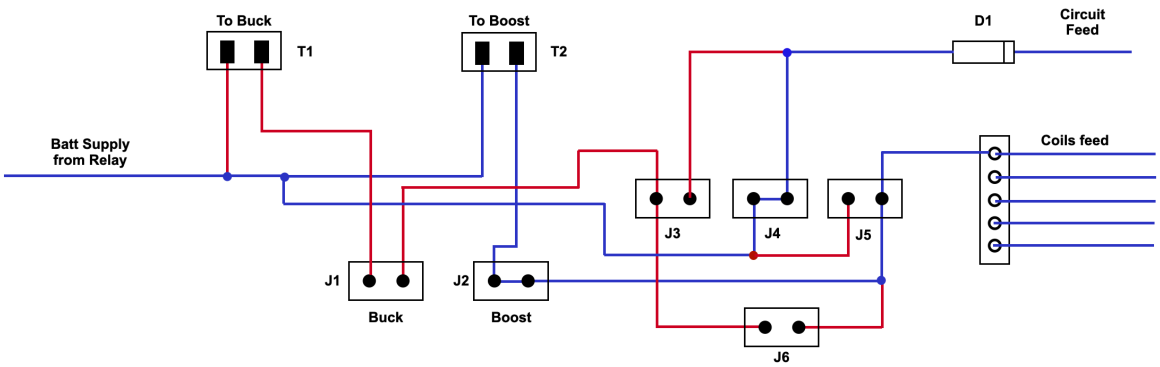

Routing for option 4

Fig 31: Circuit and coil supply pathways

Adjusting the voltage output from either of the converters just involves using a small screwdriver to adjust one of the blue trimmers on each device. Using a meter on the outputs H-15 and/or H-20 will tell you what is happening with the outputs.

The other trimmer that is usually installed on them sets the maximum current for the converter up to its limit. The devices do not need to deliver more than a few amps and so are not working at the limit of their typical capacity of around 5A for this small type of unit.

POTENTIAL DIVIDER

A potential divider is a very useful device, though not an essential one, that allows one to view the HV pulses on a scope. This is of use to determine the peak voltage and to see if there is any bounce or other artefacts in the interval between the pulses due to effects from certain components that may not be operating as per their specification.

Since most scopes will typically have a peak input voltage of around 500V, what is required is to reduce the spike voltage to a manageable level using the standard potential or resistance divider network.

In this way, using a 10:1 ratio, a 1,500V spike will show as 150V which is well within the specification of most scopes. Ideally it will need to be calibrated using a signal generator on impulse setting, where a 5V impulse/spike signal will show as a 0.5V amplitude signal when adjusted using the potentiometer.

For those who are unfamiliar with the principles of a voltage divider, with reference to Fig 32, the full spike voltage is presented across the 'input' and the output to the scope is taken off a variable resistor at the lower (right hand) end of the divider. If the full voltage is present, with respect to Ground, at the left hand end of the string of resistors, and the right hand end is at zero volts (Ground), then the voltage measured at various points along the resistor chain will be decreasing as we move from left to right.

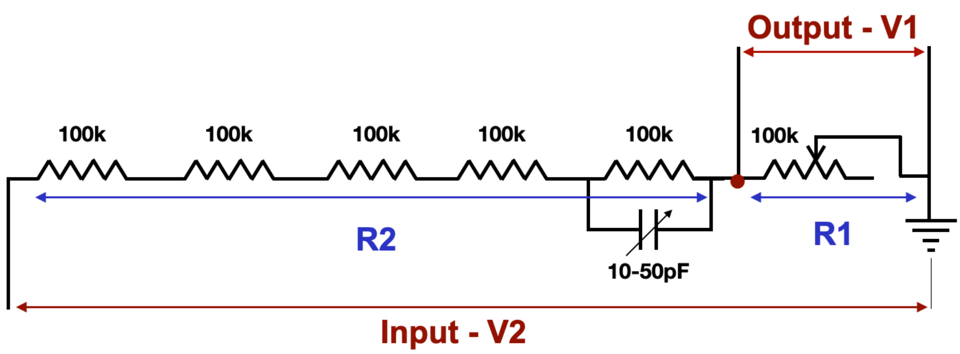

$$\frac{V_1}{V_2} = \frac{R1}{(R1 + R2)} \qquad \text{e.g.} \qquad \frac{56k\Omega}{(500k\Omega + 56k\Omega)} = \frac{181V}{1800V}$$

Fig 32: Potential divider circuit

The proportion of the total voltage we tap off at the position of the red dot is given by the equation shown and where a potentiometer allows adjustment of the proportion. In the example given the pot is set to 56k and so the proportion of 56k to the total (500k + 56k) is the ratio that is applied to the voltages, in this case 1 to 9.93.

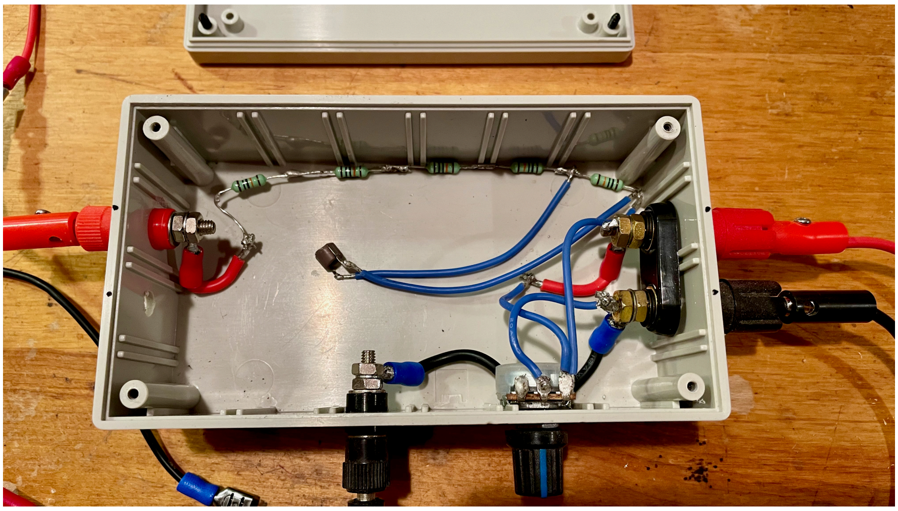

Fig 33: Potential divider circuit build

In this example, once calibrated, the scope will read 181V and so knowing the ratio of the setup, 1800/181 = 9.94, then we will know the real voltage that is being inputted to the divider when our scope reads 181V, or indeed any other voltage displayed.

For this to be accurate we need first to have calibrated the divider using a known voltage so that we can determine the working ratio. This can be done using a signal generator set to ‘impulse’ mode whereby a sharp spike pulse of 5V is outputted. Inputting this to the divider we can then adjust the potentiometer so that the output on the scope reads 0.5V, therefore setting up the divider to read with a ratio of 10:1. It doesn’t matter what the actual resistance is that you have set with the potentiometer in making the adjustment, the ratio is what is important in order determine the real voltage input. If your adjustments are such that your ratio happens to be 9.8, then that’s fine so long as every once in a while you check it when doing multiple measurements. With this particular ratio, if your scope reads a max pulse height of 114V then you know that the actual peak voltage is 114 x 9.8 = 1,117V.

The small capacitance also present in the divider circuit is to match that presented by the scope. In practice all you need to do is to adjust it until a 5V square wave input has nice sharp corners and not curved ones due to stray capacitance. After setting it up it should unlikely need to be adjusted again. Fig 34 shows the scope trace for 1,040V pulses using the divider set to 10:1.

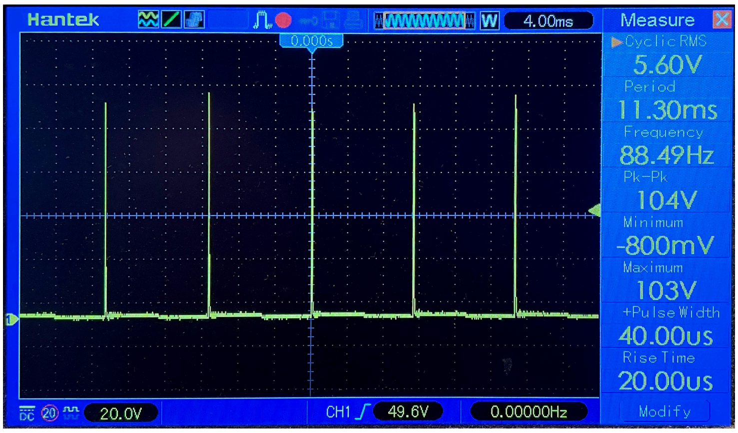

Fig 34: HV pulses with a 10:1 potential divider

While this divider may not be as accurate as required in certain industrial electronic applications, it is perfectly adequate in this context and gives us a good idea of the peak voltage we are generating. Having said that, the actual voltage figure is not that important and, as mentioned, it is the consistency and shape and the presence of any distortions or missed spikes that is of equal significance in viewing the pulse profile. However, knowing that the active device is working properly is also of value and the peak voltage figure is useful as a point of reference in a spreadsheet or report.

SWITCH ON & DIAGNOSTIC CHECKS

If you have reached this stage and come to that tense moment where you flick the 'on' switch and you see lights and signs of activity, then you should congratulate yourself and break open a bottle. If on the other hand there is a resounding silence then don't be too despondent for this is often the case. Despite the best intentions, and plenty of due diligence, any complex system being constructed is very likely to have a one or more small issues that will prevent it from operating as expected. With every component having to do its job, and every connection between them needing to be solid and as designed, it is hardly surprising that something or other will not be quite as it should. So here is the sequence of events that should mean all is well and working correctly. If not then follow the diagnostics below to identify any problem areas.

After wiring up all the terminals to the external components and connecting both batteries, or one battery and a switched on power supply, then flick the main switch. Assuming the swapper (SW2) and the trigger input are off, what you should see is one or other of the two green LEDs beneath the main relay light up. If that occurs you will then need to see the other battery LED come one, to show that the other battery supply connections are ok. For this the battery swapper system and relay (RLY1) needs to be activated.

To do this turn first set switch SW1 off (up), so that the T1 swap interval is selected, and then adjust the small R6 trimmer screw fully anti-clockwise till it clicks to its minimum resistance. This will ensure that the swap interval is at the minimum value of approximately 15sec and which is a useful value for T1 as it aids setting the relay to a chosen position. Due to the function of the two smaller relays with the chip and additional components, there is no straightforward means to set it manually and for it to maintain its position. (Please also note that if LED2 is lit to start with, then the first interval after the swapper system is activated will be twice the set value due to the chip output having to go high and then back to low for the relay to flip. This is only for the first event.)

Turn on the SW2 switch and LED3 (Swap LED) should light up showing that power is being delivered to the 4060 chip. Wait for about 30 sec for the relay to switch over to the other battery. If you are switching from LED1 to 2 then it will be the typical 15 sec. As soon as that happens, and within half a second, switch off SW2 and the relay will remain in that state until you choose to start the swapper again. If you wait longer to turn off SW2 then the relay will usually 'bounce' back to the other position, in which case just repeat the steps above.

For all my CoP tests, the swapper was turned off and B1 was lit to keep that battery (or PSU) as the supply while B2 was the receiving, pulse charged battery.

If all has gone well to this point then you should know that the battery supply connections are good and that the swapper system also works. You can adjust the other swapper interval (T2) later when needed using the descriptions in the 'Battery Swapper' section.

So whether you are experiencing that sound of ‘nothing’ or are in raptures of delight but you would still like to check that every stage is working as it should be, what follows is a detailed step by step set of diagnostics that you can logically work through. Careful logging of your results is very helpful when pondering over the readings to try and logically deduce where a problem might be occurring.

Test 1: Battery Supply & Swapper

To check: That voltage is reaching the PCB from the batteries via the relays to power the circuit and coils

Procedure: Start by putting a multimeter on H9-1 and then H9-2 together with any nearby Ground point (e.g. H6, H8, H10, H14 or H21) to check that the battery voltage is reaching the PCB via the input terminals H1 and H5. If not then check the wiring and connectors to the batteries.

Check the voltage on H7 (Out +) to see if the large relay is passing on the same voltage from either Battery 1 or 2. If not then use the schematic to check the connections are as indicated around the large relay (RLY1).
The two smaller relays deal with the HV pulse routing (RLY2) and the LED indicators (RLY3). They are wired so that when LED 1 is lit, indicating that Batt 1 is supplying the power to the large relay and onward to the circuit and coils, the HV pulses are being routed to Batt 2, and vice versa. In other words, whichever battery is providing power to the system, it is the other that is being pulsed charged and neither battery is ever supplying power at the same time as it is being charged. The batteries are not able to deliver energy, via normal electrochemical pathways, while at the same time having to ‘process’ incoming pulses whatever mechanism is occurring.

It is the upper of the two small relays (RLY2) whose output, marked ‘To Relay’ on the schematic, trips the larger relay, causing the internal contacts to switch over with a maximum rating of 16A which should be adequate for the level of external power that the generator might provide.

The timing of the swap interval is set using the two trimmers R6 and R7 as described in the ‘Battery Swapper’ section and above. After setting T1 to about 15s then, with the SW1 switch (Swap T1/T2) up, RLY1 and the swap LEDs should change at the preset swap interval. Looking inside the large relay from the top, the central movable contact should be towards its internal coil and the battery LEDs when LED1 is lit, and away from its coil when LED2 is lit.

To clarify the various diagnostic stages, the details are summarised in a table for each set of checks.

Terminals/ Components used	Function	Activity
H9-1 & H21 (Gnd)	Displays Batt 1 voltage	Checks unimpeded connection from battery 1 to board
H9-2 & H21	Displays Batt 2 voltage	Checks unimpeded connection from battery 2 to board
H7 & H21	Displays supply voltage	Checks unimpeded connection from one or other battery to power output terminal via large relay RLY1
SW2/1	Swap 'off/on' and T1/T2	Checks operation of timer and indicator LEDs based on R6/7 trimmer settings

Table 9: Diagnostics checks for the supply

Test 2: PWM trigger signal and Device driver

To check: that the square wave signal input is reaching the MOSFET/IGBT Gate.

Procedure: whether you are using a PWM module to produce your square wave input or a Hall sensor or other device, the input to the board is via H37-1. From there it goes to the base of a transistor Q2 and the signal can be scoped at its Collector at TP-H13 or at the input to the Driver (U2) at H16.

When the power to U2 is on, that is with switch SW5 (FET Driver Off/On) down, then an output should be measurable at TP-H17 (Gate In). This Driver output signal is then routed to the Gate/Base of the FET/IGBT depending on which of the Jumper headers for H28 or H32 is on. Using these test point (TPs) the signal can be checked from the PWM module output right up to the Gate/Base of the active device.

Terminals/ Components used	Function	Activity
H17 & H10 (Gnd)	Input trigger signal	Checks arrival of square wave from PWM/Hall sensor or other source via H37-1
H16 & H10	Input trigger signal	Checks arrival of square wave at input to FET Driver U2
H22 & H10	Input trigger signal	Checks arrival of trigger signal to Gate/Base of active device depending on H28/H32 jumper selected
SW5-3	Power to Driver chip	Checks circuit voltage is being delivered to input of Driver chip U2-1 when SW5 is on and LED4 lit showing driver output to active device

Table 10: Diagnostics checks for the trigger signal

With the the FET driver switch down, then HV pulses should be produced accompanied by a faint buzzing sound depending upon the frequency (PRF). If the sound is louder and harsher then you probably have the HV Load switch off which is designed to divert the pulses for HV measurement with a potential divider.

Test 3: Coil Supply

To check: the supply to the coils using the various available options.

Procedure: As described in the 'Coil Voltage' section, the voltage applied to the coils can be the same as the supply battery, or optionally raised or lowered to suit different needs using either an optional Buck or Boost converter. Besides ensuring that a suitable voltage reaches the coils, it is also done to maintain a suitable voltage to the main board components whose working range is typically in the 10-18V and which will not cope with an elevated coil supply if using several batteries in series. The small losses due to using either of these converters is more than compensated for by the optimisation of performance.

One can check the voltage coming back into the board, in other words the converters' outputs, from either device on H15 or H20. Depending upon the jumper selection you are using for the coil and circuit supply (see 'Coil Voltage' section), you can also check the circuit provision at the anode of diode D1 and at the coils using any of the terminal screws on either H35 or H38. While doing that you can adjust the Buck or Boost converter output to the value required.

Terminals/ Components used	Function	Activity
H15 / H20 & H21 (Gnd)	Voltage from Buck and Boost converters	Checks the output voltage from the Buck and Boost converters, if used, and when powered via H12-1 and H18-1
Anode D1 or H28-2	Resulting supply to circuit	Supply to main board based on J1-J6 jumper selection
H33/H36 & H35-4	Resulting supply to coils	Supply to coils based on J1-J6 jumper selection

Table 11: Diagnostics checks for the coil and circuit supply

The values that you will read are no-load values and in operation they will naturally drop by 0.5-1.0V when under load but, so long as you are consistent with your figures that will not matter, but it is important in recording voltages to indicate whether they are under load or 'no load'. This detail may be required later in calculating total energy supplied, or perhaps to identify normal I^2R losses that are part of any circuit.

Test 4: HV and load switch

To check: that the HV pulses are as expected and to measure their peak value.

Procedure: When HV pulses are being produced, or more exactly when there is an output from the Driver chip (U2), then the blue LED4 will light in a continuous fashion at a rate equal to the PRF.

If the HV load switch (SW4) is down then the pulses are directed to the battery that is not supplying the load via the small relay (RLY2). If the switch is up then the pulses are directed to terminals H14 (Ground) and H19 (+) where you can connect a potential divider.

If the HV pulses are not being absorbed into the receiving battery, they will take the next easiest path which will be to the potential divider. If one is not connected and the generator is operated then the sound will change to a harsh 'buzzy' sound reflecting the fact the the pulses are not being absorbed and are trying to discharge from any suitable surface and find a way to ground or to the air.

Terminals/ components used	Function	Activity
SW4	Directs HV pulses	When on the HV pulses are directed to one or other battery via swapper
H19 & H14	HV pulse measurement	Connect to potential divider for real time observation of HV pulses

Table 12: Diagnostics checks for the load switch

Switching SW4 on will therefore select the pulses going to the battery and will normally be left down and only set to 'off' when you are taking HV test measurements with a potential divider or for diverting the pulses away from the receiving battery to a set of capacitors, for example, so that you can conduct a control experiment and which will be addressed later.

Test 5: Meters

To check: that the various meters are operating correctly.

Procedure: Although when using a PWM module the frequency is shown on its own display, it is of use to have a separate frequency meter, especially when you are using a rotor input such as from a Hall sensor.

Early work using a 555 timer chip for the trigger pulses showed that noticeable temperature drift was occurring, even over a 10min run due to the warming of the RC components. The PWM module is far more stable and also, the presence of the PRF meter will also serve to remind you if you have not 'soft pressed' the module controls after a frequency adjustment to activate the module's output. The readings are normally within ±

1Hz of each other and such precision is important since the CoP values can be affected by a small shift in PRF of less than several Hertz.

With SW3 on, the + and - supply to the frequency meter is via H31-1 & H31-4 respectively and the signal feed, from whatever trigger pulse source you are using, goes out via H31-2 (note that H31-3 is not used).

Terminals/ components used	Function	Activity
H31-1 & H31-4 (Gnd)	Supply to frequency meter	Checks voltage supply for PRF meter
H9-1 & H21 (Gnd)	Voltage for panel meter 1	Checks voltage supply for left hand panel meter for battery 1
H9-2 & H21 (Gnd)	Voltage for panel meter 2	Checks voltage supply for left hand panel meter for battery 2

Table 13: Diagnostics checks for the meters

The small panel meters serve to give a quick, but not very accurate, indication that one battery is under load and the other is charging, as indicated by a voltage drop and rise respectively. They are arranged so that from the front the left one indicates battery 1 and which, in my own case, aligns with the left hand battery when looking from the same direction. The panel meters are part of the main negative line between the batteries and the board and are also in line with the fuse block and main switch.

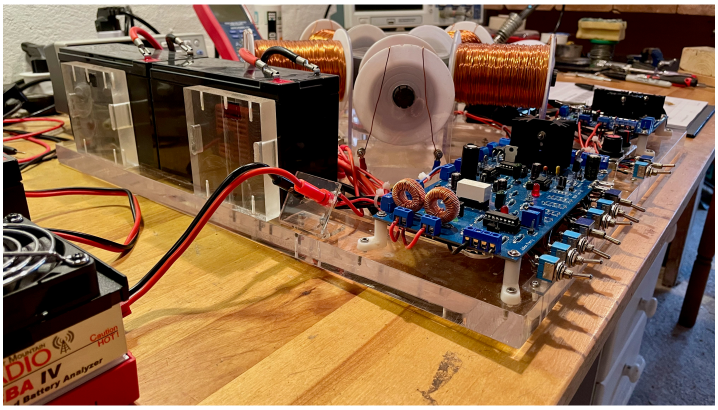

Fig 35: Connection to CBA from Batt 2

The voltages that they read are supplied directly from the battery inputs to the board and via H9-1 (battery 1) and H9-2 (battery 2). If you are going down the route of doing accurate CoP and power tests, then you will need to use some form of electronic load. For testing purposes battery 2 was dedicated to being the one receiving pulses, and therefore undergoing repeated cycles of discharge and pulse charging, and battery 1, or a power supply to replace it, was the supply or 'run' battery.

As such the swapper is turned off and therefore the Batt 1 LED1 should be on. Fig 35 shows an earlier (v2) PCB but the connection to the CBA device is the same and comes off the battery 2 positive terminal and from the Ground line on to a small perspex plug panel. Do not use an output from H7 (Out +) since that is powered by the supply source which, in this context is a PSU, serving as battery 1 and so will not show what is happening to the battery being pulse charged.

OTHER COMPONENTS

Besides the components placed on the PCB, there are various others that make up the whole device, some of which have been mentioned in the diagnostics section and elsewhere. These are listed in the Table 14 below and with a suggested UK source that should be able to be matched from wherever you are based. With most of them you can use a suitable alternative.

Component	Function	Source
Frequency meter	Displays PRF from various inputs	https://www.aliexpress.com/item/1005003066921539.html?
Panel meters	Displays live voltage of batteries 1 & 2	https://www.ebay.co.uk/itm/154574694810?
PWM module	Delivers accurate adjustable square wave trigger pulses	https://www.ebay.co.uk/itm/203966572841?
Main switch	An example of many possible varieties	https://www.ebay.co.uk/itm/123975254037?
SPST Switches	For the various switchable features	https://www.ebay.co.uk/itm/392442247174?
Wire	Silicone coated AWG18 and 22	https://www.ebay.co.uk/itm/331718921763?
Coil Spools	Flange OD~82mm, Hole ID ~19mm, length ~84mm	https://www.aliexpress.com/item/4001164937002.html?
Copper wire	0.71mm Enamelled Copper Wire on D160 Reel (4kg) Ref: Ref: SX0710D-4KG-D160	https://www.scientificwire.com/acatalog/ecwire-solderable.html#aSX0710D_2d4KG_2dD160
Ferrite rods	Mn-Zn 18mm Dia x 100mm Length	https://www.aliexpress.com/item/1005004001417754.html?
Fuse block	5A circuit protection using just one of the tabs	https://www.ebay.co.uk/itm/143706726818?
Battery (SLA)	12V-7Ah AGM Sealed Lead Acid battery	https://www.tayna.co.uk/mobility-batteries/dc-battery/dc12-7-0s/
Battery ($LiFePO_4$)	12V-7Ah with dedicated charger	https://www.ebay.co.uk/itm/124067183749?
Battery ($LiFePO_4$)	12V-18Ah with dedicated charger	https://www.ebay.co.uk/itm/113564191058?

Table 14: Other components and sources

WINDING THE COILS

The coils are the main source of the HV pulses on account of Faraday induction in conjunction with Lenz's Law, which is an electromagnetic version of the 1st Law of Thermodynamics (Energy Conservation).

As indicated previously, producing large chunky coils is not a requirement to observe CoP>1 since even quite modest coils can produce a healthy HV pulse on paper. However, as examined in depth earlier, this pulse will be clipped by the 'avalanche rating' of the active device used, an inevitable consequence of the device's role in switching the solenoids on and off.

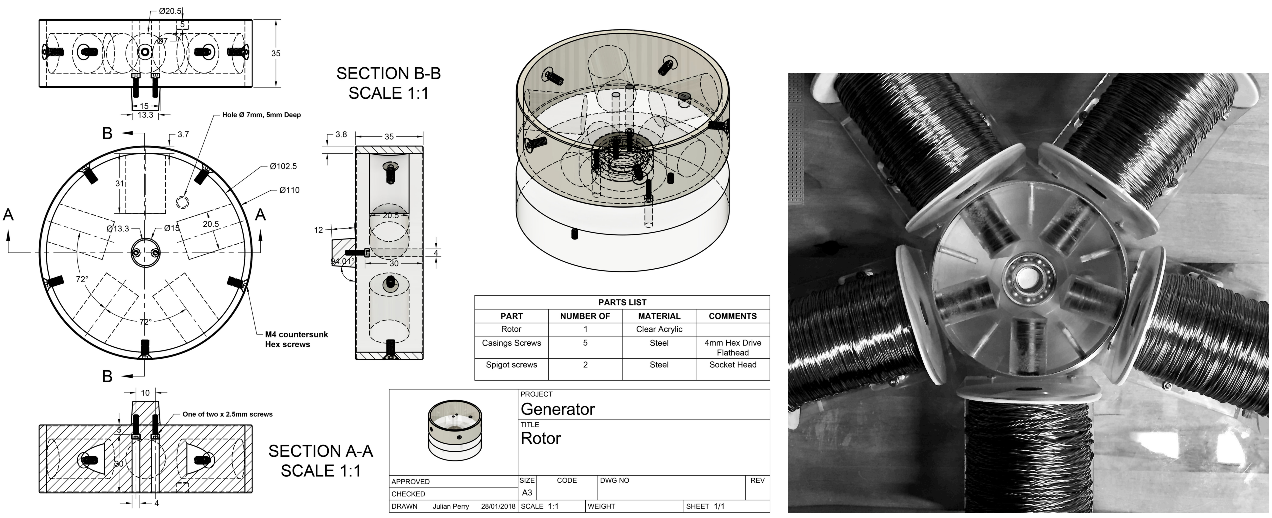

Fig 36: Rotor design and alignment

Interestingly, measurements done with the rotor removed have resulted in CoP values about 10% lower than with the rotor in place. This is assumed to be due to the enforcement of the magnetic fields in the solenoids by the rotor magnets which are sited little more than 10mm from the solenoid ends. As one of the original design drawings in Fig 36 shows, the rotor magnets are aligned precisely with the solenoid cores and this undoubtedly bolsters the magnetic field with a knock on effect on the field strength and HV pulses. However, as explained earlier, the rotor system is limited due to the fixed PRF that results and therefore, even though the magnetic field is a little weaker, the CoP results are far superior with a flexible trigger system that can be set to accommodate the specific characteristics of your battery and device.

For winding the coils, as can be seen from the 'other components' list in Table 14, the recommendation is to use 0.71mm enabled wire and, using some plastic spools, wind it on until the spool is almost full. There is no specific number of turns required but in my case each coil was approximately 2,600 turns and each took around 10 mins to wind by hand.

While you don't need a winding device with a counter, it can be of interest to know how many turns you have done. Of more importance is the need for some kind of spool carrier that will allow you to draw from a large spool of wire unimpeded. With 4kg of wire, this should be enough for 5-6 coils, and where the large spool is kept in position on the shaft with a couple of clips or a pair of mole grips on either side of the spool.

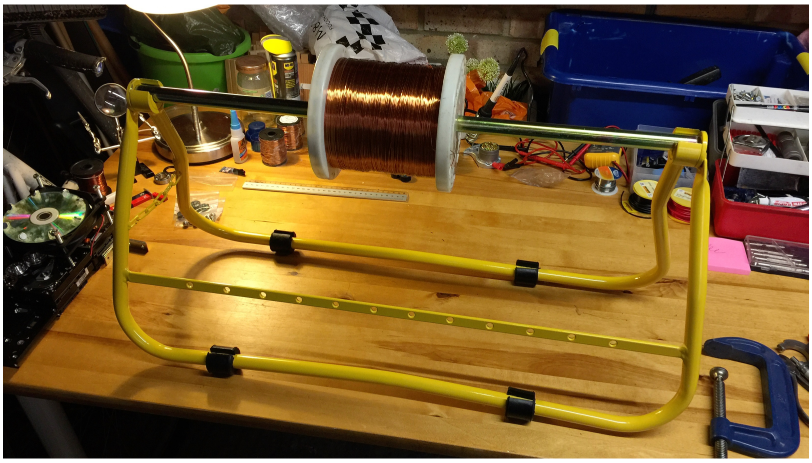

Fig 37: Spool holder

The flanges of the plastic coil spools are inclined to flex outward a little as you approach the end of the winding and, in my case, due to the need to have an accurate and stable alignment with the central rotor, I built perspex coil holders. This is not required if you are doing away with the rotor and you could even use a glue gun to fix each of the coils to a removable or non-removable plinth. Fig 37 shows the winding in progress and be sure at the start to secure about 15cm of wire over the flange so you can make the electrical connections later. Making a small notch in the flange is a good way to prevent it from moving around.

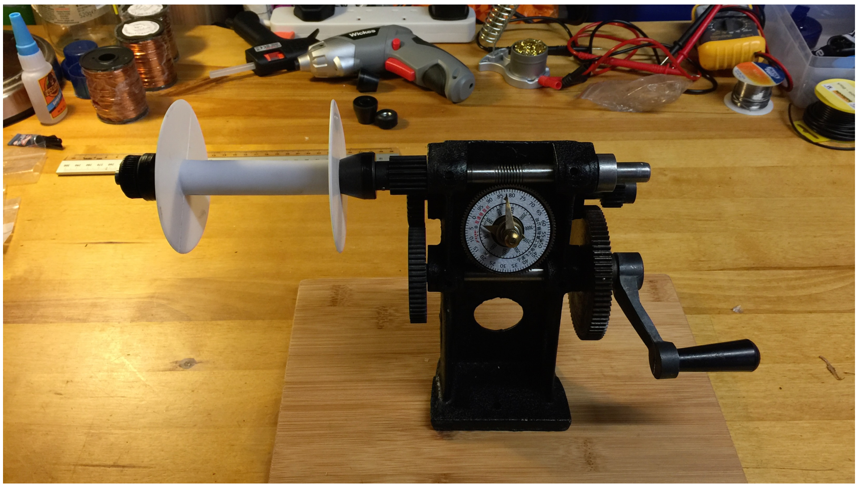

Fig 38: Manual coil winder

You should aim to finish the last layer of winding with the wire at the same end as you started with. You will also find that, even if you start the winding with neatly laid out coils, as in the photo, by the time you have completed a few layers the uppermost layer they will be 'all over the place' as it becomes increasingly difficult to lay each wire next next the previous turn in a smooth fashion. Despite the best laid plans you will end up with strands crossing over each other, however, this will not affect the coil's performance or behaviour. They will then look like in Fig 40 with two loose, but secured, wires at one end for the connections.

When it comes to wiring them up, you will need to scrape the enamel off about 4cm of the wire to ensure a good contact with whatever terminal you chose to use. I used 3mm studs onto which I could fit small washers and nuts as shown in Fig 41.

The ferrite rods should be inserted into the inner diameters and they can be secured with a little glue at either end to prevent movement and to accommodate the fact that they are not an interference fit.

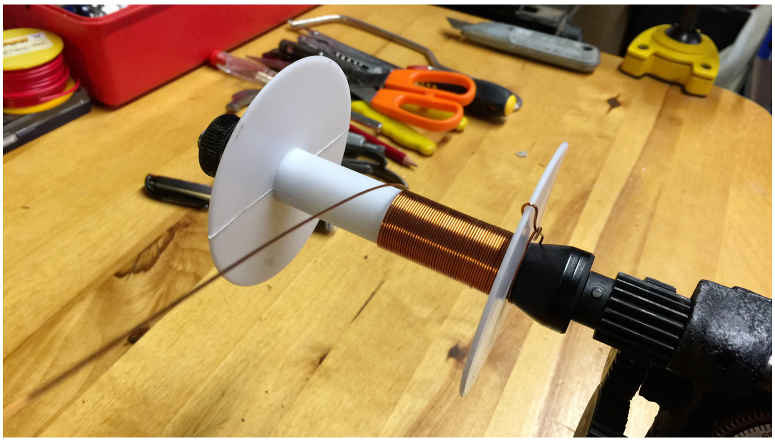

Fig 39: Coil winding

Early in 2018, one of my coils was wound using litzed wire, as can be seen in Fig 41, which is a way of twisting together 3 strands onto one coil. The aim here was to construct a trigger coil that would detect the presence of the approaching rotor magnets and serve to produce the trigger pulses. Besides the fact that I could never get it to work properly, producing a trigger pulse in this way would has the same limitations as using a Hall sensor, namely that you are restricted to the PRF that naturally results from the rpm of the rotor at any moment. This coil has since been unwound and replaced with a single stranded one, the same as the other four.

Fig 40: Wound coils

With one of the multi purpose meters shown in the equipment list (Fig 42), you should expect to measure the inductance of an individual coil at between 350 and 400mH and with a resistance of 10 - 15Ω. When wired in parallel then this will drop to 20 - 25mH and 1 - 2Ω. Given those values it is reasonable to wonder if at 12V whether about 12A would be flowing in the parallel connected coils? The answer is that the square wave trigger pulses, together with the resultant switching, substantially reduces the current that can flow and buildup within the coils. While the current will vary with the PRF and duty cycle settings on the PWM module, you can expect a current in the 0.5 - 1A range. However, if you choose to try two or three batteries in series than that will increase the current accordingly.

Regarding the number of coils to produce, counter intuitively, on doing some tests where I disconnected all five coils and then connected them back in one by one, for the first three there was the expected improvement in CoP. However, when four coils were connected the CoP dipped noticeably and, with five connected, returned to a more expected value.

Like just with every other variable in this generator, the individual properties of each build are specific and need to be optimised and yet, with regard to the number of coils five or six should be very adequate.

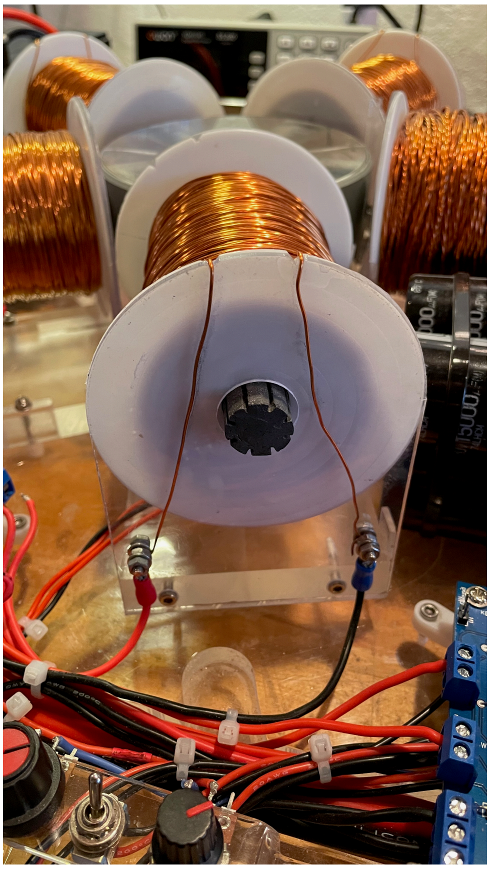

Fig 41: Coil connections

While the Vt graph of the HV spikes made with a divider (Fig 34) shows that particular aspect of the pulses, there are other aspects that are not revealed in such a scope trace. The pulses from each coil combine together in a way not easily quantifiable to create a release of charge within the receiving battery. While some of the charge is released by processes that are well understood, the CoP results clearly indicate that other processes are occurring, even though they are not yet well understood.

Such processes, involving the battery's electrochemistry, may involve what are referred to as 'Type B' energetic reactions (see 'Research papers' in the Appendices) that are subject to the 1st Law of Thermodynamics (Energy conservation) but not necessarily strictly confined by the 2nd Law (Entropy increase). However, this needs to be confirmed with detailed tests of the battery's State of Health (SOH) that will determine if the electrolyte itself is being consumed to provide some or all of the energy measured energy gain.

Using multiple coils does not change the peak voltage seen on the scope and yet using only one coil would substantially reduce the CoP values. Clearly there is more going on than just delivering a particular dV/dt to the electrolyte interface. Such a scope trace does not show the charge liberated by the pulses and the proportions coming via known and unknown sources and pathways.

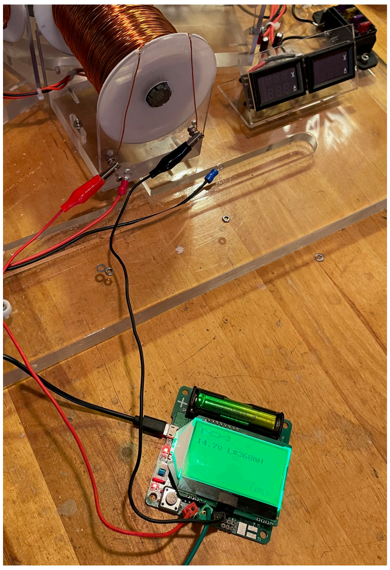

Fig 42: Measuring inductance

Generally speaking, more coils are better due to the combined effect of multiple coils in a way not directly observable on a Vt trace, even though, when they are combined in parallel, their inductance drops considerably, as with resistors in parallel. Apart from this, cost and functionality will be the over riding considerations and 5 - 6 coils should be very adequate.

TOOLS & EQUIPMENT

If you are planning use the generator with a set of suggested variables then some of the equipment described here is not required. You will likely experience values of CoP>1 but the degree to which you do that cannot be predicted. With that in mind a set a variables is offered in the Appendices, so that you will have a good chance with this particular build of receiving some good CoP results, without going down the route of extensive testing.

Equipment	Function	Source
Computerised Battery Analyser	To discharge precise amounts of energy through its electronic load and to record the voltage-time profile as the battery is pulsed charged.	http://www.westmountainradio.com/product_info.php?products_id=cba5
Recording multimeter	To automate the measurement of current supplied to the system	https://uk.banggood.com/Owon-XDM1041-USB-Digital-Multimeter-55000-Counts-High-Accuracy-Universal-Desktop-Multimeters-Meter-with-3_5-inch-TFT-LCD-Screen-p-1847470.html?
Digital multimeter	General purpose meter	https://uk.banggood.com/ANENG-AN8008-True-RMS-Wave-Output-Digital-Multimeter-AC-DC-Current-Volt-Resistance-Frequency-Capacitance-Test-p-1157985.html?cur_warehouse=CN&rmmds=search
Capacitance and induction meter	To check active components and measure capacitance and inductance	https://www.ebay.co.uk/itm/273536112922?
Power Supply	General purpose PSU to supply up to 60V @ 5A	https://www.aliexpress.com/item/1005005010926038.html?

Table 15: Equipment required

If on the other hand you are choosing to measure the CoP of the device as you make adjustments to the generator variables, then there are certain pieces of equipment that are essential to be able to do this with any degree of accuracy.

These are shown in Table 15 and with some suggested links, although you can find alternatives elsewhere.

The features of a computerised battery analyser, the CBA, with its electronic load are central to the CoP and power measurements. The West Mountain Radio CBA (model IV or V) was recommended to me and has performed admirably, even though on occasions the USB connection to the laptop was a little poor so that after a 70min ‘charge monitoring’ session, the full data was not transferred to the laptop reliably.

Fig 43: CBA device

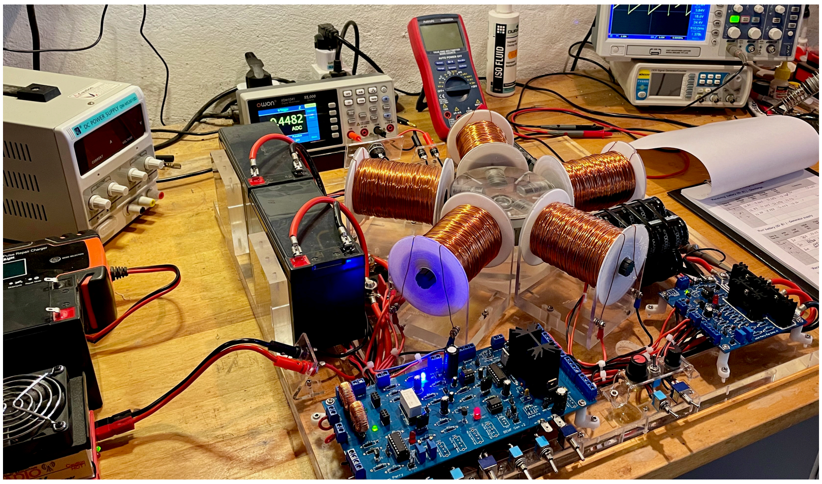

Fig 44: Connections to the CBA (v2 PCB)

This is more to do with my laptop socket than the plug provided. Taking screen grabs is highly recommended to provide some useful evidence generally, as a back up in this situation and for inclusion in reports.

If I had to make one recommendation for a modification to the CBA, it would be to change the input connections where currents of up to 10A are used to discharge the battery over an extended period. Unless your connections are solid, the device may sometimes come up with a note saying that the contacts might be dirty and the discharge trace shows as a little ‘noisy’. I removed the ‘Powerpole’ plugs on the pigtail lead and soldered the wires directly to the terminals on the analyser to ensure a good solid contact. While this didn’t stop the event entirely it was much less common.

Fig 44 shows the original ‘Powerpole’ connectors at the input to the CBA and my own banana plugs used with the output terminal that is connected to battery 2. In Fig 45 the ‘Powerpole’ connectors have been removed and the bare wire soldered directly onto the metal contacts and these were then shrink sleeved.

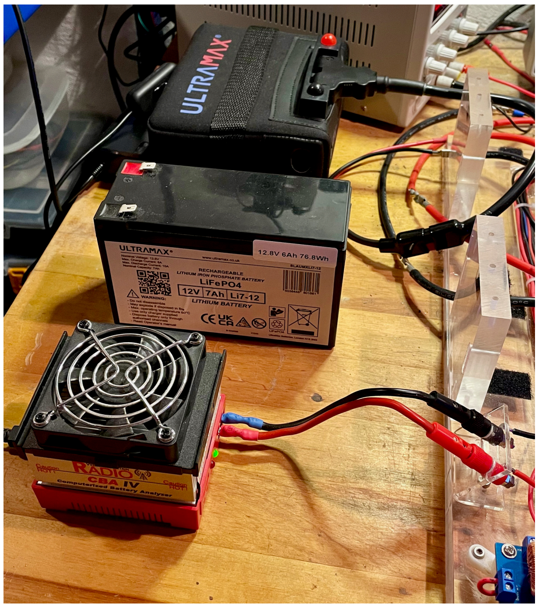

Fig 45: Modified connections

Measuring the supply current could theoretically be done manually but, unless you are especially good at multitasking and remembering to take evenly spaced readings over the duration of a test run over perhaps 10 - 20 mins, then it is much more convenient to have this done automatically.

Fig 46: Recording multimeter

With the Owon recording meter for example, you can set the sampling rate intervals from 15ms to over 2.75 hours and up to 1000 data points. The file is then exported as an XLS file for spreadsheet use and the calculation of an average in a suitable spreadsheet (one is offered in the

Appendices). The instruction and specification manual for this meter, and some other pieces of equipment, are also in the Appendices.

Most good multimeters will suffice and they are generally used to check a voltage at a point on the go and, perhaps more importantly, to use as a continuity meter in circuit diagnostics and during construction to ensure that connections are sound.

Since the majority of the voltage measurements used in the spreadsheet calculations are those taken from the CBA's 'live' readings, then the accuracy of those is more relevant in deriving the uncertainty values, as explained in the 'Uncertainty Analysis' document in the Appendices. There are also sample spreadsheets, one empty and ready for use and the other with some of my test data, that can be used in conjunction with the 'Spreadsheet Guidance Notes'.

BATTERIES

The batteries are the heart of the system in that without them, there would be no observed energy harvesting, even if you are producing substantial HV pulses.

Fig 47: 12V SLA Gel battery

A detailed understanding of what is happening at the pulse-electrode interface is not yet available, partly because the purpose of this project has been to determine if a real phenomena is occurring and not to examine the specific energetic pathways and processes. That aside, there are various factors that need to be considered when choosing and using a battery, either singly alongside a power supply, as can be done in testing, or when using two in the normal operation of the generator with the battery swapper engaged.

Much of the prior work done on 'Bedini' type devices has been done using Sealed Lead Acid gel batteries (SLAs) since they are easily available and relatively cheap. A useful and easily available one is the 12V-7Ah AGM (Absorbed Glass Mat) gel format battery as in Fig 47. Early work was done using these batteries, especially with the first build back in 2018, in getting the rotor to spin and generally observing an operational device. The same batteries were used in the first CoP tests and acted as a reference for other chemistry types and capacities tested later on.

I would recommend starting with one of these in getting the device to work, doing diagnostics and for some early tests to serve as a reference. Then it is appropriate to move on to a 7Ah Lithium Iron Phosphate ($LiFePO_4$) battery for comparison and even better results as in Fig 48. Bear in mind also that the batteries take a little while to give of their best and pulse charging of a new battery may not give quite as good results as one that has been pulse charged multiple times. While there have been suggestions that pulse charing can damage batteries, the evidence so far is that there is some loss of capacity after many hours of pulse charging but that need not interfere with the battery's function within its normal lifespan. Oxidation and other events at the electrodes are possibly accelerated by the pulses but the impact has yet to be measured in any quantitive fashion.

Fig 48: 12V Lithium Phosphate battery

The advantage of Lithium Phosphate batteries, despite being lighter but more expensive, is that they have shown themselves to be more rugged and consistent in their voltage profiles and the results have generally been better. They are also have a higher energy density (W/kg) and are able to deliver a higher stable current for longer which is helpful in the power tests where you are likely to want to use external loads drawing up to 10A or more for up to 30 mins at a time.

The small and light 7Ah $LiFePO_4$ battery, with an example source linked in Table 14 in the 'Other Components' section, can sustain a continuous discharge current of 7A for the typical duration of a swap interval, before handing over to the other Lithium battery.

Fig 49: 17Ah SLA battery

When it comes to the larger capacity batteries, for example, 17 and 18Ah batteries, then Fig 49 shows a 17Ah mobility battery and Fig 50 a $LiFePO_4$ one used used in golf carts and able to deliver a continuous 25A without any damaging effects. However, bear in mind the maximum relay capacity for any external loads used.

It all depends on your needs at the time, but from the large number of tests that have been undertaken over the best part of a year, overall the Lithium batteries have shown the best performance, although data on the actual power that they would deliver, based on the CoP results, has yet to be acquired.

So, while you are unlikely to need this type of battery early on in your build, they are experimentally very good, although quite expensive.

Fig 50: 12V Lithium Phosphate battery (18Ah) with T bar connector and Torberry lead

METHODOLOGY

If you are going down the route of testing the generator to derive CoP values, or even straight into power tests, you will need a methodology that will stand up to scrutiny, even if it's only to your own high standards.

The main value in working with a clear and repeatable method is that it becomes automatic over time and you are less likely to deviate from a prescribed process and generate results that are due to changes in the method rather than changes to the performance of the device.

For this purpose, it is helpful to lay out a method is a series of steps and, as with many other aspects of life, to create a flow of 'thought, word and deed'. In this way your rationale for how you are doing things is developed as a thought process and then crystallised or precipitated as words, or diagrams, on a page before finding expression as actions that are predictable and repeatable.

The methodology that I used for all my CoP tests was developed from a suggestion made to me and which I then fleshed out and developing further with some useful form work to show the various steps involved in a single test. By seeing the steps on a page or screen you will find it a lot easier to identify where problems might arise or where the workflow is not conducive to a smooth experimental process.

Of course, you are in no way limited to using the methods that I used but it has provided clear and repeatable results that has allowed me to tune the various parameters so that I can clearly see the outcomes of small adjustments made. As the number of tests builds up over time, having confidence in the method you used will become significant in the evaluation of your results.

I will start by presenting again (for those who have not read certain of my other documents) the rationale behind the method used.

- - - - - - - - - - -

CoP Tests

Testing the performance of a generator such as this requires a series of very accurate and precise readings of the energy state of both batteries used in the system. The battery providing power to the circuits is referred to as the 'Run' battery and the battery being charged by the whole unit is referred to as the 'Receiving' battery.

There are certain quantities that we can measure accurately in a straightforward way and there are some we cannot. Based on the hypothesis being tested, expressed simply as: '*That the generator is able to extract a quantity of energy from the environment resulting in a CoP>1*', then we cannot directly measure the total amount of energy that the receiving battery is absorbing for one main reason. Even if the hypothesis is shown to be true, we have no way of knowing the proportion of energy that is being drawn in from the environment compared to that being provided solely from the generator device. This state of affairs is represented by Fig 51.

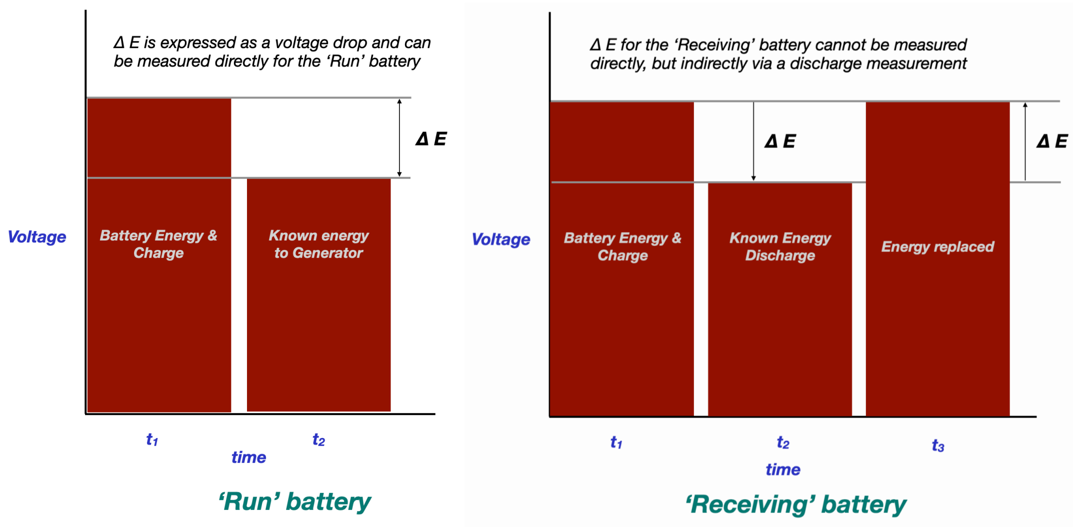

Fig 51: Principles of CoP testing

Despite this, we are able to determine very accurately the energy differential between two states of the receiving battery, as indicated by its open circuit (no load) terminal voltage. In other words, when fully charged, it will display a peak voltage and when it is at a lower state of energy and charge, it will display a lower voltage potential.

However, in practical terms, measuring voltage alone is unreliable to determine the energy differential so this is more accurately achieved by removing a precise and measurable amount of energy from the battery using an electronic load. In this way we can observe and record the total energy dissipated and lost as heat and the resulting drop in the open circuit terminal voltage. This process can be done using a Computerised Battery Analyser (CBA) and where perhaps the only device capable of doing this, along with the other required test explained below, is the West Mountain Radio CBA series of analysers.

For these tests the CBA IV will be used that will allow for a controlled and monitored energy discharge as well as record the battery voltage, graphically presented as a *Vt* profile, while the battery is being charged by the generator.

Once a precise amount of energy has been expended from the initially fully charged battery, then we know that, in order to return it to its original peak voltage and fully charged state, it will require the input of the same amount of energy. The small amount of energy dissipated as a result of the internal resistance of the battery will be assumed to apply equally during both charging and discharging. Any variations arising from the different characteristics of continuous and pulsed currents can be integrated into the associated uncertainties.

The other quantity that we can measure directly is the energy delivered to the generator by the so called ‘Run’ battery. This can be calculated by means of measuring both the input voltage, the average current delivered by the battery and the time the device is running in order for the ‘Receiving’ battery to reach full charge.

The voltage can be measured at the start and end of the test and a simple average arrived at. The average current is best determined from a series of regular, automatic measurements over the duration and a mean value derived from the data set.

The time can be retrospectively determined from the charging profile where the point of peak voltage, or where the voltage plateau starts, can be easily seen from the graph. Even if the generator is run for longer than that time, the energy calculations can be retrospectively made with the correct value of time.

So we can know the energy delivered by the ‘Run’ battery to the generator as:

$E_{(Supplied)} = V_{(av)} \,.\, I_{(av)} \,.\, t_{(Vpk)}$ J - *Equation 1*

We also know the amount of energy that is required to be delivered to the ‘Receiving’ battery, by the potentially ‘open’ system, in order to return it to its full state of energy and charge. That amount is the same as that dissipated during the ‘discharge’ phase of the test. In that case the voltage will have returned to $V_{(pk)}$ in time $t_{(Vpk)}$

$E_{(Received)}$ = Energy expended and measured in controlled ‘Discharge’

The ratio of $E_{(Received)} / E_{(Supplied)}$ is the Coefficient of Performance (CoP) for the whole device, including its ‘relationship’ to the local environment.

The testing situation can be summarised through the following diagrams:
In summary then, measuring the CoP of this type of generator requires the following stages:

9. A measurement of the energy lost in a controlled discharge of the ‘receiving’ battery from a state of full charge

10. A measurement of the energy delivered by the ‘run’ battery to the generator in operation

11. The return of the ‘receiving' battery to its original energy state and voltage in a measured time

12. The calculation of CoP as the ratio of ‘*energy returned to the receiving battery*’ divided by the ‘*energy supplied by the run battery*’.

The results can be plotted on a graph, using standard spreadsheet software, showing the variable you are changing, e.g. PRF, on the X axis and the result, e.g. the CoP value, on the Y axis as in Fig 49.

This process is then repeated for different operational parameters of the generator.

To give an example of how this might play out in practice, let's say you are wanting to find the optimum PRF for your setup with a specific battery type and capacity being charged. A sample spreadsheet will be offered via the Appendices that will contain all the necessary calculations, along with instructions on using it, to avoid spending a lot of time setting such up. Of course, you are free to amend and improve on it for your needs any way you choose.

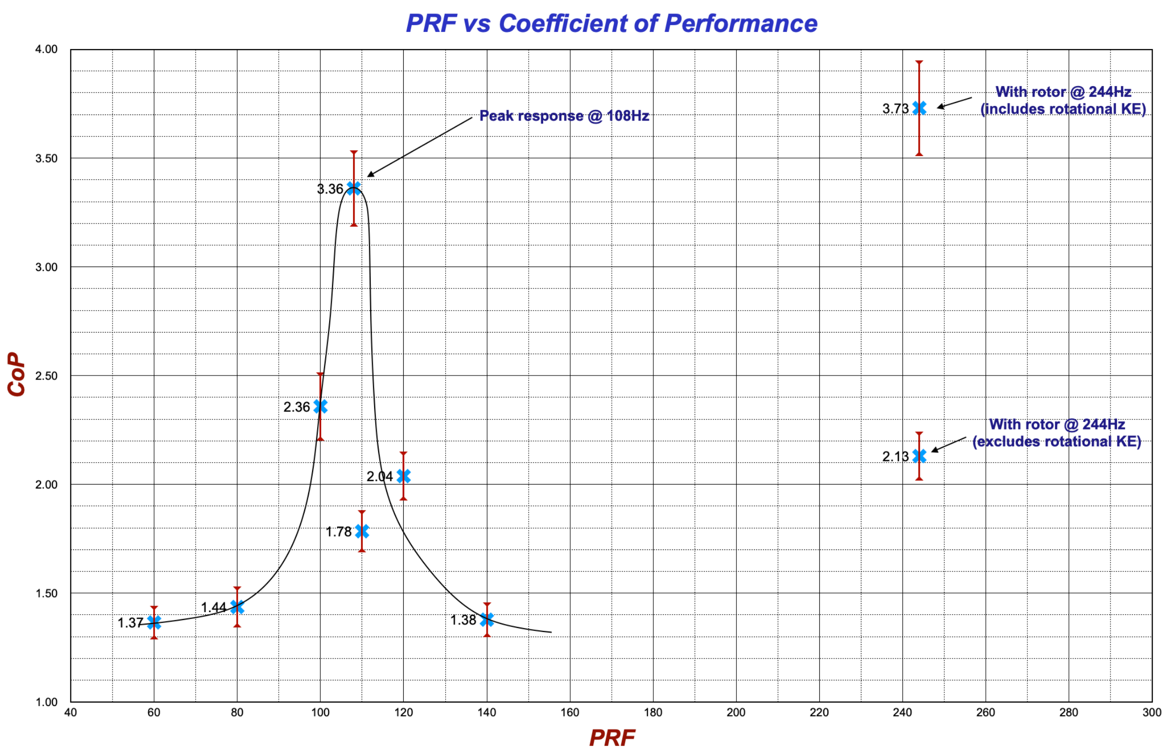

Fig 52: Plot of CoP vs PRF

Starting with a PRF of 60Hz you obtain your first reading, after following the experimental procedure and doing the calculations using the spreadsheet. Then you repeat the tests at 100Hz and 140Hz and plot the results, as in Fig 52.

With even just three points you will most likely see a trend appearing, in this case that the CoP is higher at 100Hz than at 60Hz or 80Hz and falls off at higher frequencies. This will prompt you to try some other values in between 100Hz and 120Hz and try, for example, 105Hz and 110Hz. Doing this iterative process will allow you to zoom in on the optimum pulse frequency, aided by a PWM module that will let you set precise frequencies and a duty cycle and which are temperature stable.

In perhaps as few of seven or eight tests, you will have been able to hone in on the optimum PRF for your particular setup. The same approach can be taken for all the other variables in turn, including, duty cycle, profile charging point, battery capacity. Using a repeatable and consistent method is the only way to do this without introducing other unknown variables that can influence and confuse your results.

As you find the optimum values for your system then you can fix those values in further tests with the remaining variables such that, over time, your results continually improve. This is the pattern I experienced and where at the start my CoP values were in the 2 - 5 range and then, over time, gradually increased to the 30 - 40 range as successive parameters were optimised.

Control Tests

A control experiment serves to remove from the process the one factor that, under the hypothesis, is proposed to result in an energy gain, namely the HV pulses. As such it provides a baseline result in the absence of the factor being tested.

To achieve this, the HV pulses, while being generated in the normal way by the circuit and coils, were diverted away from the receiving battery to a suitable destination that did not interfere with the circuit's operation. Rather than let the pulses discharge to air, or some other grounded path, it was decided to direct them to a set of super-capacitors which would satisfactorily absorb the pulses and, in the absence of a potential divider, avoid the possibility of any high voltage arcing on the PCB.

This was achieved by connecting a bank of super-capacitors made up of six 500F, 2.7V capacitors wired in series to give an estimated 130F at 16.2V. These are connected to the terminals dedicated to measuring the pulses on a scope (H14 & H19) with a potential divider and diverting the pulses to them using the 'Load Switch' (SW4).

The methodology is the same as for CoP measurements and, with a flat 'charge monitor' response showing no voltage increase, then theoretically the amount of energy required to return the battery to its starting voltage is infinite, hence the CoP is zero. A report showing some Control tests is in the Appendices.

Power Tests

Load testing is considered to be essential over and above other practical tests in confirming the hypothesised harvesting function of a generator of this type. Due to the fact that some factors and variables cannot be fully accounted for, or even estimated, until a live load is added to the system, and also in this case battery swapping enabled, then accurate measurements of the amount of load that can be sustained can only be done with all the elements of the generator in operation.

Testing the available power output for this type of device can be done using a variation on the so called 'loop' testing procedure. Here, the generator output is fed back into the input such that, if there is more output than is required to run the generator, with its losses, then extra energy is being drawn into the system. In this event the device will continue to run beyond the expectations of its nominal power supply and presents with a CoP>1.

In the Pulsed Flyback Generator, this process in effect occurs every 15mins or so due to the essential battery-swapping mechanism that is integral to its successful operation. Battery swapping is fundamental to the device's operation since there are no known appliances that can run directly off inductive flyback pulses and so storage in a battery or capacitor is an important part of the energy flow. Indeed, testing using super-capacitors suggests that the battery chemistry is central to the phenomenon and that the electrochemistry of the battery is a fundamental link in the energetic pathway. Without the batteries being a part of the functional chain there would likely be no available power at all.

With two batteries, at any moment one of them will be the supply or 'run' battery and the other will be the 'receiving' battery. The run battery supplies all the energy for the circuit to operate and also any external load attached to the system, while the receiving battery is being pulsed charged. Then at an interval of typically 15 - 30mins, the batteries swap over their roles and the now charged 'receiving' battery becomes the 'run' battery. Never at any time are either of the batteries both supplying energy and being pulse charged at the same time.

Given that the energy used by the run battery to charge up the receiving battery is much less than the amount arriving in the receiving battery, then there is spare energy available to deliver to an external load. For example, if the total amount of energy arriving at the receiving battery is X but the energy supplied to the generator to provide that energy return is much less, say 1X/8, then there is 7X/8 of the energy available for an external load and this would equate to a CoP of 7. The power available would equate to 7X/8t where t is the time taken for the battery to be recharged. In practice t is the swap interval or time which would be set to a value to allow the receiving battery to reach a suitable level of both discharge and recharge.

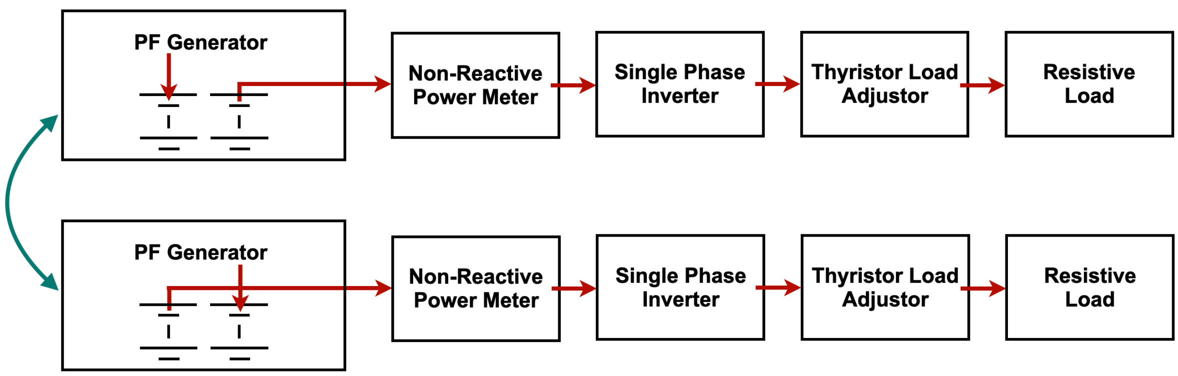

Fig 53: Load testing power train

Any energy hypothetically drawn in to the open system is first stored in the receiving battery before being used in the next swap cycle, as the supply battery, for the circuit and load. So in effect, this ongoing process is equivalent to looping the output back into the input but instead of happening in real time, it occurs with a delay of 15 mins. The energy is from output to storage, then storage to input with resulting output to storage etc. and with the receiving battery acting like a giant energy sink or sponge for the harvested energy and charge liberated by the pulses.

The CoP measurements are then in effect a form of hybrid situation in that they involve this looping process. However, there are also factors and losses that are not easy to quantify and which will affect the real live measurements of the power available from the run battery after a cycle of being charged as the receiving battery.

The proposed setup is shown in Fig 53 and consists of a non-reactive power meter connected directly to the generator output so that a reading is taken before any losses resulting from the inverter. The power meter is connected to a single phase inverter whose 50Hz output is then adjusted using a Thyristor unit to feed a series of incandescent lamps. These provide a purely resistive load ranging from 10W to 300W.

Given this setup, in order to undertake the power measurements, repeated swap cycles are undertaken as shown in Fig 54. At the start of the graph, battery 2 is the receiving battery (black line) and with a starting voltage of V1. It is then pulse charged while battery 1 (the other greyed line) is under load as the 'run' battery that starts at a voltage of V2 and while it provides power to both the circuit and the external load.

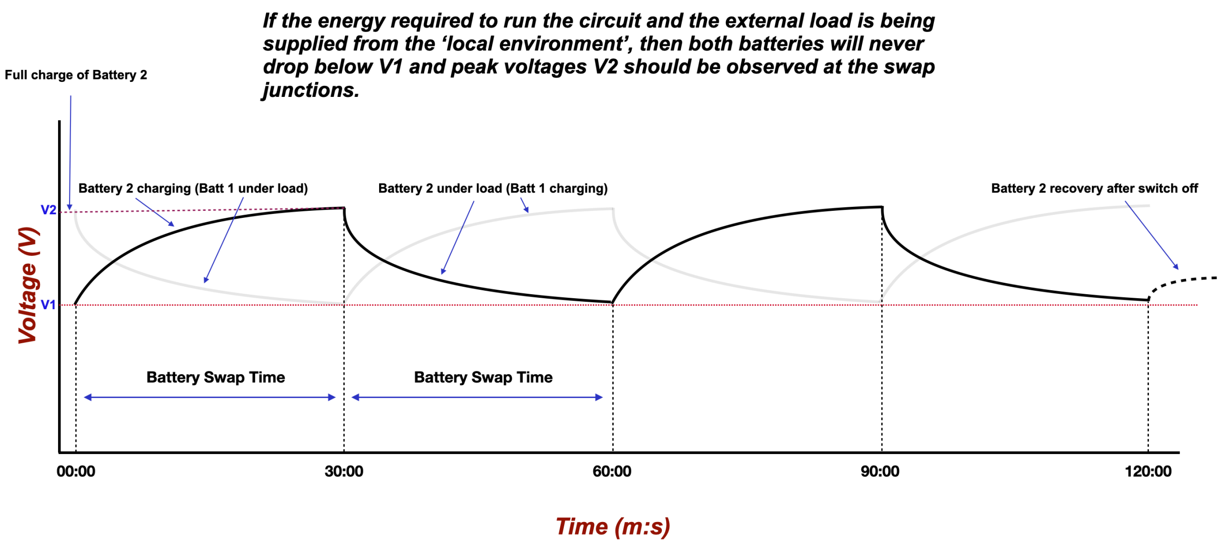

Fig 54: Battery Swapping during Load Testing

Using a swap interval of 30 mins then after that time battery 2 is now at a much higher state of energy and charge and, after the swap, becomes the run battery and battery 1 starts to receive the pulse charging.

If, for example, this swapping cycle continued for 24 hours, there would be 24 complete cycles and 48 swap events. If at the end both batteries have not dropped below V1, indicated by the horizontal red dotted line, and equally are still able to reach their peak starting voltages of V2, then that is clear evidence that energy has been drawn into the system and that the external load is not drawing down more energy than can be replenished during each cycle.

In practice, testing will involve incremental increases to the load to find the point at which a voltage drop is recorded following the recovery phase after a series of cycles. From that value the maximum power output that can be sustained is derived.

At that point, with the known external load and the dissipation of a measurable amount of energy, the net energy state of the batteries will have been maintained through an as yet unspecified process of energy influx.

OTHER RELEVANT FACTORS

The following relevant factors are those that play an important role in the device's performance but which, in addition to the issue involving the peak HV and the active device, have not been discussed in any documentation that I have come across.

Charging Profile

A battery charging profile displays the voltage rise of a battery on the Y axis and energy delivered (or time) on the X axis. They are never a straight line with a constant gradient but usually have a shoulder at both ends, more like a gentle S shape. This is particularly so with Lead Acid batteries whereas with Lithium there is a smaller shoulder at the top with a shallower, flatter gradient for much of the battery's capacity before it rises sharply towards the end and full charge.

As Fig 55 shows, if you are charging the battery on or near the top shoulder, where the gradient is shallower, then you will get lower CoP values than if you are on the stepper gradient of the main charging zone.

In practice, this means that it is a good idea to start the pulse charging from a state of 75% discharge and so move it up towards say 85 - 90% of full charge and then back down again when it takes its turn as the 'run' battery. Working in the 95 - 100% charging zone will therefore be less effective than in the 80 - 90% zone where the battery is more receptive to the charging process and whatever mechanisms are occurring that result in charge arriving at the electrodes.

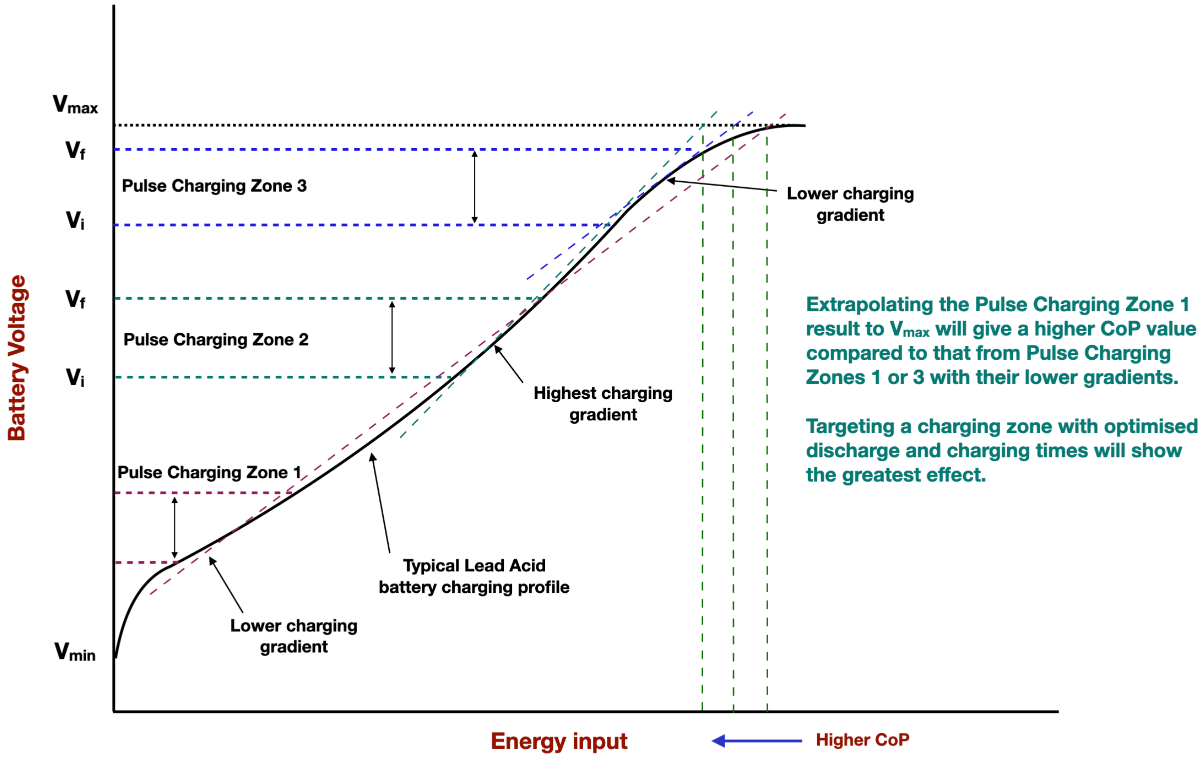

Fig 55: Effect of charging gradient on CoP

It is also worth mentioning here the ‘Surface Charge Effect’ that is an important response during pulse charging. When charge, from whatever source, arrives at the electrodes at a high rate, the battery is unable to fully assimilate and ‘process’ it as there is not enough time for the charge to undergo the necessary chemical reactions and migrate deeply in to the electrolyte, especially if it is in the gel format with its limited mobility. This is why on test runs the battery should be left for 1 hour after pulse charging to let the charge migrate and be fully assimilated by the chemical processes before taking a reading.

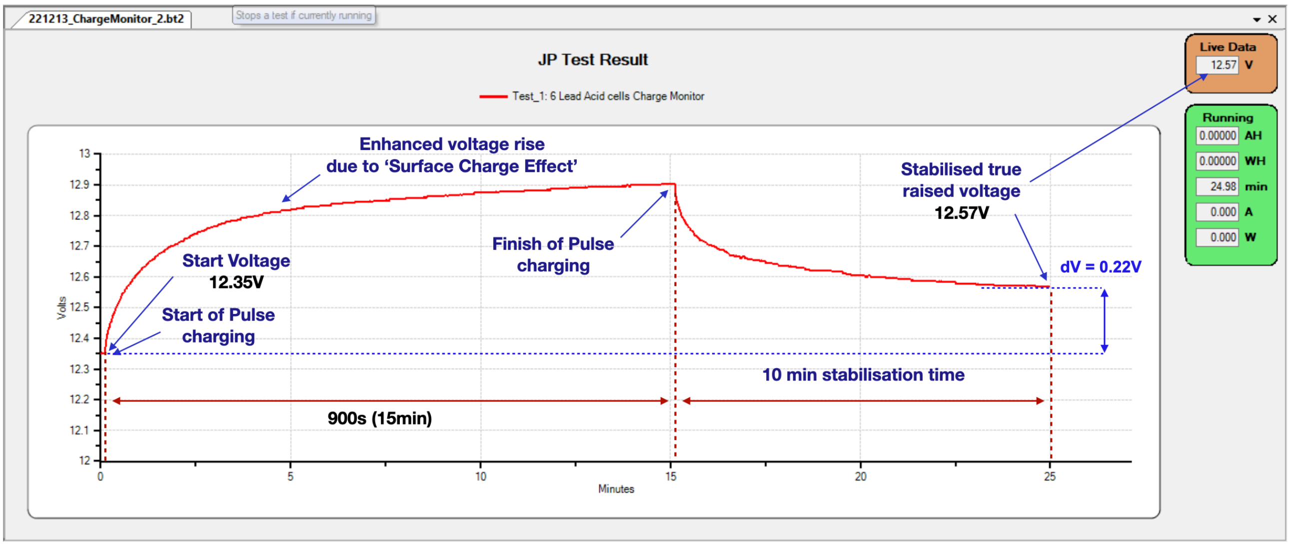

Fig 56: Stabilisation period after pulse charging

The surface charge effect is the reason for the rapid rise in the battery voltage at the start of pulse charging, as demonstrated in Fig 56 with the subsequent stabilisation period. The effect shows as a steep rise in voltage at the start and is also relevant in regular mains charging. In the case of charging a car battery in the normal way, the solution is turn the headlights on for a few minutes to sap off the excess surface charge to obtain a realistic voltage reading.

If during the pulse charging process, a reading of the voltage increase was taken live instead of a stabilised reading, then the CoP values would be hugely inflated. This is both invalid and misleading and would give a false impression of the battery response.

Swap interval

The swap interval is the period of time that the batteries are in their role as either ‘run’ or ‘receiving’ battery before they swap over. This is achieved by the swap circuit, as is described in the ‘Battery Swapper’ section and consists of the 4060 decade counter chip, some components to set the inbuilt oscillator frequency and several relays to flip the source of the power from one battery to the other, route the HV pulses and operate the LED indicators.

The swap interval is adjusted using one of the two 1M trimmers in conjunctions with one of three jumpers that connects different chip outputs to the rest of the swapper circuit. In this way the swap interval can be set between 15s and about 1 hour.

Switch SW1 is used to set either of two swap intervals, one to help check the swapper function and the other to determine whereabouts on the charging profile the pulse charging is taking place.

So what determines the swap interval depends on various factors, not least of which is the time it takes, using the current demand of the circuit and the external load, to take the battery down a suitable %Ah (capacity) to be in the optimum charging zone.

To give an example, if you wish to start the pulse charging at 75% of the battery's capacity then you need to ensure that when the battery is acting in supply mode, that the total combined current demand from the circuit and external load will take it down from the starting capacity to the optimum region on the charging profile. From there it can be pulsed charge back up towards a suitable value. This might mean oscillating between for example 75% and 90% of the battery's capacity.

To illustrate this with some numbers, if we are using a 7Ah battery and the total load current is 6.5A and we are looking to drop from 100% of the battery capacity to 75% during the supply stage, then this equates to 7Ah x 25% = 1.75Ah. In discharging 1.75Ah at 6.5A then this will take a time of Ah/A = 1.75 / 6.5 = 0.269hr = 16.15min = 16min 9sec.

So if the swap interval is set to change the batteries over after 16min 9 sec, then this should place the supply battery at close to 75% of its capacity before it becomes the receiving battery and starts to be pulse charged. While you won't usually need to be this accurate, over numerous cycles the %Ah may drift slightly and take your away from the optimum charging position for your battery and setup. This may or may not be an issue depending upon the way you are choosing to use the device, but it is worth having the information to hand.

If you have found that you need to operate your battery between 90% and 70% then to start your cycles off you would need to discharge 10% of the capacity from a state of full charge (100%) using the CBA's discharge feature. Then you would start discharging the battery in the supply role using the above type of calculation.

The time it will take for the battery to charge back up towards your chosen peak capacity will depend on several other factors that have been mentioned, in particular the PRF. The time taken will be a figure that is derived from each test run based on extrapolation or interpolation of the graph used to derive the total energy supplied to the generator to reach back to the start reference voltage. Using this figure is how the theoretical power output is calculated and is part of the 'in-house' calculations of the supplied spreadsheet.

If the variables are such that, over approximately 16 minutes the battery does not reach up to your preferred value, then over time the battery voltage at the end of its discharge stage will drop lower and lower. What this simply means is that for the power that you are demanding, the pulse charging is insufficient to replenish the energy dissipated in the previous cycle. Small adjustments will then need to be made to either adjust the power demand or modify the various settings to achieve a better charging response, assuming there is still some available to be had.

These details are what the power tests, and the associate methodology, will enable one to determine and what is the maximum amount of power that can be used while still maintaining both batteries at their starting voltages. For this the optimum settings of both the generator and the batteries are first determined by CoP tests.

So to get the best performance from this device, a fair amount of experimentation will be required, even if you find that you are seeing CoP>1 results with minimal adjustments based on some estimated settings or those suggested in the Appendices.

This is uncharted territory and, to quote John Archibald Wheel again, this is a strange thing and it needs to be explored by experimentation as that is the mapping process for this new realm of discovery.

CoP TEST SEQUENCE

What follows is a step by step guide to undertaking a CoP test based on the rationale laid out under the 'Methodology' section. Data can be recorded using a 'measurement proforma' sheet in the Appendices. This was originated in Mac Pages and can be copied or adapted as required.

The key steps in the test procedure are:

1. Externally charge the 'Receiving' battery using a regular mains battery charger and, after at least a 1 hour period of stabilisation (or overnight), record the 'start' voltage using the CBA as part of the setup for the following 'Discharge' stage.

2. Using the 'Discharge' feature on the CBA, dissipate a specific amount of energy (Wh) from the 'Receiving' battery using a discharge current appropriate to the battery, for example 3,000mA for the 7Ah capacity. Record the discharge profile with the X axis as energy (Wh) expended. [D*ischarging 1.0Ah (3,600 C) will result in approximately a 14% reduction in capacity and take about 20 mins at 3A*].

3. Rest the 'Receiving' battery for 60 mins and, if necessary, record its voltage recovery using the CBA. Any subsequent voltage increase from pulse charging will be due to the effect of the pulses and not any residual and continuing voltage recovery after the discharge stage.

4. Switch on the power supply and the generator at the main switch and the Recording Digital Multimeter (RDM) and set it to Auto for a suitable measurement interval e.g. 60s. Set the PWM unit output to 'on', start the CBA 'Charge Monitor' function and the start the pulses using Driver switch SW5)and a stopwatch simultaneously. Run the generator for the chosen time (t seconds) and switch 'off' at the main switch but leave the 'Charge Monitor' function running for a further 10 mins.

5. The CBA monitors and records the voltage continuously and displays it on a Vt graph as the charging monitor profile. Recording for a further 10 mins allows the surface charge to migrate and stabilise before a final value is read.

6. During the charging process, the current supplied by the PSU (acting as the 'Run' battery) to the generator is recorded typically every 60 sec. The average current $I_{(av)}$ is determined as the mean of a set of readings after being exported as an XLS file.

7. For the voltage, a 'no load' DVM value is taken at the coil + terminals (H35/38) before the FET/IGBT switch is turned on. If you are using different coil and circuit supplies then they can be taken from H15 or H20 depending on whether you are using the Buck or Boost converter.

8. This allows calculation of the total energy in Joules supplied to the generator by the PSU during its run time t seconds as:

   $E_{(Supplied)} = V_{(av)} . I_{(av)} . t$ J (*Equation 1)*

9. At the end of the stabilisation time the CBA live voltage value is recorded as the final incremental and stabilised receiving battery voltage.

10. The energy delivered to the ‘Receiving’ battery after a single charging period is unknown since the battery’s voltage is not expected to return fully to its pre-discharge level and, even if it did, the true voltage is masked by the ‘surface charge effect’. Instead, a plot of ‘Receiving battery voltage' vs $E_{(Supplied)}$, is plotted and extrapolated (or interpolated if the voltage has risen above the starting reference voltage - $V_{(pk)}$) to give a reading of the energy supplied to return the battery to its original peak voltage of $V_{(pk)}$.

11. Determine the extrapolated value of $E_{(Supplied)}$ J as *Value 2*

12. The energy delivered to the receiving battery = ‘Energy Discharged’ J as *Value 3*

13. Calculate the Coefficient of Performance (CoP) as: *Value 3 / Value 2*
    (Total energy supplied to ‘Receiving’ battery / Total energy supplied to generator)

14. Calculate the uncertainties using an appropriate statistical method (see the ‘Uncertainty Analysis’ document for details) and repeat with other variables and, where necessary, plot CoP as a function of the variable.

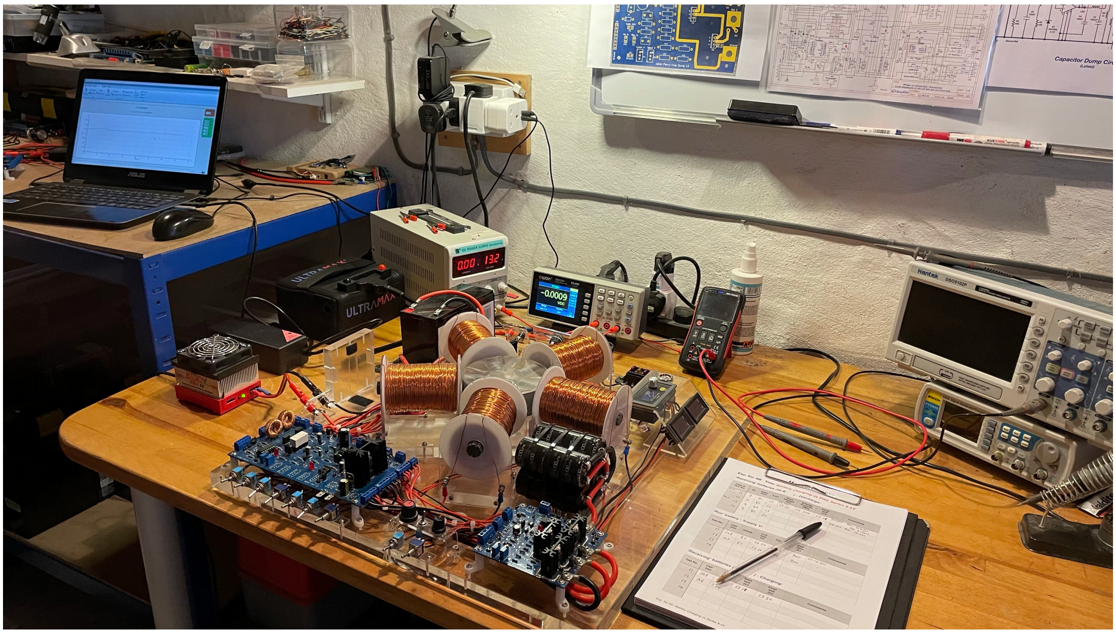

Fig 57: CoP test measurement setup (pre-replication v2 PCB)

The general testing and measurement arrangement is shown in Fig 58, although different tests required different battery connections and configurations.

POWER TEST SEQUENCE

[As power tests are yet to be undertaken, this section will be completed during 2023 along with some indications of the power levels obtained.]

DATA RECORDING

If you have chosen to undertake some test procedures on this device then you will need to engage with data recording. Without an organised system that keeps the data and readings coordinated, you will quickly come undone on account of the 2nd Law of Thermodynamics - namely increasing entropy or disorder. While living systems may display negative entropy at a local level, having a look in your waste bin should convince you that you are generating an increase in entropy at a surprising rate. The same applies with information, especially when it is tied to physical and electronic recording systems, digital or otherwise.

The mechanisms available for data recording are numerous and in this project include a digital recording meter, graphical plots of voltage or energy against time, as well as paper based systems. Perhaps depending upon how many full moons have passed since you were born, then you may have dispensed with paper systems altogether and rely purely upon a smart phone, an iPad or similar.

So long as you are detailed in your records, it does not matter and indeed, I continually use a smart phone for keeping a record of tests that need to be done using a form of shorthand, alongside paper based data recording of the actual tests. I have shared it below on the basis that, for some, it might prove a useful aid when designing experiments without having to have all your data around you. Using Google 'Keep' or similar, I used a form of shorthand for the test parameters that would help in setting up sequences of tests to, for example, identify the optimum PRF or the effect of different coils voltages on the results. Seeing them grouped together was helpful in designing sequences of tests.

The first bolded line indicates the contents of the lines beneath and in this example we have the coil voltage set, the degree of discharge at the start of a test, the % capacity discharged, the charging time (which equals the swap interval when swapping is enabled for regular use), the PRF, square wave duty cycle, battery chemistry, capacity and the pulse HV. The % just indicates that the test had been completed and the figures in a line are bunched together so it would fit on to one line of the note editor. You can of course devise other categories, as is shown next for doing some tests with super-capacitors.

Coil V (%Ah, Swap) (PRF, %Duty) (Batt, Ah) (kV) % Done

12.5(85,20,10)155,65 (7Li)(1.7) %

12.5(85,20,10)120,65 (7Li)(1.7)

Coil V (Time, Vmax) (PRF, %Duty) (kV) % Done

12.0(15,15.0) (60, 65) (1.7) %

12.0(15,15.0) (65, 65) (1.7)

RELEVANT EQUATIONS

For those who would like to know a little more about the physics, here is a selection of equations relevant to some of what is happening in the 'regular' parts of the system. What is occurring to induce the energy gain has yet to be resolved. There are also published scientific papers in the Appendices that consider some of the options and the general area of the 2nd Law of Thermodynamics and under what conditions it can be 'violated'.

$$\frac{Total\ Energy\ Out}{Total\ Energy\ In} = \eta \ \left(assumes\ a\ closed\ system\ \leq 1\right)$$

Efficiency

$$\frac{Total\ Energy\ Out}{Total\ Energy\ from\ User} = K \ \left(assumes\ an\ open\ system\ 1 \leq K < \infty\right)$$

Coefficient of Performance (CoP)

$$B = \frac{\mu_0 IN}{L}$$

Magnetic flux density in an air core cylindrical solenoid (T also 1Wb/m² = 1Tesla)

$$B = \frac{\mu_r \mu_0 IN}{L}$$

Magnetic flux density in a cylindrical solenoid with a ferrite core (T)

$$V_L = N\frac{\Delta\Phi}{\Delta t} = K\frac{\mu N^2 A}{l}\ \frac{di}{dt}$$

Self-induced EMF in an air core solenoid [K=Nagaoka coefficient] (V)

$$V_L(t) = -\frac{\Delta\Phi}{\Delta t} = -L\ \frac{di}{dt}$$

'Induction EMF' in an air core cylindrical solenoid (V)

$$\epsilon = -N\frac{d\Phi_B}{dt}$$

EMF induced in a coil of wire with N turns as a function of rate of change of magnetic flux B [Faraday's Law] (V)

$$\phi_B = B\,.\,A = BA\ \cos\ \Theta$$

Magnetic flux calculated from flux density [T] and area A (Wb)

$$\Delta G = \Delta H - T\Delta S$$

The Gibbs free energy as a function of enthalpy H and entropy S at a temperature T (kJ mol^{-1})

$$\Delta H = \Delta E + P\Delta V$$

Enthalpy H as a function of the internal energy E and the temperature at a constant pressure P (kJ mol^{-1})

$$S = k_b\ ln\Omega$$

Entropy is a measure of the number of possible microscopic ways that a macroscopic state can be realised. (J.k^{-1} mol^{-1})

HEALTH & SAFETY

At this juncture it is appropriate to make the normal health and safety and other obligatory statements in the context of what is presented in this document.

The work described in this manual is advisory and presented only to inform and educate. Any practical work inspired by it is undertaken at readers' own risk and the author accepts no responsibility for any injury or other detrimental consequences of building what is described within its pages.

Having said that it is pertinent to state the obvious, that working with high voltages, even if they carry little conventional current, still requires caution and sound practices to avoid unnecessary accidents or mishaps, let alone damage to sensitive electronic devices.

This is an area of investigation where being fastidious, highly organised and diligent are worthy and desirable traits.

May positive and fascinating insights emerge from your endeavours.

APPENDICES

The appendices provide a range of additional documentation in support of the build, data recording and analysis, as well as a selection of relevant published papers. It is available in the 'Appendices' sub-folder on the same Mega and Dropbox storage links:

https://mega.nz/folder/YUM0nLoT#bYpLIazqMM5K2IrEQjghDQ

https://www.dropbox.com/sh/td55b8675vvqtbg/AADzPSKMOI8q_YM1cFUT2T07a?dl=0

File name	*File name*
Author Biog	Measurement Proforma (editable)
Avalanche Guidelines	Owon XDM1041
Basics of digital multimeters	PCB v4 Schematic (also in 'PCB files' folder)
CBA_IV_Manual	PCB v4 2D view (also in 'PCB files' folder)
CD4060BE	PCB v4 Net view (also in 'PCB files' folder)
CoP vs HV Table	PF Gen V4 Component List (PCB)
Costings	Redox reactions in Lead Acid Batteries
DHG1-i1800PA	Relevant equations
DSEI 12-12A	Relevant published papers (listed below)
Hall effect sensor	Sample spreadsheet (with data) Excel/Numbers/PDF
HANTEK DSO5000 Manual	Sample spreadsheet (empty) Excel & Numbers
Heat sinks	Spreadsheet Guidance Notes
HFD2 Relay	STH12N120k5
Interim Report 2 (Using caps)	STH12N170k5
Interim Report 3 (Control)	STP20N90K5
IR2121	Suggested settings
$LFePO_4$ 18Ah battery	Switch connections
Lithium 12V 7Ah battery	Test Run checklist (editable)
LP12-7.0 battery	Uncertainty Analysis

It includes component specifications and information sheets as well as recording forms and sample spreadsheets, with separate guidance notes, so you can the process test run data and derive useful results with uncertainty values. Additionally, there is a subfolder of useful research papers on related areas of enquiry and which will be continually added to as more relevant material becomes available.

	Relevant Published Papers
1	*Beyond the second law of thermodynamics* by Danial Sheehan
2	*Beyond the Thermodynamic Limit: Template for Second Law Violators* by Danial Sheehan
3	*Chemical Thermodynamics - A chemical wonderland* by Rubin Battino and M Letcher
4	*How batteries store and release energy: Explaining basic electrochemistry* by Klaus Schmidt-Rohr
5	*Energy Renewal: Isothermal Utilisation of Environmental Heat Energy with Asymmetric Structures* by James Lee
6	*Non-invasive yet separate investigation of anode/cathode degradation of lithium-ion batteries due to accelerated aging* by Pouyan Shafiei Sabet et al.
7	*Polarisation of Vacuum* by Stanislav Konstantino
8	*Practical Conversation of Zero point energy* by Thomas Valone
9	*Quantum Thermodynamic devices: from theoretical proposals to experimental reality* by Nathan Myers and Sebastian Deffner
10	*Sustainable energy and the second law of thermodynamics: An introduction to the special issue* by George Hathaway and Daniel Sheehan
11	*Type B energetic processes and their associated Scientific Implications* by James Lee
12	*Vacuum radiation induced by time dependent electric field* by Bo Zhang et al.

The information in this document is provided freely on a non-commercial basis for the use of interested parties to further scientific knowledge and understanding of the world. If you would like to use quotes or data from it then please include the following citation:

Perry, Julian. (2023). Pulsed Flyback Generator - Assembly & Guidance Notes. Kerrow Energetics, Penwith, Cornwall, UK

Contact: jp@kerrowenergetics.org.uk