

Five ways to save money in hydroponics

Commercial hydroponics can be extremely expensive, given the technological complexity and supplies required for a successful crop. The biggest costs are usually related with the initial setup but subsequent crops can also become very expensive, especially if you are using boutique fertilizers or additives that can get very expensive very quickly. Today I want to talk about five ways in which you can save money in a hydroponic crop from a crop-cycle perspective.

Avoid buying liquid concentrates as fertilizers. Liquid fertilizers have some intrinsic advantages – like their homogeneity – but they contain a lot of water, which means that you will need to ship more than one pound of water for every pound of fertilizer you get. This will increase the cost of the fertilizer significantly, even if you're buying fertilizers in bulk for a commercial crop. When buying single bulk or blended fertilizers make sure you always buy solids to greatly save on these costs.



Prepare your own blend of fertilizers for macro nutrients. The most complicated part of fertilizer preparation usually deals with the micronutrient portion of fertilizers, if you want to be as simple and cost efficient as possible you can actually

buy this portion – some companies specifically sell the micro part – and then prepare all the macro fertilizer blends yourself. You can then hire a consultant or read the scientific literature to get a formulation you can then use to prepare your macro portion from bulk commercially available fertilizers (which are extremely cheap).

Prepare your own foliar treatments. Foliar spraying can greatly reduce problems and increase crop yields, so it is usually a no-brainer to make sure you use foliar sprays within your crop cycle. Some of these foliar additives can be very expensive though, but it can be very cheap for you to prepare your own additives if you have the proper know-how.

Use a recirculating nutrient system. Drain-to-waste nutrient setups are extremely wasteful. If you want to have a crop that is as cheap to run as possible you will need to go to a proper recirculating setup. Once you do this you will be able to use your recirculating solutions for weeks before having to carry out a nutrient change and, even then, there are some techniques that might allow you to keep your nutrient solution for even longer. Imagine if you only needed to prepare/change nutrients once every blue moon.



Make sure you use silicon additives. Many growers fail to use silicate containing additives within their crops and generally suffer from a far greater chance of having losses due to fungi. Potassium silicate is extremely cheap and with it you

can make your own silicon containing additive that you can use to greatly fortify your crop against fungal disease. A small additional expense can save you a lot of loss and heartache down the line. You can save a lot of money by avoiding commercial hydroponic silicate products and instead making your own silicate additive yourself from potassium silicate.

When implemented, the above changes can help a commercial operation save tens to hundreds of thousands of dollars per year in nutrients, additives and crop losses. Even only implementing a couple of the above can help a mid sized operation save a ton of money in just fertilizer if, for example, in-house macro fertilizers are used, or if a recirculating system with proper nutrient management is established.

Of course, the above steps are not trivial so I would recommend anyone attempting to do them for the first time to get someone with experience in the hydroponic industry to guide their hand through the process. That could either be me or any other highly experienced consultant in the field of commercial hydroponic growing and nutrients management. If you have enough time and the inclination to do so you could also try to learn the above things yourself from scientific literature and online resources, but if you choose to do so I would advise you try to implement what you learn in smaller crops before scaling to larger projects.

**Using electro-degradation to
enhance yields in**

recirculating hydroponics

The efficient use of nutrient solutions is a very important topic in hydroponics. Although some commercial growers use run-to-waste systems where solutions are not recirculated, the economics of fertilizer use often demand re-circulation in order to enhance nutrient utilization and maximize growing efficiency. However one of the biggest problems found when circulating nutrient solution continuously is the build-up of plant exudates, which can be toxic and detrimental to plant growth.




Image taken from [this article](#)

Several solutions for this have been studied historically, most commonly the use of filtration systems – such as activated charcoal cartridges – to capture these exudates and prevent their accumulation. The problem with this approach is that activated carbon – or other filters – are not neutral to some of the components of nutrient solutions and might disproportionately and efficiently capture metal chelates and eventually cause nutrient deficiencies. There are some ways around this – such as changing the formulations or replenishing solutions after filtering – but both are far from

ideal.

More recently [a paper](#) has been published showing how electro-degradation can actually alleviate this problem by destroying these exudates – which are commonly organic acids – in nutrient solutions. The paper talks about how they used this technique to treat recirculating solutions in strawberry, eliminating autotoxicity and increasing fruit yields substantially.

The technique is very simple, basically using either a DC or AC current passed through an electrode that the solution circulates through, destroying the problematic molecules in the process. The first image in this post clearly shows how not renewing the solution causes important problems with yields that are completely removed by the use of the AC based electro degradation.

Table 1

Changes in mineral nutrients after application of electro-degradation of nutrient solution in no plant experiment. Electro-degradations were applied in 10l of 25% standard “Enshi” nutrient solution with $400 \mu\text{M L}^{-1}$ benzoic acid for 24 h. (Experiment II).

Electro-degradation	NO_3^- (ppm)	P_2O_5^- (ppm)	K^+ (ppm)	Ca^{2+} (ppm)	Mg^{2+} (ppm)	Fe^{3+} (ppm)
Control ^a	687	37.5	7.9	49.9 a ^d	16.2	3.5 a
DC-ED ^b	658	35.8	7.6	41.6 b	13.8	2.2 b
AC-ED ^c	669	37.5	7.2	52.6 a	15.4	3.4 a
Significance	NS	NS	NS		NS	

^a Electro-degradation was not applied.

^b Electro-degradation was applied using “Direct Current”.

^c Electro-degradation was applied using “Alternate Current”.

^d Means within a column followed by different letters are significantly different and NS indicate non-significant according to the Tukey's test at $P < 0.05$.

Image taken from [this article](#)

Another advantage of this technique is that – contrary to filtering techniques – there is little loss in the amount of nutrients in solution when performing the AC electro-degradation. Since the oxidation/reduction of the metal chelates used is highly reversible, the actual concentration of these elements in solution remains practically the same after treatment. You can see this in the image above, where

there is no statistically significant change for the concentration of nutrients in solution.

The paper concludes suggesting a treatment of 24 hours (for 300L in the experiments) every three weeks, to completely recover from the exudates present in solution. For this AC application they used a frequency of 500Hz at 14V with an electrode area of around 53 square centimeters, made of titanium metal. For this process you need an inert metal or conductive material that will not react at the potential values used. You can buy titanium metal tubes – which are not expensive – to build an anode/cathode pair to carry out this experiment. *Note that the frequency and voltage characteristics are vital so using a proper power supply to generate them is of the highest importance.*

The above technique is novel and easy to build for treating commercial hydroponic solutions. It is far easier and economic compared with filtering techniques and can be applied from smaller to larger scale growing operations.

Using machine learning control methods in hydroponics

Properly controlling the hydroponic environment is perhaps one of the most challenging tasks the modern grower must face. Either with a small grow room or a big green house, it is difficult to properly control variables such as temperature, humidity, pH and nutrient concentration, ensuring they are all kept in tight ranges with the proper controlling actions always being applied. Today we're going to talk about some of

the research done into advanced control systems and how using these could help you boost your crop yields.

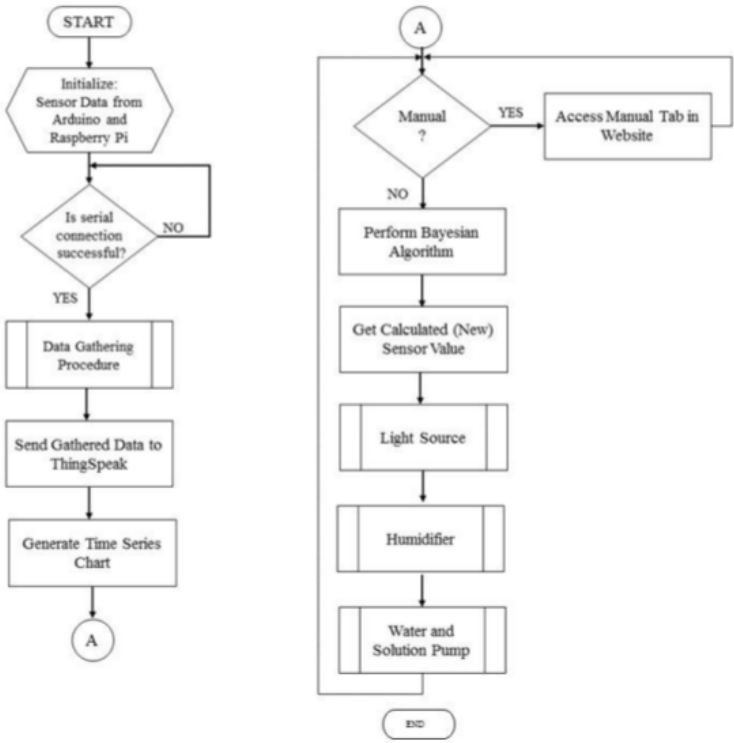



Fig. 2: System flowchart

Control flow algorithm taken from [this paper](#)



(a) Manual Control (b) Automatic Control

Fig. 7: Actual yielded crops matrix

Crop result comparison taken from [this paper](#)

Hydroponic crops are dynamic systems, with plants continuously affecting their environment and demanding control actions in order to keep conditions constant. For example plants will tend to transpire water and absorb carbon dioxide during their

light cycle, so in order to keep humidity and carbon dioxide concentrations constant you might need to turn on humidifiers, dehumidifiers, carbon dioxide generators, etc. Knowing what action needs to be taken is not trivial and naive control implementations – like turning on humidifiers, AC systems, etc when some thresholds are reached – can cause problems where sensors fight each other (for example a sensor trying to increase ambient humidity and another trying to raise temperature) or even fail to trigger.

In order to provide better control, researchers have created systems that rely on machine learning – systems that can learn from examples – in order to learn what control actions are needed and execute them in order to provide ideal control to a hydroponic setup. A machine learning system will be able to anticipate things like the lag between turning an AC unit on and the temperature decreasing, so it will be able to be both more efficient and more accurate in the way it controls your environment. This use of automated control guided by machine learning is also known as “smart hydroponics”.



Fig. 11. Overall process of the System

Control flow diagram from [this article](#)

For example you can read [this paper](#) where growers were able to increase the yield of a crop by 66% just by ensuring they could maintain proper environmental conditions the entire time using machine learning. In this case the researchers use a probabilistic method where the system determines the probability of an action – like triggering a sensor – will cause a desired effect. As data is accumulated the system basically executes whichever action has the highest probability to lead to the desired outcome.

There are other papers on the subject. In [this one](#) a deep learning neural network is used to perform a similar control role, although the quantification of improvements in this paper is not sufficient to claim that the control method would have been an improvement over a traditionally managed hydroponic setup, as the comparison is made between a soil control, not a hydroponic control with no automated environmental management.

[This paper](#) uses a simple IoT sensor control system and a multivariate regression approach in order to control the environment in a hydroponic greenhouse, this system was created with the aim to be cheap and usable in developing countries.

Although there are now several different demonstrations of this being done in the literature there still does not seem to be a commercially mature technology to carry out this task and the implementations seem to still be tailor made to each particular situation. However the modeling techniques used are not exceedingly complex and even modest commercial growers could – nowadays – afford to setup something of this nature.

With a computer, some arduinos, raspberry pi computers, sensors and time and effort a grower could definitely setup a very nice, machine learning based control system to benefit from the above described technologies.

Calcium's behavior in hydroponics

Calcium is often one of the most puzzling elements in hydroponic culture due to its ability to respond fairly non-linearly to nutrient concentrations in solution. This behavior is the result of its transport dynamics and its relationship to other elements that may antagonize it very effectively when they reach higher concentrations. On today's post I will talk about calcium behavior in hydroponics and how we can play with both environmental and chemical properties to change its concentration in leaf tissue.



A leaf showing symptoms of a calcium deficiency. This is usually NOT caused by a lack of calcium in the nutrient solution.

Imagine you've had a good growing season up until now, your plants are looking great, you've been doing everything properly. Suddenly, you start to see what appears to be a calcium deficiency on your leaves. You proceed to analyze your tissue and notice that your Ca levels are way below what you

expect, yet your nutrient solution seems to be very Ca rich, at almost 150-200 ppm. You panic and increase the Ca level to 250 ppm, your following test results come out even lower. You're not alone, you've just misunderstood calcium transport.

Most calcium deficiencies are actually not the result of Ca missing in nutrient solutions but they are caused by faulty Ca transport, which is often related with environmental issues. Calcium transport depends substantially on transpiration, so the solution to Ca deficiencies can be as simple as increasing your vapor pressure deficit (VPD). Ca is also absorbed more effectively when its at a lower concentration than at a higher one, so often increasing Ca will decrease its transport to leaf tissue. [This study](#) on tulips and its bibliography illustrates this fairly well, the increase in tissue with Ca follows a parabolic trajectory, where the largest Ca concentrations actually lead to lower Ca in tissue. You can push Ca to the leaves with higher VPD though, as [this paper](#) on Ca fortified lettuce shows, with the lettuce grown at Ca at 300ppm at 28C showing the highest Ca accumulation.




Fig. 1: Relationship between sesquiterpene lactone content and calcium content in the flower of *C. coronarium* L.

Taken from [this article](#). It shows that terpenes can increase at larger Ca content in tissue in some flowering plants.

At lower levels Ca can actually start to show the inverse behavior and start to accumulate very heavily in tissue as its transport can become extremely favorable. If you notice a large increase of Ca in your leaf tissue at your plants ideal VPD then you might actually want to increase Ca in solution rather than decrease it, as an increased concentration in solution might actually make transport less favorable. However the most determinant factor in Ca absorption is water transport, so this excessive Ca might just be indicating you that your VPD is too high (so reduce temperature or increase your relative humidity).

This big influence of VPD explains why results of ideal Ca concentrations and Ca:K ratios are significantly disperse in the scientific literature. A [recent paper](#) for strawberries shows this ideal ratio to be around 1.3-1.4 but some papers (like the one on tulips shown before or [this one](#)) has it way closer to 1.0. The ideal ratio for your crop will also be dependent on the water transport your plants are forced to assume due to your ambient conditions so you will likely need to optimize this variable for your particular growing conditions.

A good place to start for flowering plants in terms of K:Ca is usually a ratio of around 1.2 however, you will need to do tissue analysis to figure out whether the Ca absorption is at the ideal point or whether you want to increase/decrease your Ca. Just bear in mind that increasing Ca in solution might reduce it in your leaf tissue so to reduce Ca in tissue try to play with your VPD first before you play with the concentrations in solution. Chances are that if you're getting too much or too little your issue is that you're too deviated from your ideal VPD situation.


Better understanding pH dynamics in hydroponic culture

Plants require specific conditions within their root zones to give optimal yields. The negative of the logarithm of hydronium concentration (a.k.a pH) is a very important measurement that tells us if the root zone is too acidic or too basic for our plants. However pH is affected by many different variables, making the calculation, theoretical understanding and modification of this variable difficult – and often very puzzling – for a significant percentage of growers. Today I want to go into how we can understand these pH changes from a qualitative perspective to get an idea of how pH can be expected to change in solution if we modify one variable within our hydroponic crop.



The final pH in the root zone of a hydroponic crop will mainly depend on the following factors: Chemical species in solution, plant absorption/secretion and media reactions. Basically we need to consider the chemistry of the solution, the way plants absorb nutrients and the chemistry of the media. The size of these effects relative to one another is also very important to consider.

In a normal hydroponic crop with a nitrate rich solution plants will tend to absorb more nitrate than any other ion, which will tend to increase pH as plants will excrete bases when they take anions in. Since the effect of nitrate in pH is very small the net effect is a substantial increase in the pH of the nutrient solution. Anions like citrate can also cause increases in nutrient solution pH due to a similar reason, they can be strongly taken in by plants and generate upside drifts.



Model of a citrate ion molecule.

The effect of anion absorption can be offset by acids present in solution. Phosphorus is most commonly present in hydroponics at a pH of 5.8-6.2 as H₂P₀4(-) which is a weak acid that is able to react with the bases excreted by plants. This means that higher P containing solutions will tend to drift less towards the upside than solutions that contain more P. This is a common reason why yields are usually better for less experienced growers when higher P solutions are used, because they are more forgiving to pH drifts due to their higher buffering capacity towards the upside at the pH used in hydroponics. When we add P the effect of the buffer is way bigger than the effect of P absorption, as P is not absorbed in very high concentrations by plants.

When you increase or decrease the concentration of an ion in solution you need to consider if it will be absorbed, how

strong its absorption will be relative to other ions and how this relates to the chemistry of this ion in solution. For example citrate provides a relatively strong buffer towards both upside and downside moves at a pH of 5.8-6.2 but its absorption by plants negates this effect and generates a net drift towards the upside in solution in most cases.

Species that are not absorbed so significantly – like bicarbonate ions – are better at buffering moves towards the upside, since they provide a weak acid to react with plant secretions without providing a very significant source of anions for the plant to absorb and exchange for more basic species.

This exact same logic can be applied to positive ions that are easily absorbed by plants. Adding ammonium generally causes strong decreases in the pH due to plants exchanging ammonium with hydronium ions. However this effect can be compensated by adding things that cause upside drifts either through absorption or through buffering effects. Adding ammonium citrate to a nutrient solution causes a rather balanced effect due to both the ammonium and citrate absorption/buffering compensating.

The media is also very important to consider as non-neutral behavior will tend to strongly drift pH due to the mass of the media within the hydroponic setup. To prevent this problem it is not uncommon to treat media with a buffer before use, although this can become prohibitively expensive quickly if large volumes of media are used. It is more common to deal with the quirks of the media using nutrient solution chemistry. To get an idea of how your media affects pH you can let your media soak in nutrient solution and notice how pH evolves as a function of time. If the pH increases or decreases you'll get an idea of what you should expect and then think about how you might want to handle it.



Peat moss

Media that decompose – particularly peat moss – will tend to acidify and become more and more acidic with time. This effect can be so strong that the only reliable solution is to amend with some low solubility base that can offset this effect with time (such as limestone). This is however problematic since the base will tend to run out towards the flowering stages of crops, where its help is most needed to offset potassium absorption. For this reason it's usually more manageable to use a media like coco coir, which has a more stable pH profile and will not tend to "fight growers" through the crop cycle.

As you can see pH in hydroponics can be a complicated issue. When it starts to drift you need to immediately think about what the source of the ions is. If the pH increases you have a contribution of basic ions that exceeds the strength of the weak acids in solution, while the exact opposite is the case if the pH decreases. What is the source of these ions? Is the plant absorbing more of something? Do we need to add something to counter the effect? Is the media the culprit? Start asking and trying to answer these questions and it will become easier to understand why your pH is drifting and what you might be able to do about it.

Five important things you should know about humidity in hydroponics

Humidity is one of the most important metrics in hydroponic culture. In order to have a successful crop it is critical to understand what humidity actually is, how we measure it, why

we need to control it and where we want it to be. Today I want to share with you five things you should know about humidity in hydroponics. These pointers should be equally useful to newer growers and those who want to get even higher performance out of their current hydroponic crops.



There is a big difference between relative humidity and absolute humidity. Although we generally understand humidity as the amounts of water in the air in reality what we usually measure in hydroponics is “relative humidity” (RH) which does not measure how much water there is in the air but what percentage of the available capacity we are using. A measure of 70% tells you that the air currently holds 70% of the water it could hold, but it doesn’t tell you anything about how much water there actually is in the air. Absolute humidity, on the other hand, tells you how much water you actually have in the air. As warmer air can carry significantly more water it is important to realize that a 70% humidity at 80F implies there is way more water in the air than a 70% humidity measurement at 50F.

Relative humidity meters are not reliable instruments. We often buy instruments to monitor ambient conditions without much thought about how they work or how good or bad they are. Relative humidity is a tricky measurement and cheap, semiconductor based relative humidity meters have not been very reliable or accurate. Usually a humidity meter will have an error of +/- 5% and it’s accuracy will not be on point if

it has been exposed to very high humidity values (if there was ever any condensation on the sensor it probably was damaged to a significant extent). I often recommend buying at least 3 instruments with different chipsets. Having just one – or even several meters with the exact same chip – can be a recipe for disaster. The chipsets below are setting a new standard for precision and accuracy, so I would recommend you give sensors with those a try if you're looking for more accurate RH measurements.

Company	Humirel	Sensirion
Part Number	HTU21D	SHT21
Product		
Humidity Accuracy	±2%RH	±2%RH
Temperature Accuracy	±0.3°C	±0.3°C
Operating Range	0—100%RH	0—100%RH
Power Supply	1.5—3.6V	2.1—3.6V
Interface	I2C	I2C
Power Consumption	400uA	300uA
Response Time	5s	8s
Package	3x3mm DFN	3x3mm DFN

It is often better to go over than to go under. Although higher or lower relative humidity values are both sub-optimal for plants, it is often better to go with higher humidities rather than lower humidities in terms of crop yields. Lower RH values will tend to stress plants more – especially if the temperature is high – while higher RH values are often easier to deal with for tropical plants (which are the kind we often grow in hydroponics). Although higher humidity can definitely cause important issues – such as fungal diseases – we have ways to deal with this that are more effective than our ways to deal with stress caused by low relative humidity. Of course extreme values will be very detrimental to plants either way, so when I say high consider I don't mean 95% RH at 80F.

Relative humidity can change a lot depending where you measure it. It is important to place RH meters at different places that represent what the plant is actually being exposed to. If you place RH meters in a greenhouse, far away from plant canopy, you will get a very poor representation of what the plant is actually experiencing and you might try to increase humidity substantially when this might not be needed. Ventilation is critical to alleviate this issue but in order to stay on top of it I always advice having meters within plant canopies in order to know for a fact how much your humidity values diverge between different places in your growing environment. Humidity is always bound to be higher closer to plants – as they transpire – but we should know how different it is. A big difference is a strong hint that there is not enough ventilation around the plants.

Humidifiers are often needed during the winter. Most crops that grow in warehouses during winter times require humidifiers, since the plants often do not evaporate enough water to compensate for the complete lack of water within the dry winter air (this is specially true if the volume of the growing warehouse is large). If you want to have a successful crop during winter times it will be paramount for you to have adequate humidifiers to ensure that your RH values are within what's ideal for your plants. Depending on the size of your crop this might require significant planning and investing so make sure you always consider this when designing your winter growing cycles.

I hope you have found the above pointers useful. There are certainly many other important aspects of humidity, such as [relative humidity meter calibration](#), judging ideal humidity using [vapor pressure charts](#), or ensuring that [plants have enough defenses](#) if they happen to get exposed to exceedingly high humidity levels within their canopy. We will certainly discuss some more of these within future posts!

DIY Warm white LED lamp PAR measurements, not so exciting after all!

If you read my last few posts about DIY LED lamps built using 150W warm white LED cobs (which do not require an independent driver) you might have been excited by some of my claims. I previously stated that you could probably get around a 1000W HPS equivalent using just two of these lamps, which meant an energy saving of around 60% relative to the HPS equivalent. However to really verify these claims I wanted to get new PAR and lux meters to perform proper PAR and lux measurements. The results my friends, are disappointing.

42	56	20
53	220	76
43	62	24

Previously I thought that these lamps were close to half of an HPS equivalent based on initial lux measurements. At the same distances, directly below the lamp, I could get around half the lux equivalent of your average HPS lamp, I thought from the warmer spectra of these white warm cobs that the PAR contribution would be significantly higher than that of a regular HPS but it seems that – due to the inefficient use of a white phosphor to produce the spectra – basically the PAR

efficiency is equal to that of an HPS lamp.

The PAR (Photo-synthetically Active Radiation) basically measures the number of photons that can be used in photosynthesis that you get per square meter per second off a given light source. I will write a more in-depth post about PAR in the future, but it basically tells you the plant-usable photon flux you get. It is therefore measured in $\mu\text{mol}\cdot\text{s}^{-1}\cdot\text{m}^{-2}$.



I performed classic PAR measurements with a 150W lamp 15 inches above a target center with measuring points around a 4 square feet area (to compare with the variety of HPS measurements you can find [here](#)). The results, in the first image in this post, show you the map of PAR values across the 2 feet by 2 feet area. This shows that the lamp is basically giving you $1466 \mu\text{mol}\cdot\text{s}^{-1}\cdot\text{m}^{-2}$ per 1000W at its highest point, which is below the PAR/watt of even the poorest HPS models. With this lamp model using 150W cobs you will therefore need at least 7 lamps to reach the same equivalent of a 1000W HPS in terms of actual photo-synthetically active radiation.

Not only that but without any focusing or dispersing elements the PAR decay as a function of light distance is much more dramatic than for regular, reflector mounted HPS lights. **With all these information it now seems clear that these warm white**

light LED cobs are NOT a good HPS replacement.

However the idea of the zip tie lamp is not dead! I found out that there are actually “full spectrum” LED cobs that are specifically designed to be grow lights (so basically a combinations of red and blue LED lights). These cobs come in 20, 30 and 50W formats and they should have a much more favorable PAR than the 150W warm white LED cobs. I have now ordered some of these cobs ([here](#)) to rebuild my zip tie lamp and see if I can indeed get a much better PAR/watt and watt/dollar compared with normal HPS lights.

Cheap DIY high power LED grow lights: Introducing the Zip-tie lamp

Make sure you also read [this post](#), where I studied the PAR of these lamps and realized they are not as good as I thought!

Several months ago [I wrote a post](#) about using high power LED cobs that do not require an external driver in order to build a high power DIY LED lamp. However I hadn't built a practical lamp using these cobs at that particular point in time so I just gave a general idea of why I would use these diodes and how the particular lamp setup would work. Today I want to talk about how to build one of these lamps in practice using an aluminum heat sink, a 150W warm white LED cob, a fan and some zip ties. The setup lacks the use of any adhesives and should provide you with roughly a 40-50% equivalent of a 1000W HPS. With two of these lamps you should be able to run the equivalent setup to 80-100% of a 1000W HPS in terms of PAR with around 60% less power consumption.



The idea of this post is to help you build a very affordable DIY lamp. However please note that this lamp involves work with mains voltages which are dangerous. Please familiarize yourself with all the precautions needed when working with high voltages. **All the information herein is provided as-is for educational purposes with absolutely no guarantee, either expressed or implied.**


To build this lamp – showed above – you will need these materials (note that if your country uses another voltage you will need to buy the appropriate pieces for the voltage in your country):


- [Warm white 150W LED cob](#)
- [200x60x30mm aluminium heat sink \(2 needed\)](#)
- [110V-120V AC fan](#)
- [Nylon zip ties 30cm](#)
- Cable and wall connector
- Thermal compound (optional)

Initially I wanted to build a lamp using a high power warm white LED cob by gluing the cob to the heatsinks using a thermally conductive glue. However the problem with this is that these glues very permanently bind the cob to the heatsink so if for any reason the cob fails you would lose the heatsinks because the cob would be bound to them. For this reason I decided to use zip ties instead, which provide an

easy way to secure the entire ensemble and allow you to easily replace any failing part rather quickly. I used nylon zip ties but you can also use stainless steel ones if you want the setup to be more resilient (although things will be harder to cut if you make a mistake).


To assemble the lamp I basically used 4 zip tie lines two horizontal and two vertical. For the lines that go the width of the heat sink I just had to use one zip tie but for the other two lines – that also go above the fan – I had to use two zip ties for each line (you can connect one zip tie to another to have a larger zip tie). You need to tighten the zip tie very hard to ensure the cob is in direct contact with the aluminum along all its length, you can also use some thermal compound (like the one you use for CPUs) between the cob and the aluminum heat sink for maximum heat transfer. The pictures below show you a bit better how I performed the entire assembly. *When putting the fan on top of the heat sinks make sure the airflow is towards the heatsink (flow arrow in the case pointing down) and that the fan can spin freely).*





Finally I connected the cob directly to the AC line by soldering the appropriate live/neutral cables to the connectors at the left side of the cob (in the above picture). I then covered the soldered spots with silicon glue to ensure that the connections were as electrically isolated as possible. Make sure you solder as small portion of wire as possible and make sure the wire makes absolutely no contact with the aluminum heat sink or you will have a short. I also soldered the fan cables to the live/neutral as well since the fan can be driven directly by AC as well.

Since you have the zip ties you can also use them to hang the lamp, you can also add screws to the fan screwing ports and use those to hang the lamp from the ceiling. When I turned on this lamp its power consumption was around 220W – measured directly from the wall – meaning that it consumed a bit more power than what was advertised (which is not uncommon for these cobs). Since my voltage is a bit higher than 110V – which is the minimum rating for this cob – I actually get a slightly higher light/heat output than someone using it at a lower voltage. The fan – which takes around 12-15W on its own, also contributes to this consumption level.



When you power on this lamp – image above (sorry about the camera not being able to handle the light intensity) – you’ll immediately notice how the heat sink starts to heat up. I have tested the lamp through 2 hours of continuous operation up until now and the heat sink reached a stable temperature of around the 120°C ($\sim 250\text{F}$), the final temperature you reach will of course depend on your ambient temperature and how well you assemble the components. It is however very important for you to test each one of these lamps for 12/24 hours to ensure that they don’t heat up excessively. *Nylon will melt at around 220°C so you definitely don’t want your lamp to ever reach even close to that temperature (to be safer you can use stainless steel zip ties)*. However it is very likely that the LEDs will burn out way before this happens if your temperature rises too far. You can also add a second fan or use a larger heat sink if your temperature is too high.

In the end the setup is extremely simple to build and you can get roughly 40-50% of a 1000W HPS with one of these lamps. With two of these lamps you will run at around 450W which is 55% less power than an equivalent HPS setup. Although heat generation is no joke here, it is indeed much less than the comparable heat output of a 1000W HPS. With a cost of less than 80 USD per lamp you will be able to build these lamps at a far lower cost than the very expensive grow lights you can get online (which can often go for thousands of dollars for a

single 1000W HPS equivalent). *If you read my earlier post you will notice that I previously thought you needed 4 cobs to reach the equivalent of a 1000W HPS, turns out you only need two 110V cobs running at 120V!*

I have made some PAR, lux and temperature measurements but I want to keep those for a future post where we will look at some of the spectral and thermal characteristics of this lamp vs other lamp types.

Potassium concentration and yields in flowering plants

From the different nutrients that are needed by plants we have known for more than 4 decades that potassium is of critical importance to flowering/fruiting plants. Potassium is one of the most highly limited nutrients in soil due to its high mobility and great increases in yields have been achieved with both potassium fertilization in soil and the use of properly balanced nutrient solutions containing enough potassium in hydroponics. But how important is potassium and what is its ideal concentration in hydroponic nutrient solutions when growing flowering plants? Today we are going to take a look at the scientific literature about potassium and what the optimum levels of potassium for different flowering plants might be in order to maximize yields.




Fig. 1: Effect of potassium concentrations on quality traits in strawberry cultivars.

There are many studies in the scientific literature dealing with the effect of potassium on various flowering plants. Earlier evidence from the 1980s pointed to optimum concentrations of potassium being close to the 160-200 ppm range. The book “mineral nutrition” by P.Adams ([here](#)) summarizes a lot of the knowledge that was available at the time and shows that for the growing techniques available at the time using greater concentrations of K was probably not going to give a lot of additional benefit.

However newer evidence from experiments carried out within the past 10 years shows that optimum potassium concentration might depend on a significant variety of factors, from which media, other nutrient concentrations and growing system type might play critical roles. For example study on strawberries in 2012 ([here](#)) showed optimum concentrations of K to be around 300 ppm for strawberries and the optimum media to be a mixture of peat+sand+perlite (image from this article included above).

Table 3: Effect of cultivar on yield of tomato and ascorbic acid concentration in fruit.

Cultivar	Yield (kg/plant)	Ascorbic acid (mg/100 g FW)
Avinash-2	2.00a ^a	25.69a
Pant T-3	1.74b	22.80b

^aValues in a column followed by the same letter are not significantly different, $P \leq 0.05$, Fisher's LSD.




Figure 3: Yield, pooled over cultivars as affected by supplied K concentration. The regression was $Y = 0.0024(X) + 1.1733$ ($R^2, 0.871$).

Evidence from experiments on tomatoes ([link here](#) and image from this article above) also shows that for tomatoes the actual optimum concentration of K might actually be larger under some condition with the optimum in this study in terms of fruit quality and yields being 300 ppm. In this last case the tomatoes were grown using a nutrient film technique (NFT) setup. However there have also been studies under other growing conditions – like [this one](#) on reused pumice – which shows that increasing K concentrations to 300ppm can actually have detrimental consequences. In this case tomatoes fed at 200, 290 and 340ppm of K had very similar results when using new substrate but the old substrate heavily underperformed when high K concentrations were used.

Papers published on the effect of different K concentrations in melons ([here](#)) and cucumbers ([here](#)) also point to optimal concentrations in the 200-300 ppm range and for the optimum N:K ratio to be between 1:2 and 1:3 for these plants. This is probably the reason why you will often find suggested nutritional guidelines for flowering plants – like those below taken from [here](#) – mostly suggesting K concentrations in the 250-350ppm range. However you will often find that they

directly contradict research papers, like this guideline suggesting K of 150 ppm for strawberries while we saw in a recent paper that 300ppm might be better. This is most probably due to differences in the sources used which might have used different growing systems or plant varieties which responded to other conditions better.

CROP	N	P	K	Ca	Mg
Tomato	190	40	310	150	45
Cucumber	200	40	280	140	40
Pepper	190	45	285	130	40
Strawberry	50	25	150	65	20
Melon	200	45	285	115	30
Roses	170	45	285	120	40

Concentration in mg/l (pap)

All in all the subject of K concentration in hydroponics is no simple one. Using low K will limit your yields tremendously but increasing your K very high can also harm your plants, especially depending on the type of media you are using. In general aiming for a K concentration between 200-250 ppm is safest but in many cases increases to the 300-400ppm range can bring significant increases in plant yields. A careful study of the available literature and the actual growing conditions that the plants will be subjected to will be key in determining what the best K concentration to use will be. Alternatively carrying out adequately designed experiments under your precise growing environment will help you carry out an evidence-based decision about what K concentration to use.

Five reasons why a dedicated hydroponic testing room is a great idea

Most commercial hydroponic setups completely lack testing environments. The most common reason for this is that commercial crops are meant to produce revenue and a testing environment means dedicating space, time and money into something that might not be as productive as the rest of the production facility. Furthermore a testing room implies that you will need to create a completely independent setup and hire someone who knows how to do research in order to ensure it is both adequate and fruitful. Although many people believe this not to be worth it today I want to talk about the five most important reasons why I consider that a testing room is something incredibly useful to have as a part of your commercial growing facility and why getting one will probably pay off greatly for you going forward.



Testing product changes. Perhaps the first and most direct benefit to having a testing room is to ensure you can test

product changes. It may be the case that your supplier for some particular fertilizer product or additive has ran out and you now want to test a new product to replace it. It may also be that you want to test how a product does compared to what you generally use but you don't know if it does better or worse. Most growers are afraid of change because making facility-wide changes that won't work could have huge financial consequences. A testing room ensures you can test safely and then roll-out changes slowly without having to risk your entire crop cycle to find out.

Optimizing what you currently have. Change is very rare across commercial facilities because growers understandably want to preserve their current results, even if some better results by making some change would be possible. This constraints growers from making incremental changes that might make their crops significantly more productive. By having a testing room you can optimize the setup you already have by making adequate research into optimizing things such as environmental or nutritional factors.

Trying potentially game-changing modifications from academic research papers. There are many papers published each year on how to increase the yields of hydroponic crops. Some of these papers offer somewhat risky and controversial changes that might not transfer well across species. However if something gives you the potential to increase your yields by say, 50%, it might definitely be worth trying across a testing setup. Obviously these things are too risky to try across an entire facility but a testing room would be perfect to help you try these new and exciting modifications, potentially giving you a huge edge versus all the other people who will never try this.

Try new plant varieties. Usually growers try new plants without having a clear idea of how productive they are going to be under their growing setup. This means that you introduce a new variety with a huge question mark regarding its productivity and potentially financial benefit or cost. A

testing room provides you with a risk-free way to test how a particular plant variety will perform under the exact conditions in your facility, potentially allowing you to make far less risky decisions when it comes to making planting changes in your facility.



Research new ideas. A final benefit you can get from a testing room is that you can research your own new ideas. With adequate experimental design even a small room with just 10 plants can be used to test some ideas to see how they affect plant growth. This means that you can develop your own in-house growing modifications that will make it much harder for others to compete with you. For example if you developed a secret foliar additive in your growing facility it would allow you to only use this for your own crops, without the industrial secret ever being used by your competitors.

Of course there are many other advantages to testing rooms but the above are just some of the wonderful things you'll be able to do if you have a testing room and someone trained in scientific research who can help you design experiments and get the most out of it. A testing room also doesn't need to be huge and even starting out with 10 plants can be a huge step in taking your commercial growing facility to the next level.