Nutrient problems and foliar sprays

Nutrient related issues are common in hydroponic crops. They can happen due to a large variety of issues, including pH drifting, EC drifting, lack of proper nutrient ratios, humidity issues, temperature issues and root damage. The fact that an issue is of a nutritional nature will be evident within a leaf tissue analysis, but its correction by changing the nutrient solution's composition might not be evident, since transport problems imply that a deficiency in tissue might happen for a wide variety of reasons different than the concentration in the nutrient solution being "too low" (read more here). In today's post I will talk a bit about why the quickest path to recovery might actually be to perform foliar sprays instead of only attempting to change the chemistry of the nutrient solution.

Let's first talk a bit about nutrient transport in plants. A foliar analysis might be showing you a low level of an element like K in tissue, but this does not necessarily mean that the plant doesn't have enough access to K in the nutrient solution. All we know from a foliar analysis is that K has not been able to go into the leaves, but this doesn't automatically mean that K in solution is too low. This problem

can happen if the temperature of the room is too high and the relative humidity is too low – very high VPD conditions – in which calcium and magnesium will be uptaken very aggressively and the plant will be deprived of potassium significantly. You can see this in studies like this one where it is clearly shown that the concentration of potassium in tissue is proportional to VPD more aggressively than to K concentration in nutrient solution.

The real fix to a problem like the problem above would be to lower the VPD of the environment – by reducing temperature or increasing relative humidity, depending on what's wrong – but choosing to just increase the amount of K in the nutrient solution would only lead to a minor response from the plant (because that's not the problem in this case). *If the grower makes an assumption and that assumption is wrong, then significant time would have been lost in the fixing of the problem and the leaf tissue analysis will reflect very limited progress.*

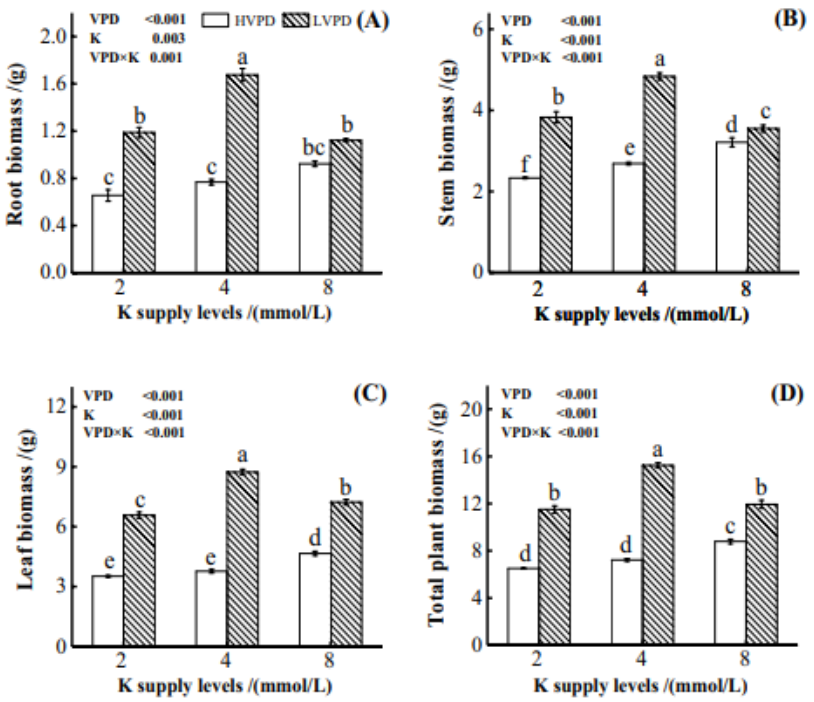

Fig. 5 Effects of vapor pressure deficit (VPD) and K supply on **a** root; **b** stem; **c** leaf; and **d** total plant biomass on day 30. Data represent the mean ± SE ($n = 5$). Different letters indicate significant differences as determined by Tukey's test ($P < 0.05$). A two-way ANOVA was performed to test the effects of VPD, K supply (K) and their interaction (VPD × K)

Image taken from [this study](), showing the relationship between VPD conditions and K

This is where foliar spraying comes into play. In order to "hedge our bets" in the fixing of a nutritional problem, we might want to increase the supply of the nutrient available to plant leaves by applying that nutrient to leaves directly while we figure out what is wrong with the environment or the nutrient solution. This will alleviate the issue because we will be delivering the nutrient directly to leaf tissue, regardless of what the actual root cause of the problem creating the blockage in nutrient transport is. That way, if we are wrong about the fix, we will already have made some progress in fixing the problem by delivering the nutrient that we're failing to transport where it is more strongly required.

Granted, there are a couple of caveats here. *The first is that we must have leaf tissue analysis so that we are sure about what needs to be applied (no guessing)*. The second is that we

still need to look into what the root cause is and solve the issue, otherwise the foliar spraying will eventually reach a limit and be unable to completely get the plants back to full health. Think of the foliar sprays as the CPR you can give your plants while the ambulance is on the way, the plants won't be able to survive from the CPR forever, but it will help them stay alive while the true solution for the problem arrives.

Table 3
Zinc and iron concentrations (μg g^{-1} on dry weight basis) in different parts of two cultivars of tomato plants grown at different levels of zinc with or without foliar application of zinc sulphate

Zn concentration in nutrient solution (μmol l^{-1})	Zn concentration sprayed to the leaves (mmol l^{-1})	'Blizzard'			'Liberto'		
		Leaves	Fruit	Root	Leaves	Fruit	Root
Zinc							
0.15	0.00	11 a[a]	16 a	46 a	16 a	14 a	43 a
	0.35	56 c	25 b	63 a	75 c	23 b	41 a
	3.50	541 d	32 c	65 ab	630 d	30 c	71 b
7.70	0.00	33 b	30 c	162 c	36 b	22 b	96 c
Iron							
0.15	0.00	122 d	87 c	2333 a	111 d	74 b	1521 a
	0.35	93 b	77 b	2335 a	84 b	80 bc	2364 b
	3.50	72 a	61 a	6187 c	73 a	60 a	4660 d
7.70	0.00	97 c	85 c	3561 b	98 c	84 c	3565 c

[a] Within each column, same letter indicates no significant difference between zinc treatments (HSD at 99%).

Table taken from this study showing how effective foliar applications of Zn can be in delivering the nutrient to leaves in tomato plants

To design a foliar spray to alleviate a deficiency, first read my post about some important considerations when using this technique. Second, make sure you start with lower concentrations, to prevent further stressing plants that might already be subjected to a significant degree of stress. Third, make sure you test the foliar spray on a small group of plants so that you know what the response of the plants will be before applying to the entire crop. Under some circumstances using this method might cause additional issues, so it's important to make sure the plants can take the spray before subjecting a larger number of plants to it. When doing a foliar spray to alleviate a deficiency I suggest carrying it

out only once a week initially and moving to two times per week if necessary until the root cause is fixed and the applications can be stopped.

If you are currently facing a nutrient deficiency problem and would like my help in formulating a foliar fertilizer for your specific case feel free to use the contact form or book an hour of consultation time so that we can further discuss your issue and help you fix your crop's condition.

Five things to consider when trying to copy commercial hydroponic nutrients

There are hundreds of different formulated hydroponic fertilizers out there and most of them are very expensive. Due to these very high costs, growers will often want to copy a set of hydroponic products they are very familiar with or a set of products that other growers – ideally growing under similar conditions – have had success with. However, the process of copying a commercial hydroponic nutrient with raw inputs is not as straightforward as many would like it to be and the procedure to do this accurately can be complicated due to both the nuances of the fertilizer industry and potential measures manufacturers might take to make reverse engineering of their products significantly harder. In this post I want to talk about five things you should consider before attempting to copy a hydroponic nutrient formulation, so that you can be very aware of the potential issues and problems you might find along the way.

The labels are often not accurate (enough). A fertilizer's

label contains the minimum guaranteed analysis of the fertilizer. Depending on the legislation, this usually means that the fertilizer must contain, at a minimum, this amount of every one of the specified nutrients, but there is no problem if the fertilizer contains *more* than what the label discloses. If a company is selling a fertilizer that has an NPK of 12-12-12 they can actually register that fertilizer as a 10-10-10 fertilizer and sell it as if it was a 10-10-10. The fertilizer will in reality be a 12-12-12, but the manufacturer can be sure that it will always be above the 10-10-10 specification. This is often not done out of malice, but out of the fact that the fabrication process itself might create a significant amount of variance within the composition of the actual fertilizer being produced and the manufacturer always wants to be above the minimum. This means that if you want to get the true mineral composition of the product, you'll need to send the actual fertilizer you want to copy to the lab. *Never rely on the label when copying a fertilizer.*

Label of a very popular hydroponic fertilizer. Trying to copy this fertilizer directly using this composition and "derived from" information, would lead to substantially higher costs, manufacturing problems and errors. This is common to a very large array of commercial hydroponic products.

Not everything that can be claimed is claimed. When a manufacturer decides to create a fertilizer product, it might decide to leave out a specific nutrient within the formulation that is there, but that they do not want to claim to prevent reverse engineering. This is often not illegal – you're getting more than what you paid for from the point of view of the regulators – but it does mean that you're going to be completely missing something if you just copy what the label says. This is a very common trick that is done with micronutrients, where a manufacturer will claim, for example, that the fertilizer has Fe and Mn, but will make no claims about Zn, B, Cu or Mo. A person copying the label would be

missing these nutrients, so their plants would end up dying from deficiencies.

The “derived from” is usually not what it’s derived from. Usually a hydroponic product will contain a list of the inputs that were “in theory” used for its fabrication. This will be a list of commonly available raw fertilizers, but more often than not, fertilizer manufacturers might include a product from which the composition might be derived, that is significantly more expensive than the raw inputs that the fertilizer is actually derived from or add unnecessary inputs to the list. A simple example would be a fertilizer that is made with potassium sulfate, magnesium sulfate, and monopotassium phosphate. The manufacturer might choose to say it’s derived from potassium sulfate, monomagnesium phosphate, potassium carbonate and magnesium sulfate. You can probably derive the same final composition from both salt mixes, but the monomagnesium phosphate is a very expensive input compared to the monopotassium phosphate and the potassium carbonate is unnecessary in this product and will generate pH issues. This is a very common trick, designed to make reverse engineering attempts more expensive and to difficult manufacturing for people who try to copy using this information.

Inputs with non-fertilizer components. A fertilizer can often have nutrient ratios that appear to be impossible to get to given the “derived from” section they have given. This often happens when there are inputs within the fertilizer that contain non-fertilizer components that are not reflected within the label, or even within an analysis of the nutrient solution. For example a manufacturer might decide to create a calcium supplement containing calcium nitrate and magnesium nitrate and then the label might say it has way more Ca than what is possible from just the calcium nitrate. This means there is another source of Ca present but, what is it? In this case, the manufacturer might be using something like calcium chloride, which they completely neglect to mention within the

label. However you should not make assumptions about what these things are, but actually perform an analysis to try to confirm your suspicions. Often assuming the “missing part” is something like calcium chloride can lead to you formulating something that is actually toxic to plants.

Additives that are not part of the mineral makeup. Many fertilizer formulations will also contain additives that do not have any mineral content and that therefore are completely avoided within the label. This is very problematic, since the effect of some hydroponic formulations might be largely related with some of this non-mineral content. The reason why a formulation might work significantly better than another of very similar nutrient composition might be the use of some additional substances within the formulation, such as undisclosed plant growth regulators, gibberellin inhibitors or other substances with very strong effects on plants. Even things as simple as non-ionic surfactants – which can significantly increase the wetting in media like rockwool – can make a big difference between two fertilizers with the same mineral composition. Knowing that these substances are there and copying them can be quite complicated and requires a lot of relatively expensive analysis to figure out.

As you can see, copying hydroponic nutrients is not just a matter of reproducing something that mimics what the label specifies (that would be very easy). It generally requires chemical analysis of the actual fertilizer to determine its mineral composition, judicious evaluation of the available raw inputs to evaluate which ones might be appropriate to reach the required composition and special consideration about the possibility of other additives that might be present within the product and the analysis to find out what these additives might be.

Five things you can learn from leaf tissue analysis

Lab results are incredibly useful in hydroponics, as they give us a quantitative view of what's going on within our crops. From the potential array of analysis that can be carried out, few give us as much information as leaf tissue analysis. Despite this fact, few growers ever routinely carry out this analysis, as it's often perceived as unnecessary unless problems are showing up within a crop. In this article I want to talk about five different pieces of information that leaf tissue analysis can give us that can be very useful to hydroponic growers, not only when problems are showing up within the plants but as a routine measurement carried out at several different points within a plant's growing cycle.

Are the plants facing bad vapor pressure deficit (VPD) conditions. Leaf tissue analysis can tell you whether environmental conditions are pushing the plants in the wrong direction by showing you how the ratios of elements like Ca/Mg and Ca/K are skewed. Whenever a flowering plant is grown under a hydroponic solution with a Ca/N close to 1 and the VPD of the environment is very high, the amount of Ca will tend to increase a lot relative to K. This is mainly because the transport of Ca ions is controlled in a bigger proportion by the vapor pressure deficit of the environment, so plants grown at high VPD values will tend to show high Ca in tissue. See this paper, where it is clearly shown how VPD is directly proportional to Ca in tissue. At lower Ca concentrations, the difference tends to be greater between high/low VPD values.

Calcium content (mmol kg⁻¹ dry matter) of leaf margin (m) and centre (c) at two calcium levels and two vapour pressure deficits

calcium		vpd (kPa) 0.75 (l/l)	0.43 (h/h)	mean
16%	m	367	277	322
	c	429	390	410
64%	m	783	689	736
	c	941	920	931
mean	m	575	483	529
	c	685	655	670

VPD strongly affects Ca in tissue. Results in cucumber at two different VPD and Ca concentration levels.

Is there any heavy metal contamination going on. Growing plants for human consumption that contain a significant amount of heavy metals is usually unacceptable. This means that the early detection of heavy metal accumulation is important. Leaf tissue analysis can offer some early insights into heavy metal accumulation within leaves, in order to protect growers from getting end-products that contain large amounts of heavy metals. A plant that contains a significant amount of heavy metals in leaves before the flowering stage is not completely lost, given that heavy metals can be significantly hard to move within plant tissue. If this is detected the problem can be dealt with and inputs can be analyzed to figure out where the heavy metals are coming from. Waiting for the end-product to get a heavy metal test can be a significant waste of valuable time.

Are things where they are supposed to be. One of the reasons why it’s important to carry out leaf tissue analysis routinely is that they can provide you with an idea of whether things are where they are supposed to be or not. Comparing leaf tissue analysis from a plant this crop cycle with plants from past crop cycles can give you an idea about whether things are progressing as planned or whether there are significant deviations from the past. This might be particularly important if changes are being tested or implemented and can provide an early warning about plant stress or issues that have to do with nutrient or environmental inputs.

How nutrients are changing as a function of time. When a plant shows clear visual symptoms of a nutrient deficiency, the problem is already well underway and damage to the crop's yields have already happened. In order to stay on top of things and make sure the plants are not experiencing any problems, leaf tissue analysis can help us assert whether plants are able to transport all ions adequately. Drops in elemental levels as a function of time in tissue can signal that a problem is imminently going to happen unless the situation is evaluated and measures are taken. Weekly leaf tissue analysis of a crop is a very powerful tool to track nutrient uptake and potential issues, especially if all the data is properly logged and comparisons can be easily drawn. The change in the amount of total solids within leaf tissue can also be tracked and can be used as a way to gauge whether a plant is being exposed to excessively dry conditions.

Are your silicon supplements doing their job. Silicon is very hard to transport by most plants – especially plants like tomatoes and other commercially grown flowering plants – so ensuring that the silicon you provide your plants is reaching tissue becomes important. Potassium silicate applications can often be useless if the are not being done correctly, as the life of silicate in solution is very short once the pH is reduced to the level generally used in hydroponics (5.8-6.2). At this point silicate turns into silicic acid, which readily polymerizes to form insoluble silica chains. Doing leaf tissue analysis looking for silicon generally reveals if the applications of this element are being successful and how successful the assimilation is through the entire crop cycle.

The above are some of the ways in which leaf tissue analysis can help you improve your crop results, although they are by no means the only uses for these quantitative results. In general, leaf tissue analysis should be treated like very valuable information and judicious records of all nutritional and environmental conditions should be kept in association

with them. A consistent history of leaf tissue analysis is extremely valuable in a growing facility, it helps avoid problems, carry out effective changes and quantify the real results of experimental interventions.

What is the ideal amount of media per plant in hydroponics?

When designing a hydroponic crop, the amount of media is a crucial variable to consider as it will determine a lot of the capital costs involved as well as play a key role in determining how irrigation is setup and how big the plants can get. However, how can we figure out what the ideal amount of media in a crop actually is? In today's post I am going to talk about the amount of media per plant in hydroponics, which factors play into deciding what size to use and what different choices will affect other aspects of your crop, such as irrigation frequencies and plant densities.

The first question we need to ask ourselves is, why do we need the media? The function of the media is to provide the root system with structural support and environmental protection. Plant roots cannot generally survive in the open air, so the media provides a cozy home where the roots can prosper and give the plant the water and nutrients it requires. The volume of media you provide will determine the size of this "safe space" and the actual media choice will determine how "safe" the space actually is. Plants require media to allow for enough air – because nutrient uptake requires oxygen – but it also requires the media to allow for some water retention in

order for water and nutrient uptake to actually take place. How optimum this oxygen/water/nutrient relationship is for a given media choice, will determine how big the media needs to be in order to sustain the plant.

Plants that are large also require a lot more water/nutrients, so the media and root system will need to provide enough absorption. A small amount of media will demand more from the root system – every cubic inch of root will need to work more efficiently – and it will also demand more from the irrigation scheduling, because ideal conditions will need to be more closely monitored since the root system will affect them quicker. You can sometimes see huge plants grown in 6′ x 6′ x 6′ rockwool cubes, these offer a small amount of volume (0.9 gallons), so to support a big plant, ideal media conditions need to be maintained all the time, which means very judicious monitoring of water content and frequent irrigation periods. As the cubes are irrigated the plant quickly uptakes water/nutrients, so the cube needs to be irrigated again. However, irrigate too frequently and oxygen content will drop and the plant will start to suffer as the root system won't be able to cope to maintain the plant's needs.

An evaluation of the media volume therefore requires an evaluation of other growing conditions. Consider when irrigation cycles will happen, how is monitoring going to be done and how does the media need to be managed to reach ideal conditions. *More media, means bigger costs but more forgiving*

root zone conditions, so less experienced growers can often do better with larger amounts of media. Novice growers will often fail when attempting to grow plants using less media, because they lack the experience to maintain the conditions needed for this to happen. *When growing larger plants, media volume per plant in the order of at least 5 gallons is recommended for people who don't have a lot of experience* or for conditions where close monitoring of the plants and automated irrigation is not going to be a choice.

Take this study on tomatoes grown using different volumes of media, the authors were able to achieve the same results with either 10L or 15L containers, but they got lower yields when moving to smaller container sizes. Someone starting out under these conditions would be better off erring on the higher side – using more media than less – in order to avoid reducing their yields due to insufficient volume being present for the irrigation conditions used. This might mean a higher expense, but a successful crop is always preferred to a crop with lower yields/failure. It's easier to plan for more media and then reduce it than the opposite.

If you are already growing and you want to lower the amount used per plant, you need to consider whether your media will allow for this or not. Only media that allows for significantly high water retention will allow for this to happen under intermittent irrigation, while media that do not retain water very well will only be able to do this under basically constant drip irrigation. If you're already doing 10+ irrigation cycles per day in intermittent irrigation with adequate dry-back between periods, then the media might already be reaching its limits in terms of what the root system can do in that volume. *Watching how the water content changes as a function of time will help you assess whether your media can be pushed harder or not*. If run-off EC/pH values are getting too extreme, this might also be a sign that you're reaching extreme regions in your media.

Remember that plants need to uptake the same amount of water/nutrient per unit of time to sustain growth. This means that a plant that requires 3 gallons of nutrient solution per day will still require this amount, regardless of whether the volume of the root zone is 1 gallon or 5 gallons. If you go from 5 gallons to 1 gallon then the drybacks will be significantly faster, so you need to adapt in terms of irrigation frequency.

In summary, media volume is a complex topic and requires a careful examination of different factors. Think about what ideal conditions are like for the media you chose and whether the irrigation system can provide enough oxygen/water/nutrients for the root zone in a given volume to fulfill the plants needs per day. *When in doubt, use more media*. If media reductions are being considered, remember that this will mean quicker dry back periods and therefore more frequent irrigation required. This means much higher stress for plants if irrigation cycles are missed or if problems in the root zone arise (for example problems with solution pH). Less media used means a more technical approach with more judicious monitoring will be required.

Getting all the data to evaluate a problem in a hydroponic crop

Problems are an inevitable part of being a hydroponics grower. Even experienced growers will sometimes face issues when moving between environments or plant species as things change and new challenges arise. A big part of being a good grower is

to be able to think about these obstacles, find out their causes and successfully respond to them. In this post I want to share with you some information about the data you should gather in order to properly diagnose a problem in your hydroponic crop. This is important as not having enough data often makes it impossible to figure out what's going on, while simple measurements can often give a very clear view of what's happening with the plants.

Take detailed, well documented pictures. What you see is a very important portion of what describes a plant's status and issues. The first thing you should do is document what you're seeing – take pictures of the plants showing the problem – and write down the symptoms you are observing. This documentation process should be organized, give each plant an ID, take pictures under natural light or white light of the new leaves, old leaves and root zones (if possible). Take pictures across different days showing the evolution of symptoms. Have all this information so that you can then better interpret what is going on. Also remember that symptoms do not necessarily mean deficiencies and deficiency symptoms does not necessarily mean more of a nutrient needs to be added to a nutrient solution (for example a P deficiency can show under low nutrient solution temperature even if P in the solution is actually very high).

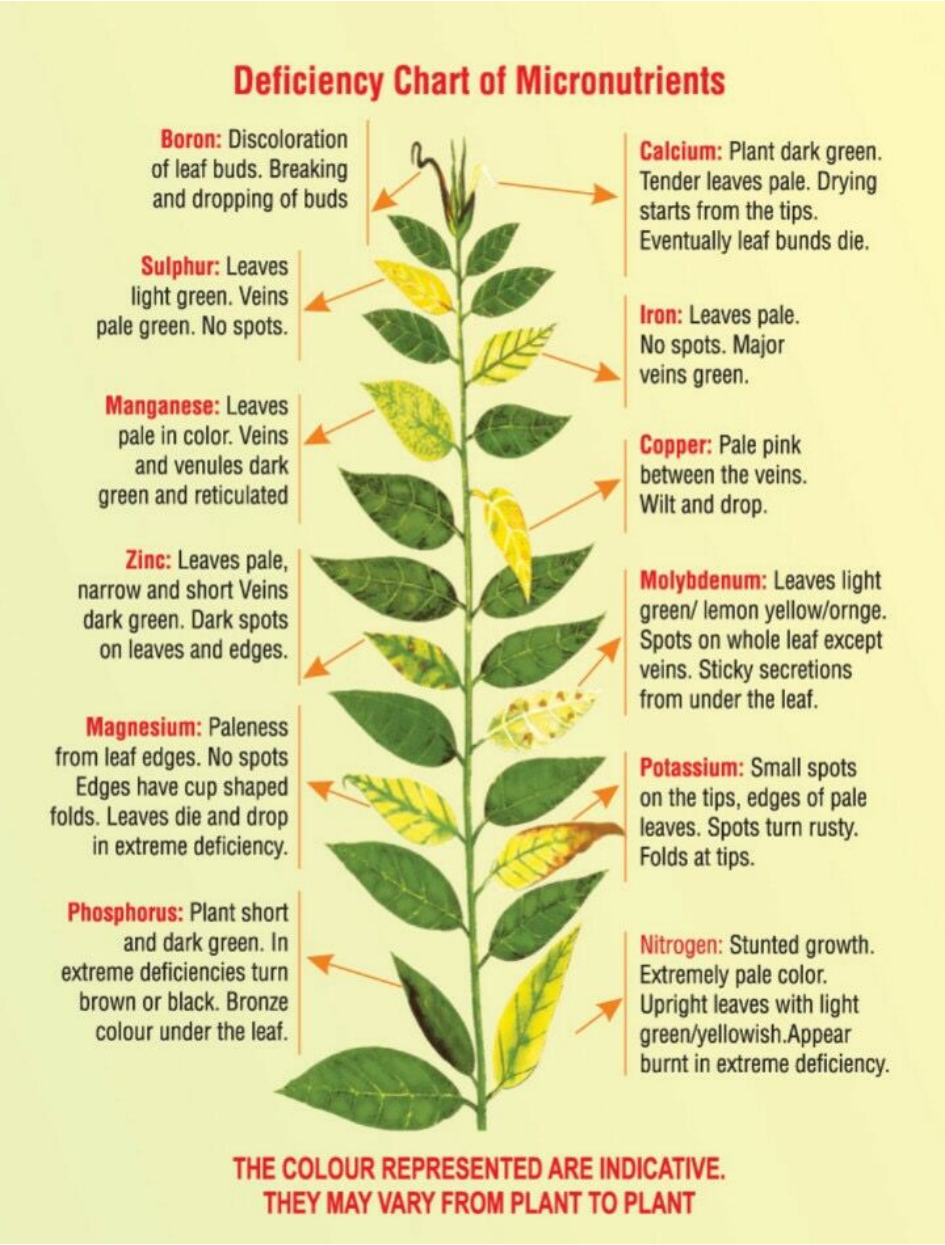

Taking detailed pictures can help assess whether a nutrient deficiency is present by gauging the changes in a plant as a function of time. However these should be confirmed with leaf tissue analysis as some of these symptoms can have causes not related with a nutrient deficiency.

Record all environmental data. When a problem happens, it is often related to the environment the plants are in. Having recorded data about the environment is a very important part of evaluating the issue and figuring out what went wrong here. Getting a good view about the environment usually involves having measurements for room temperature, temperature at canopy, relative humidity, carbon dioxide concentration, nutrient solution temperature, PPFD at canopy, and root zone temperature. All of this data should be recorded several times

per day as they are bound to change substantially between the light and dark periods.

Get nutrient solution analysis. Diagnosing a problem is all about having a complete view of what’s going on with the plants. The nutrient solution chemistry can often be a problem, even without the grower knowing a problem is brewing there. Sometimes nutrient solution manufacturers might have batches with larger errors than usual, or the input water might have been contaminated with something. There is also the potential of human error in the preparation of the solutions, which means that getting an actual check of the chemistry of the solution can be invaluable in determining what’s going on.

Get leaf tissue analysis. Even if the nutrient solution analysis does not reveal any problems, there are often issues with plants that are related with interactions between the environment and the solution that can go unnoticed in a chemical analysis of the solution itself. Doing a leaf tissue analysis will show whether there are any important nutrient uptake issues within the plant, which will provide a lot of information about where the problem actually is.

Critical nutrient foliar concentration for Blueberry (source: Penn State University)

Element	Deficient	Below Normal	Normal	Above Normal	Excessive
N (%)	1.65	1.7	1.9	2.1	>2.1
P (%)	0.05	0.06	0.1	0.18	>0.18
K (%)	0.35	0.4	0.55	0.65	>0.65
Ca (%)	0.35	0.4	0.6	0.8	>0.80
Mg (%)	0.18	0.2	0.25	0.3	>0.30
Mn (ppm)	45	50	250	500	>500
Fe (ppm)	65	70	200	300	>300
Cu (ppm)	4	5	11	15	>15
B (ppm)	29	30	40	50	>50
Zn (ppm)	14	15	25	30	>30

Critical nutrient foliar concentration for Brambles (source: Cornell University)

Element	Deficient	Below Normal	Normal	Above Normal	Excessive
N (%)	1.80	2.00	2.50	3.00	>3.00
P (%)	0.23	0.25	0.35	0.40	>0.40
K (%)	1.45	1.50	2.00	2.50	>2.50
Ca (%)	0.57	0.60	1.70	2.50	>2.50
Mg (%)	0.27	0.30	0.70	0.90	>0.90
Mn (ppm)	45	50	150	200	>200
Fe (ppm)	48	50	150	200	>200
Cu (ppm)	6	7	30	50	>50
B (ppm)	28	30	40	50	>50
Zn (ppm)	18	20	35	50	>50

Critical nutrient foliar concentration for Strawberries (source: Cornell University)

Element	Deficient	Below Normal	Normal	Above Normal	Excessive
N (%)	1.50	1.80	2.00	2.80	>2.80
P (%)	0.20	0.25	0.35	0.40	>0.40
K (%)	1.20	1.50	2.00	2.50	>2.50
Ca (%)	0.60	0.70	1.50	1.70	>1.70
Mg (%)	0.25	0.30	0.45	0.50	>0.50
Mn (ppm)	40	50	150	250	>250
Fe (ppm)	50	60	150	250	>250
Cu (ppm)	5	7	10	20	>20
B (ppm)	20	30	60	70	>70
Zn (ppm)	15	20	35	50	>50

Expected nutrient ranges for leaf composition of different species. Leaf tissue can often help tell whether there are some important abnormalities in progress and may help the grower assess which causes to look at.

Take well documented pictures of tissue samples using a microscope. A microscope can be important in determining what's going on with plants, because it can show developments in roots/tissue that cannot be seen with the naked eye. Microscopes can often reveal very small insects or fungal structures that would have otherwise gone unnoticed. For this reason, a microscope and the taking of microscopy images can

be of high value when dealing with a problem in a hydroponic crop.

With all the data mentioned above, most hydroponic crop problems will be much easier to diagnose. Some of the biggest failures in dealing with problems in hydroponic crops come from not gathering enough data and just guessing what the problem might be given how the plants look. Sadly plants can show similar responses to a wide variety of problems and – in the end – nothing replaces having the data to actually diagnose what's going on in order to deal with the issue appropriately. *Lacking an evidence-based picture is often the biggest difference between success in diagnosing/fixing an issue and failure or even worse problems caused by taking actions that have nothing to do with the real problem at hand.*

Building a DIY control infrastructure for a hydroponic crop: Part one

Controlling an entire hydroponic crop using electronics is not a trivial task. This includes everything from the automated control of things like relative humidity and ambient temperature, to other variables, such as lights, solution pH, conductivity and temperature. Many paid solutions exist in the market, but, in my experience, none of them offer enough flexibility to accommodate all potential environments, as all the ones I know are closed source and do not allow users to readily modify the firmware/software used to fit the user's particular needs. Through the past 5 years I have setup control infrastructures across several different crops and

have usually done so using an entirely DIY infrastructure that focuses on flexibility and power for the end user. In this post I want to talk about how this setup usually works and why I came to these design choices.

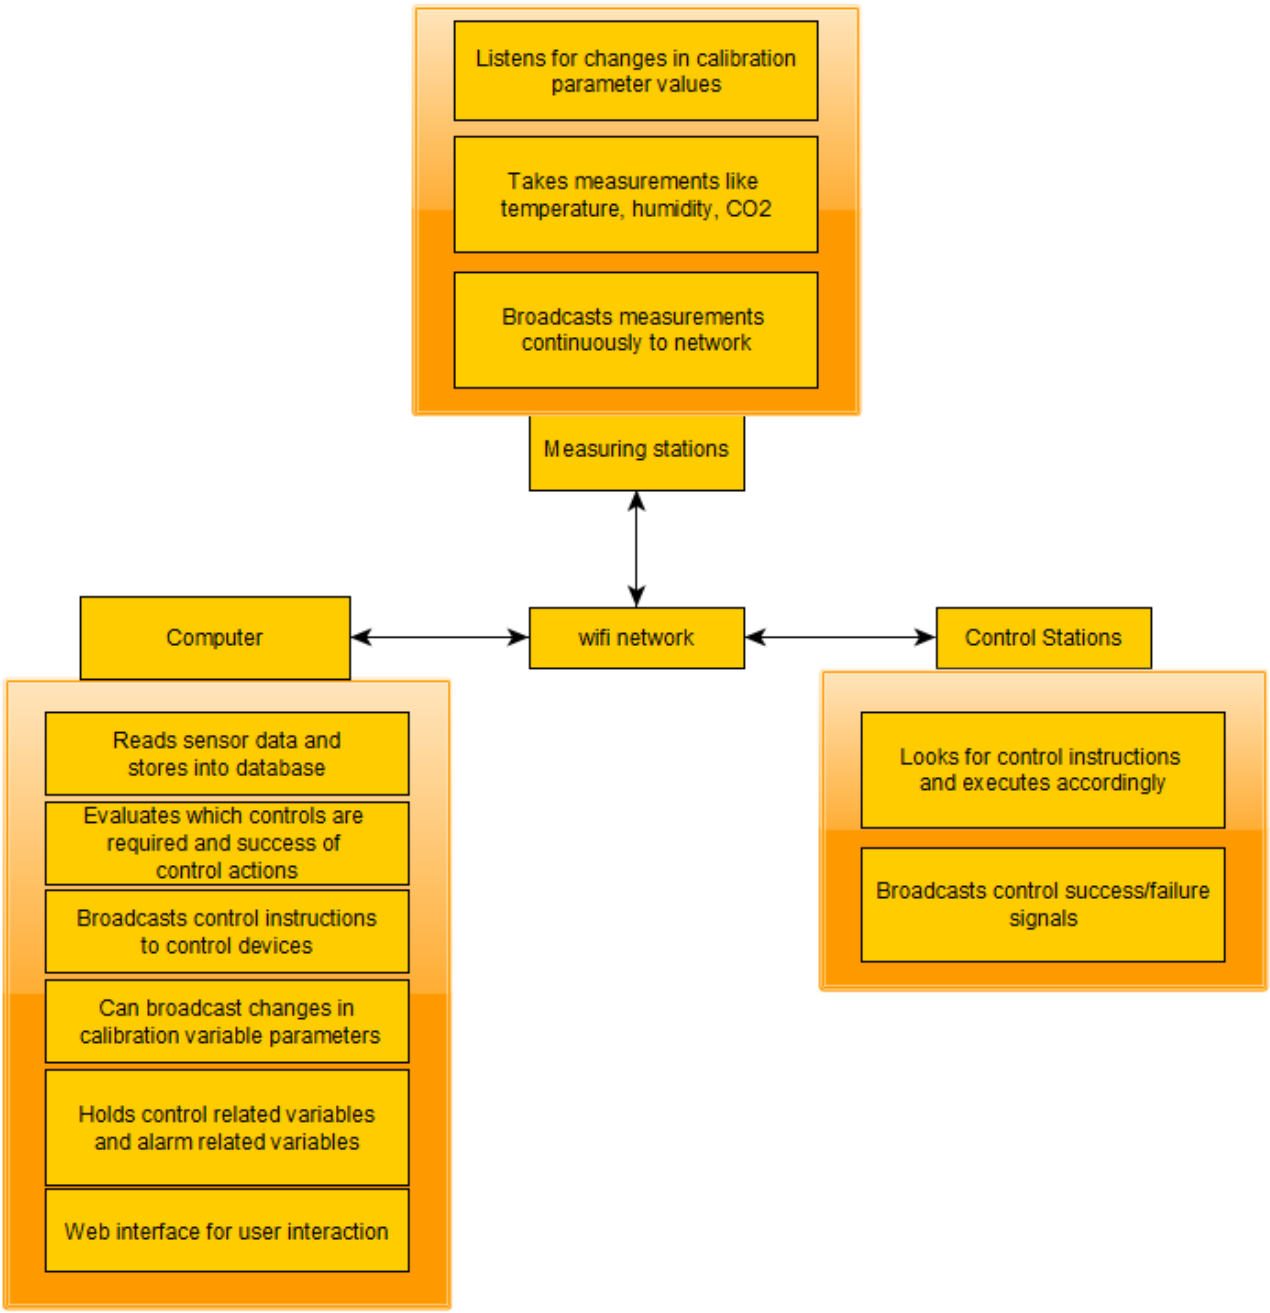

Usual network configuration I used to built electronic monitoring/control infrastructures for hydroponics

In general the infrastructure I setup relies on the use of wifi for the communication of the devices. This is because it's usually the easiest to setup, although it might not be the most power efficient or the most desirable in all cases. I

generally divide devices into three camps. There is a main device – which is usually a capable computer – which serves as the “central hub” for the entire setup. This computer contains the main database that stores all information about devices, sensor readings, calibration variables, alarms, etc and is in charge of deciding which control actions to take given the sensor reading it is receiving and the control devices it has access to. This central computer usually hosts a website as well, where the user can easily modify things, issue manual control actions, add new devices, set up alarms, etc. The computer can be duplicated as well, to prevent this from being an important point of failure. In several cases we have used a raspberry pi to play this role.

The second and third group of devices are usually Arduinos whose main role is to either take readings (measuring stations) or execute control actions (control stations). Some arduinos might actually serve both purposes as an arduino can often be fit with things like pH/EC probe readings as well as relays that control peristaltic pumps that are used to push pH up/down or nutrient solution into a solution tank. It is worth noting that the decision of what to do for control is never taken by any control station but all they do is interpret control messages from the computer and then try to execute those actions and then give back some response of what’s going on. Measuring stations, on the other hand, are only trusted with the task of taking some measurement from the environment and broadcasting it to the network, the only thing they might listen for are messages issued by the computer to modify their calibration, whenever this is required.

The arduino nano includes wifi capabilities and offers a very convenient low-power core for measuring stations that do not require high power to operate sensors

Measuring stations can be fully customized to have as many reading as the user desires and can be implemented to do all

sorts of things, from measuring temperature and humidity, to measuring air-flow, to measuring media water content. This allows for dozens of different temperature and humidity reading spots using different kinds of sensors, to monitoring things such as irrigation flow and solution ORP and dissolved oxygen values.

The entire setup relies on the use of the mosquitto (mqtt) protocol in order to have each device brodcast a specific topic with a specific reading that other devices can subscribe to. The computer will listen to all the devices it sees within its database and it is therefore easy for a new device to be added by a user, since it only requires the inclusion of the device into the database. The measure/control stations can subscribe to the specific topics their interested in for calibration or control actions and can act whenever they receive these messages. All the devices are automatically added to a web platform and alarms can easily be set for any of the measurements carried out by the measuring stations.

A big advantage of this approach is that control actions can be made as complex as the user desires. This includes doing things like implementing reinforcement learning based controls for things like temperature/humidity allowing the computer to use a wide array of measurements in order to take control actions, not relying solely on the measurement of one limited sensor to make these decisions. This allows a computer to use information such as outside temperature to decide how much air it wants to get into the facility for control, or how long it wants to turn on humidifiers in order to allow the desired level of humidity within a grow room.

With all this said, DIY control is definitely not the easiest route to take. Implementing something like the above will require the purchasing of a lot of different electronics – which are sometimes expensive depending on what the user wants – and does require a lot of time programming firmware and deploying software so that the desired outcome can be

achieved. With that said, the unparalleled level of control is often worth it and can be accompanied by substantial gains in the information available to the user, which often leads to improvements in yields and the significantly quicker catching of potentially important problems.

On the next part of this post, I will talk about some of the practical aspects of this project, such as which arduinos and sensors I usually use and how these are setup to communicate with the central computer. If you want to learn more about how I can help you set this up for your crop please feel free to contact me using the website's contact form.

Five common misconceptions around nutrient management in hydroponics

After many years of experience as a consultant in the hydroponic industry and interacting with dozens of different customers growing different plants with different systems, there are some common misconceptions that become apparent as time goes by. As a chemist, the ones I remember the most are related with the management of nutrient solution and the diagnosis and treatment of nutritional problems in plants. In today's post, I want to talk about some of these misconceptions and hopefully shine light into what the more accurate interpretation of these phenomena actually is.

The EC is increasing, my plants are not feeding! One of the concerns I most commonly address is that plants are "not feeding" because the electrical conductivity (EC) of the nutrient solution is not decreasing, but actually increasing

after the solution goes through the plants. Many growers think that EC measures nutrients in a solution, so if a plant feeds on nutrients, then the EC should naturally decrease as the plant feeds. This is wrong because the plant consumes both nutrients and water and EC is a proxy for nutrient concentration and not for the absolute amount of nutrients in the water. *As a plant feeds it will absorb both nutrients and water but significantly more water than nutrients*. Remember, plants are mostly made out of water and also use water to regulate temperature, humidity and nutrient uptake, so they will take way more water than nutrients, increasing the EC as they feed. As a plant grows larger it's nutrient and water demands grow, but the water demand grows significantly more than the mineral nutrient requirements, meaning the plant will progressively increase the EC more and more as it feeds more and more aggressively.

The plants are yellowing, there must be a nutrient deficiency. As soon as plants start to show signs of yellowing, a significant amount of growers will immediately look and try to interpret this as a sign that there is some form of nutritional deficiency. Most that subscribe to this belief will look for pictures of deficiencies online and do their best to match what they see with a deficiency and then proceed to supplement the solution with some fertilizer that contains

the “missing element”. More often than not, this is actually not caused by the composition of the solution at all but by some environmental factor that is not being properly managed. In run-to-waste systems this is most commonly related with a significant pH drift in the media – reason why it is always necessary to measure pH/EC of the run-off – but it can also be related to unnecessarily harsh VPD conditions or even a lack of enough air circulation. I would say that 5/10 times, problems with the plants have virtually nothing to do with the nutrient solution at hand.

If you want more X, then increase X in the nutrient solution. The relationships between the concentration of elements in a solution and the concentration of nutrients in plant tissue is not linear. Sometimes, increasing the concentration of an element in solution can actually lead to *less* of that nutrient being present within plant tissue. An example of this can be phosphorous, a plant can suffer from a phosphorous deficiency due to the formation of insoluble iron phosphate compounds in tissue that appear when the concentration of these two elements goes above some threshold. As more of either is added, more of these insoluble compounds are formed and less of P and Fe actually gets to the plant. Another example can be Ca, where the amount of Ca in tissue is more dependent on VPD than on the concentration of Ca in solution, changing the VPD by 20% will affect Ca in tissue significantly more than adding 20% more Ca to the solution in some plant species. In these cases you might add 20% more Ca but your VPD drops 20% and you actually see a decrease of Ca in tissue. *Sadly nutrient dynamics are not simple and often a more holistic picture needs to be used to approach nutritional management!*

Plants need aggressively more phosphorous when they flower. Most commonly used fertilizers in soil tend to have higher P/K values when they target “flowers”, this is because, in soil, phosphorous is not highly available and the supplementation of highly available phosphorous during flower can be very useful

to plants. However, flowering plants in hydroponics always have access to significant amounts of soluble P and most actually do not require an increase from this base level when they go into their flowering periods. Many commercial hydroponic solutions used for tomatoes will – for example – keep their P values at 50 ppm through the entire growing period, only increasing K during the flowering period, but not P. Experiments across various commercially grown flowering species have shown that levels in the 50-65ppm range are ideal for many plants during their entire life cycle, this matches the experience of growers in the horticultural hydroponic industry.

There is a perfect nutrient solution. Many growers go on a "holy grail" quest to find the "perfect" nutrient solution that will give them the absolutely best yields. Many commercial fertilizer producers also call me asking to formulate "the best possible formulation" to grow a given type of plant or – even worse – to grow a wide variety of plants. The truth is that the ideal solution to feed a plant will depend on the genetics, the environment, the irrigation system, the growing media, etc. Due to the large amount of variability between growing setups, plant genetics and growing methodologies, more often than not, the nutrient optimization process needs to be carried out for every grower. Don't get me wrong, a base formulation will probably get you 80% of the way to your maximum potential yields – nutrient solutions are not miracle generators, they are just food – but conquering that final 20% will require a lot of additional effort that will most likely be limited to your particular conditions. This is because most environments are limited by different factors and using the nutrient solution to help overcome some of these limitations will modify the solution in a way that's probably detrimental for other environments.

I hope the above misconceptions show that the world of nutrient solutions and plant management is not so simple and

that there is a lot that goes into understanding how nutrients interact within a plant and how a given growing environment needs to be modified in order to improve crop results. My goal is to help you expand your knowledge about hydroponics and better reach your goals by overcoming some of these misconceptions and tackling some of the true problems within your hydroponic crops.

Five tips to successfully manage your nutrient solution in a recirculating hydroponic setup

Although a significant portion of hydroponic growers use run-to-waste setups – where the nutrient solution is ran through plants and then lost – the industry is now moving towards the implementation of recirculating hydroponic systems in order to reduce both water usage and the unnecessary dumping of fertilizers into sewage systems. A recirculating setup has many advantages and can provide better yields than run-to-waste setups, provided the solution is properly managed and changed through the growing cycle. In this post I'm going to talk about five tips that will help you successfully manage your nutrient solution when using this type of system.

Ensure the volume of the reservoir is at least 10x the volume necessary for a single irrigation. The total volume of a reservoir is key in a recirculating setup because you want the bulk of the solution to be unaffected by whatever nutritional changes are caused by the plants during each feeding. This

means that you want most of the solution to be inside your tanks and not inside the media when every irrigation is done. A simple rule of thumb is to make the volume of your initial reservoir at least 10x the volume that it would take to carry out a single irrigation of your entire crop. If you do this the water and nutrient absorption effects will happen slowly and will give you time to manage your solution without any harm coming to the plants.

A recirculating hydroponic tomato system using dutch buckets

Circulate your solution until your pH and EC are constant. After an irrigation cycle starts, the solution will first mix with the remnants of the last irrigation cycle within the media, which will make the pH and EC of the return different from those of the main tank. In order to ensure that the plant's root system is being subjected to the desired nutrient concentrations, make sure you carry out the recirculating process until the EC and pH of the tank remains constant and matches the return pH and EC. Once this happens you know that the conditions within the media have now been equalized with the larger body of solution and you can stop the irrigation process. Constant monitoring of the pH and EC within the tank are therefore necessary within this type of setup.

Add water and *not* nutrients when the EC increases with every irrigation. In a normal recirculating setup the EC of the solution in the main tank will tend to increase with every irrigation while the total volume of the solution will decrease. This happens because healthy plants *always* absorb more water than nutrients, which means any measure that's proportional to concentration – such as the EC – will tend to increase as the amount of water goes down. You want to add enough water to bring the EC down to the desired EC but you do not want to add nutrients with this water and this would increase the EC or contribute to nutrient imbalances within the solution. Note that you will need to add *less* water than

the amount that was absorbed by the plants, because the plants also take some nutrients with them, meaning that the amount of water needed to reestablish the EC to what it was before will be lower. If an initial solution has 1000 gallons, the volume might go down to 950 gallons on the first irrigation but you might only need to add 20 gallons to bring it back to the original EC, making the total in the end around 970 gallons. *Make sure the pH of the tank is also corrected after every irrigation and water addition*.

Replenish water with nutrients when volume is down 40%, use this as an opportunity to shift the solution. As discussed in the last tip, the volume of solution will go down with time, even if some water is added to return to the original EC. At some point more than 40% of the volume will have been spent and it is at this point where you should fill the tank back to its full volume with water plus nutrients. You can also use this opportunity to change the nutrient ratios and skew them in the direction that you want your plants to follow nutritionally. For example in a flowering crop it is common to increase the amount of potassium during the blooming stages of the plant, so this can be done as nutrient solution is replenished after it's consumed by the plants. *Note that this process cannot be carried out indefinitely because both nutrient imbalances and plant exudates will accumulate within the main solution*. Most recirculating crops will fully change the solution every 3-4 weeks to avoid these problems although the life of the solution can be extended further when chemical analysis is done – to customize nutrient replenishing – and adequate filtering is implemented to remove substances contributed by plants.

Add in-line UV filters and carbon filters. It is fundamental to ensure no microorganisms contaminate your nutrient solution. For this reason, online UV-filters are necessary to keep the nutrient solution as sterile as possible. Carbon filters are also very useful as they remove plant exudates

that can contaminate the solution and cause problems within the crop itself. Many of these exudates are food for microorganisms, others are plant hormones that might cause unwanted responses in the plants. However both carbon filtration and UV filters can cause some issues – such as the destruction and adsorption of heavy metal chelates – so it is important to use chelates that are more resistant to UV and have less affinity for carbon filters to alleviate these problems.

There is certainly a lot more to the management of recirculating hydroponic solution than what I have detailed above, *it is important to note that some of these tips are simplifications and much better tailor-made solutions are possible with a proper analysis of each situation*. These are just some simple tips to hopefully make your change towards the use of recirculating systems a lot easier and should greatly increase your chances of success in the world of recirculating hydroponic setups.

A new conductivity model in HydroBuddy

On my previous post you can read about how I ran experiments to develop a conductivity model using empirical data in order to improve our ability to predict EC values from the concentration of individual nutrients in a hydroponic nutrient solution. In this post I will now talk about how this was finally implemented in HydroBuddy, what form it took and what kind of result can be expected from it. The implementation discussed in this post has already been updated to the HydroBuddy github along with all the experimental data used to

derive this empirical EC model.

Given the amount of data and the nature of the problem at hand, the easiest and most accurate way to build a model was to use a simple linear regression algorithm. As previously shown this model was able to give great results within the data, even when performing random training and testing splits. I have added a [jupyter notebook] to the github repository, along with all the data we measured in order to allow you to see how all the calculations were done, how the model was created and the sort of accuracy the model got within the set of experimental results. You can also play with this notebook to develop your own models or analyse the data any further if you wish. You can also try to reproduce our experiments and help verify our results. The linear model was translated into FreePascal and added to HydroBuddy although the program still retains the ability to estimate conductivity using the previously available LMC based model.

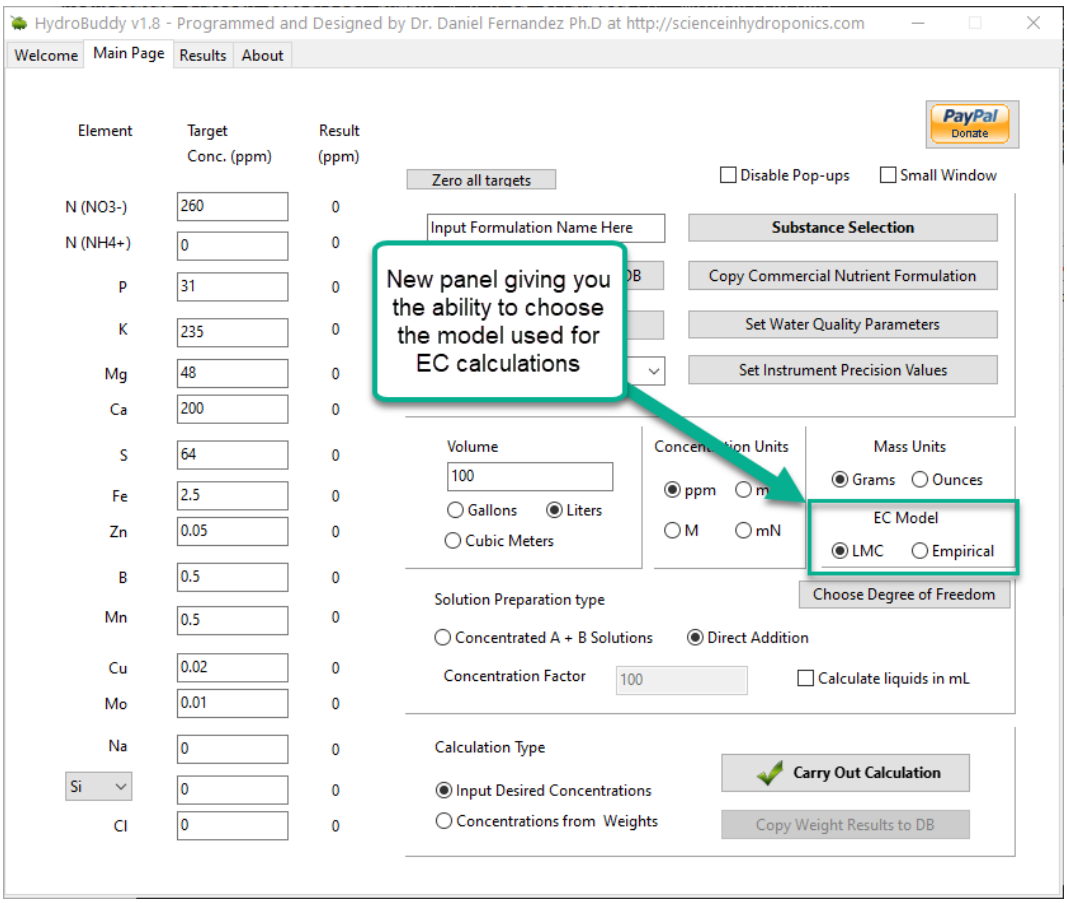

New hydrobuddy implementation now including the ability to

choose between LMC and empirical EC models.

The fact that we were able to create a model to accurately determine conductivity within this experimental space does not mean that this model will work to magically determine the conductivity of any hydroponic formulation. These experiments were designed using five salts – calcium ammonium nitrate, ammonium sulfate, potassium sulfate, magnesium sulfate and monopotassium phosphate – which means that although our model is able to greatly describe conductivity in this space, the model is likely to run into trouble when attempting to describe a space that deviates too strongly from the one described above. This will be most evident whenever there are some cations or anions that are not present at all within these experiments. For example when silicates, chlorides or other such salts are used or when strong acids or bases are added to the solution.

Another important issue is the way these ions are paired. In our experimental process the concentration of Ca and N as nitrate always increased at the same time, meaning that the linear model implicitly carries this assumption. A setup were magnesium nitrate or potassium nitrate are used as well, will contain deviations from the current model that it is likely not very well prepared to deal with. A similar problem might happen when salts such as ammonium monobasic phosphate are used, since our model only contained a single example of a phosphate salt (monopotassium phosphate). While it is not easy to predict how much accuracy will be lost in these cases, we do expect the model to be significantly more inaccurate as other salts are used.

Additionally, our experimental setup did not contain any corrections of pH values, so the conductivity values described include a pH drift related with the amount of acid contributed by the potassium monobasic phosphate, which was not neutralized by a base. This will also cause differences with conductivity, if the conductivity is measured after the pH of

the solution is corrected to the proper range used within the hydroponic process. Although at the concentration values used in hydroponics this should not be a big issue, it is still something worth considering.

As I mentioned above, the model is already implemented within the github repository – if you want to compile the program yourself – but the binaries won't be updated to v1.8 until later this week. I look forward to your feedback about the model and hope it can help – at least some of you – to dramatically improve the estimations of conductivity of your hydroponic nutrient solutions.

Keeping plants short: Using day/night temperature differences (DIF)

In this article series about "keeping plants short", we have explored the reasons why short plants are desirable and how this can be achieved using gibberellin inhibitors. However this is not the only effective way to control plant height and several other ways – some using no chemical means – can be used to keep plants short. In this article I will be talking about the use of day/night temperature differences in order to control plant height, what the research about this says and how it can be effectively applied by growers to achieve shorter plants.

The idea of using day/night temperature differences to control plant height can be traced back to the late eighties and some research done by people at Michigan State University (1). This research in easter lilies showed how plants grown at a

constant night temperature (68F, 20C) but subjected to even lower day temperatures or simply drops in early morning temperature could grow drastically shorter. The results surprisingly showed that a 14F temperature drop during the beginning of the day – first two hours – could actually cause the plants to receive the same effect as if the day temperature was lower during the entire day, yet the plants remained highly productive. This technique of reducing temperature during a few hours during the way was referred from this point on as “DIF”.

Taken from this 1986 article.

Experimenters then began testing across other plant species and found the results to be mixed. In this paper (2) chrysanthemum, poinsettia, begonia and kalanchoe were all tested in a -6 C DIF experiment and while chrysanthemum and begonia both responded in the expected manner, the kalanchoe actually responded in the opposite way and showed stronger elongation of the flower stems. In all of these cases the use of growth regulators – gibberellin inhibitors – was still needed to ensure plants stayed at the required height. This was one of the first studies that pointed to the fact that the DIF technique is tremendously crop dependent.

During the nineties it was established that DIF did work for several common crops, for example cucumber and tomatoes showed to be sensitive to the DIF effect, particularly when the first two hours of the day showed a temperature drop. In this case the effect reduced both the inter-node distance and was directly proportional to the difference in temperature. It was also established that some plants prefer pulses of cold temperature during the end of the day, while others might prefer this pulses even in the middle of the night. It was also showed that strong negative DIF treatments caused negative effect related with a reduction in chlorophyll production, resulting sometimes in even plants showing signs of chlorosis. Plants grown in negative DIF were also shown to have lower total dry weights although depending on the magnitude of the DIF, limited or sometimes even positive effects on weight and yields could be seen. You can read more about the above in this review from the late nineties which also contains a lot of literature references for early DIF research (3).

Stem elongation effects of DIF in peas, taken from this article

More recent research from 2013 on tomatoes, eggplant and sweet pepper (4) has shown that a variety of different day/night temperature treatments can be effective in minimizing vegetative growth while having a limited effect on yields. In this case the strongest effect was seen for a 15C/25C day/night temperature cycle. This paper also looked at nutrient absorption and noticed that Ca/Mg/K concentrations were actually highest in the 15C/25C temperature treatment, which suggests that changing the day/night temperature did not adversely affect nutrient absorption. The conclusions of this research were then reproduced and matched when looking at cucumber, melon and watermelon (5). However other research using positive as well as negative differences in temperatures has shown that for tomatoes, the ideal day/night temperature

difference is positive and in the order of +6C if yields and plant growth are given the highest priority (6).

The DIF method has shown to be a reliable way to control the height and vegetative growth of many different plant species, although for some it does not work very well. In general the researchers who apply negative DIF methods for reducing stretch tend to have the most success with a -10C (-18F) increase in night over day temperatures. If testing on a new plant the recommendation would be to start with a 2 hour temperature drop in the day temperature of this magnitude for the first 2 hours of light – starting the drop 30 minutes before sunrise – and see which results you can get. This is likely going to be the cheapest in terms of both climate control and potential disruptions in yields caused by this technique.