SUGGESTIONS TO HELP ACHIEVE A CoP >1

This guide aims to share some technical details with regard to the successful operation of a Pulsed Flyback Generator or equivalent. My research has shown that while a CoP>1 can be achieved using a set of build specifications and 'standard' settings, the best performance requires the optimisation of a range of factors, not least the battery, so that the pulses have the best chance of creating the desired effect at the electrodes.

Battery Swapping:

Battery swapping is essential to the operation of this type of device for the simple reason that a battery will not cope with acting as a power source, for the circuit and any load, at the same time as being required to process incoming pulses. For this reason the power required to run the circuit and the load are supplied by one battery (the 'run' battery) while the other 'receiving' battery is being charged by the HV pulses.

Every 10-20 mins or so the batteries swap over their roles and the now more charged receiving battery becomes the run battery and vice versa. In this way the energy harvested is outputted indirectly instead of in real time. This being so then the source of the energy influx cannot be due to the effect of the pulses on releasing chemical potential energy since the pulses are only used to charge a battery and not cause it to release energy. The energy that is released from the run battery has been stored within it in the normal manner, as chemical energy (Gibbs free energy, enthalpy etc.) and the conversion of this energy back to electrical energy is also well understood. What is not well understood and overlooked in the opinion of many is where the extra energy brought into the device is coming from or, if we are sure of where it is coming from then we are not clear on the precise mechanisms.

So a battery swapping circuit is essential to observe the effect of HV pulses on a battery. If you don't have one then there is a good working circuit design in the files (see link at bottom) and which I have integrated into my main PCB with the relays that are integral of it. However, it can also be made as a stand alone unit if required and would be straightforward to produce a PCB from.

Just to be clear on a few terms and concepts (that often seem to get confused in this area of work by some enthusiasts) this device has an estimated efficiency of around 85% but I have yet to measure if accurately. You can't have an efficiency more than 100% because of the way the term is defined and for a closed system. So talking about a system having an efficiency of over 100% is not only incorrect but damaging to the work of many as it shows a lack of awareness of some very basic scientific concepts. Of course in an open system, like this type of generator, LENR etc. you can have more energy output than the energy you, as the operator, put in and so that is why we have and use the term Coefficient of Performance (CoP), although in mainstream science this term tends to be reserved mostly for heat transfer process. There is a slide in the ‘Graphics’ folder illustrating this.

Pulse Repitition Frequency (PRF):

The rate at which HV pulses are delivered to the receiving battery is the same frequency as the signal to the MOSFET gate, i.e. the square wave frequency from the 555 oscillator chip, or a PWM module or from the Hall sensor and rotor system if one is using that method. The most effective and flexible I found to be the PWM since the output frequency and duty cycle is not only easy to set but temperature stable, whereas the 555 based timer will often drift a little during operation due to the effect of temperature on the RC component values used.

Using the rotor method will only give you a fixed PRF at the max rpm (at 3000 rpm this will be 250Hz) and you have no way to adjust it to suit the battery configuration or receptivity. Also the energy expressed in the spinning rotor can’t be extracted easily without slowing it down and so changing the PRF.

Undoubtedly the best way to trigger the FET and coils is to use a PWM module that was under £10 on eBay or Aliexpress etc. Adjustment down to the nearest Hertz is required to get the best performance from the system. With my system the optimum PRF to use with the small 7Ah AGM batteries is 108Hz and 65% duty cycle whereas with an 18Ah $LiFePO_4$ battery it’s 155Hz and the same duty cycle.

MOSFET & Driver:

It is surprisingly common for the ubiquitous IRF840 device to be used in these types of generator as it’s easily available and cheap. However, it is very difficult to get a CoP>1 with it for the simple reason that the maximum spike voltage you will see at the Drain is around 500-600V. This is because the limiting factor for the peak HV is not in fact the coils but the ‘avalanche rating’ of the MOSFET. Any voltage over and above that rating will cause a breakdown and the grounding of the spike so that in effect clips the voltage to that limit.

If you calculate the theoretical flyback voltage from a coil when switched off in say 10µs by the FET, the voltage will come in at over 10kV. However you never see that at the Drain for the reason explained above.

To get some decent CoP results you should aim to generate flyback pulses of at least 1,000V and for that you can use a MOSFET such as the STP20N95K5. Also, for the Diode that sits across the Source and Drain, that grounds the pulse that is created when the coils switch on, rather than the one we want when it switches off, then the DSEI12, which has a reverse breakdown value of about 1,200V, is suitable.

If one is going for the higher voltages made possible by the STW12N120K5, the STW12N150K5 and the STW12N170K5, with peak Drain voltages of 1.2kV, 1.5kV and 1.7kV respectively, then the DGH10i1800PA diode should be used which can manage up to 1,800V (see ‘Peak Spike Voltage’ below)

As I have yet to implement these other higher voltage devices I can’t yet report on how effective they are and my results so far are based on the STP20N95K5/DSEI12 alone. I plan to start with these other in November and will report back on how effective they are.

Cap Dump Circuit

There has been much written and promoted about using a capacitive discharge system in ‘Bedini’ type circuits. Often referred to as ‘Cap Dump’ circuit, they function so as to receive and store the energy of the HV pulses in a bank of capacitors first before discharging the energy in the form of a high current intensity pulse to the battery. This has been referred to as ‘HV to high impedance source and then to a low impedance source’.

I spent at least three months redesigning and testing the circuit in a ‘high sided’ format and then integrating it into my build for the express purpose of seeing if the claims made about it giving good performance were to be upheld. That was just part of doing a complete job on the testing.

The device is set up to discharge its stored energy when the capacitor bank voltage reaches 24V and then stop the discharge when it drops down to 17V. On the charge monitoring graph one could clearly see the little ‘bumps’ as each high current pulse (~100A) hit the battery but the effect on CoP was minimal.

The short answer is that the results were very disappointing and gave results in the 0.8 - 1.5 range. This reinforces the notion that it is dV/dt that is important and not dI/dt for the energy harvesting phenomenon to take place. So my advice is not to bother building one and to focus your efforts elsewhere on other more productive aspects of a build.

Again this is just my own experience and you might find things different or be driven by a healthy and admirable desire to find out for yourself. To flesh out my opinion, I have added an additional report to the files on my testing of the cap dump circuit.

Other Factors:

While the PRF is certainly a very important variable, there are a variety of other factors that can very significantly affect the CoP value and therefore the available power for an external load. These include the number of coils, the peak spike voltage, the coil voltage (the number of batteries used in series), the battery chemistry and battery capacity (Ah), the placement on the charging profile where the charging is taking place and the swap interval.

The only way to determine the optimum settings for all these is to also have a way of measuring the CoP with each test run, using a method such as that explained in detail in the earlier linked report (Interim report 1 (CoP). Without doing this you are taking pot luck with the settings and, while you might get a CoP>1, it will likely be in the 1-3 range instead of >10 or more.

Let's take each of these factors in turn:

Coils and their number:

It is generally safe to assume that the more coils you have the better the battery response but I noted that the CoP measured with 4 coils was lower than with 3 and then rose again with 5. As with some other factors, there is a balancing act with the energy input requirement for running the coils being set against the resulting effect on the battery.

Each of my coils was wound on an easily available plastic spool and consisted of about 2,600 turns of 0.71 enamelled wire and with a ferrite rod in the middle. Each coil had an inductance of 370-400mH and resistance 13-15Ω. When these are wired up in parallel they gave a measured total inductance of 20-30mH and a resistance of 1-2Ω. (There are some pictures of the coil winding process in the folders on the link at the end)

Peak spike voltage:

The harvesting phenomenon appears to be almost directly related to dV/dt, the rate of change of voltage of the spikes, and so the higher the spike voltage seen at the FET Drain the higher this value (for much of my testing dV/dt was around 2×10^8 V/s). As described, this is limited by the 'avalanche rating' of the active device and not the coils, which will normally be creating a flyback voltage (as indicated by Lenz's Law) much higher than what you see at the Drain.

I suggest starting with the STP20N95K5 and DSEI12 Diode and then change once you have some data under your belt. Incidentally, the device format for this FET is the TO-220, and also for the STW12N120K5, but for the two higher ones (150K5/170K5) it is TO-247. So that I could swap in and out these different pin formats, I devised a socket that could take both types and, after taking off the sockets wings so it could fit inside the heat sink I already had. There is a pic of the original socket I purchased from Aliexpress in the folders.

These higher voltage devices are currently hard to come by with the general shortage of silicon devices after Covid but a good site I found, that will hunt down supplies of them for you, and other hard to get components, is: https://www.jinftry.com
Using this site I was able to source these FETs when the likes of Farnell, Mouser, RS and others were devoid of stock. However, as I have yet to implement these higher voltage devices I can't yet report on how effective they are and my results so far are based on the STP20N95K5/DSEI12 alone. I plan to start with these other in November and will report back on how effective they are.

Using this socket also required that I create a custom made device symbol and footprint for the PCB construction process. After my paper is completed, I do plan to make freely available the Gerber files for a PCB so that anyone can get it printed rather than having to design one yourself. I will write more about this at the end.

It is invaluable to have a device that will enable you to measure the spike voltage without damaging a scope. For this you will need a voltage divider that is relatively simple to make. The circuit is in the folders and also a few photos of it. Rather than direct the HV pulses to a battery, they are directed to the divider and the output to a scope. Using a 10:1 ratio, a 1,500V spike will show as 150V which is well within the tolerance of most scopes. Ideally it needs to be calibrated using a signal generator on impulse setting, where you input a 5V impulse/spike signal and adjust the pot to give 0.5V on the scope.

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, so 24/25V and 36/37V by adding batteries to both sets in series. This has the effect of increasing the voltage across the coils and also the whole circuit. The latter will 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 require the use of a buck converter to power that part of the circuit. Either that or find a whole new set of components that operate at the higher voltages. Using a Converter is far simpler and I used one as part of the testing regime to see how 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 add a Boost Converter to be able to stabilise the coil voltage and currently I can choose which to use to set my variables; but this is unlikely to be required as well in a regular generator.

However, CoP tests I have done on two batteries so far (using three have 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 is that when you are powering an external device with 24V compared to 12V then the current demand is halved. This is kinder on your swapper relay that in my first build was only a Hongfa HFD2 relay capable of only 3A. I have since revised the board using a 20A 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, then I suggest sticking with one pair of 12V batteries (see Capacity below) for the whole system, and which will switch over after the swap interval, and for testing you can use just one together with a power supply for convenience.

Battery chemistry:

Traditionally sealed Lead acid batteries have been used as they are cheap to purchase and the AGM type can be used in any orientation. Most of my early work has been with the 7Ah AGM type but I have found much better performance using a 7Ah LiFePO4 battery and which is commonly available. They are more expensive but lighter, can undergo many more cycles and have a higher energy density, being able to discharge a higher continuous current for the duration of the swap interval, say 15 mins, before the other battery takes over. They have a different and flatter charging profile as discussed under Charging Profile.

Battery capacity (Ah):

Larger capacity batteries have a larger surface area for interaction with the pulses than the lower capacity ones and have been found to give better results. While doing a lot of work and tests runs using the 7Ah ones, and which are quick to charge up to the starting voltage for the beginning of a test, the 17/18Ah give better results in terms of CoP and power availability.

I suggest doing the basic checks etc to get the system going using two 12V AGM SLAs and then move to 7Ah LiFePO4 and then ultimately, if you can afford them, to 18Ah $LiFePO_4$ batteries.

Bear in mind that for most of the testing, and certainly for CoP tests, you only need one battery to act as the receiving battery. The run battery can be replaced by a power supply that gives you much more flexibility to test different supply voltages. It also doesn’t require being charged up every test run and which allowed me to undertake two test runs a day. Again some details are shown in the photos on my Mega account via the link at the end.

Charging profile:

This is not an issue that I have ever seen mentioned but it plays a significant role in the results. 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 and usually have a shoulder at both ends more like an S shape. This is particular so with Lead Acid batteries whereas with Lithium there is a small shoulder at the top and a shallower gradient for much of the battery’s capacity before it falls off sharply of a cliff at the end.

As a graphic in my files shows, if you are charging the battery on the top shoulder, where the gradient is shallow, then you will get lower CoP values than if you are on the stepper gradient of the main charging area.

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% 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 the 80-90% zone where the battery is more receptive to the charging process and whatever mechanisms are going on to deliver charge to the electrodes.

It is worth mentioning here the ‘Surface Charge Effect’ that I addressed in my interim report. Delivering lots of charge to the electrodes at a high rate means that the battery cannot fully assimilate it all and there is not enough time for the charge to migrate deep in to the electrolyte, especially if it is a gel with limited mobility. This is why on my test runs the battery is left for 1 hour after pulse charging to let the charge migrate and be fully assimilated by the chemical processes before taking a reading. This effect is also applicable in regular mains charging and, in the case of a car battery being charged with a mains charger, one simply has to turn the lights on for a few minutes to sap off the excess surface charge to obtain a realistic reading.

If you took a reading of the voltage increase during pulsed charging based on the live values, then my CoP values would be in the 100s. This is not valid and misleading and yet I have seen it done before and where the relevant parties are clearly unaware of the surface charge effect.

Swap interval:

The swap interval is the period of time that the batteries are in their role as either ‘run’ or ‘receiving’ before they swap over. This is achieved by the swap circuit, that is my case is integrated into the main PCB where the connections to the battery come in. It consists of a 4060 decade counter chip, some components to set the inbuilt oscillator frequency and one or more relays to flip the source of the power from one battery to the other.

I have placed a circuit design in the files showing the latest design I’m using and which is capable of handing up to 20A. For some this can be a separate module in between the batteries and the trigger and FET circuits but any PCB files I produce in the future will be made available for free access and use, and where the swapper will be an integral part.

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 1 hour.

I will put together another document explaining the calculations for what oscillator and specific pin output frequencies relate to what interval times. This will allow you to set any interval you like quickly by adjusting and reading the relevant frequency on a scope; but it’s a bit too detailed to go into here.

My own setup has a switch that I can set to one or other of two times. One is set at 15s to observe and check the function of the swapper and relays, and to set the relay to a desired position, while the other is set to my chosen value for actual operation of the system.

What determines the interval depends on various factors, not least of which is the time it takes, using the current draw (circuit and external load) that you have, to take the battery down a suitable %Ah (capacity) to be in the optimum charging zone as explained above.

So if you started the process at 75% capacity and are drawing a total of 5A, then given the battery capacity, e.g. 7Ah, if you want to allow it to drop 10% (0.7Ah), then you would set the swap at 8.4mins. After that time, at 5A you will have used 5A x 0.14hr = 0.7Ah and which will have taken you to an appropriate region of the charging profile. Then the battery will start to charge up back towards 75% capacity, although that does depend on how well you have set up the other factors to allow for the optimum charging effect - PRF etc.

Methodology:

As previously stated, if you want to find the optimum settings for your device, then you need a robust and repeatable method to undertake CoP tests. The methodology I developed was based on that used by Gary Hammond and others from the Energy

Science Forum and I added various features and form work to make it easy to use. When it comes to the all important load testing, I was not aware of any system available and so have come up with my own as described in the ‘Load Testing’ document.

This requires suitable equipment such as a CBA and is explained in the interim report. Of course when adjusting a variable like PRF, you need to keep all the other factors the same or your will have no way of knowing what is causing any changes - this is basic experimental protocol. To help keep the method on track, I used sheets showing all the steps so they eventually became automatic. Using such aids in the early stages of testing is very helpful as well as giving you some data to look back on.

It is good practice to make lots notes of what you are doing and what changes you make to the circuit or method. When problems arise, as they inevitably will, a logical approach, together with your detailed notes, will help you eliminate the possibilities until you find the resolution to the problem. If nothing else this type of work is an exercise in perseverance and diligence!

Files:

On the link below I have added a selection of files under the folders including docs, circuits, build, graphics and pics. I will probably add a few more files in the coming months.

Part of the process of my writing a scientific paper on this research is that others should be able to replicate the work based on what is in the paper. For that to be able to happen I will need to have revised the circuit and remove what doesn’t work well based on my finding. Also the PCB will need to be redesigned and other materials prepared to allow anyone reading it to potentially be able to repeat my findings.

This means that, probably towards the end of the spring ’23 I will have a revised the main PCB and details of design changes, such as getting rid of the rotor system and the cap dump circuit. I can let you have all that information via my Mega account folders.

So for those who are just starting this journey, you may want to wait til then to start or, if you are unstoppable or have already started, then I hope the information I have given here will be useful to you to make adjustments as you see fit to get improved performance.

I feel it is worth starting here that I am doing all this because I feel this information is of use and value to everyone and also because, as a scientist, it is part of the process of finding gaps in our knowledge and to continually push the boundaries of what we know and understand about the Universe.

As I’m retired, I have no commercial motivation for doing any of this and it is ‘curiosity driven’. As such I am providing this all free and with no expectation of return of any kind. However, if you get some promising results and progress then I would be interested to hear about it. That is reward enough.

That being said then I need to make this obligatory statement here:

This information is provided on the basis that I accept no responsibility for any consequences resulting from any of its use and doing so is a decision solely for the user. It is offered on the basis of open-source knowledge and with no desire or expectation of a return of any kind.

Here is the link: https://mega.nz/folder/xF12XBKK#FBieZOk74f3ZeFiAa4g-3A

Julian Perry
November 2022