


Creating a pH/EC wireless sensing station for MyCodo using an Arduino MKR Wifi 1010

There are multiple open-source projects available online for the creation of pH/EC sensing stations for hydroponics. However, all of the ones I have found use a single Arduino or Raspberry Pi to perform the measurements and store any data, making them unsuitable for applications where more flexibility is needed. For example, a facility using multiple different reservoir tanks for nutrient storage might require multiple pH/EC sensing stations, and single-board wired setups would be unable to accommodate this without a lot of additional development. In this post, I am going to show you a simple pH/EC sensing station I built with an Arduino MKR Wifi 1010 that can communicate with a MyCodo server using the MQTT protocol. Multiple sensing stations could be built and all of them can communicate with the same MyCodo server.



My Arduino MKR wifi 1010 based sensing station, using uFire pH and EC boards in a small project box.

This project makes use of the small pH/EC boards provided by uFire, which have a lower cost compared to those provided by companies like Atlas, but do have adequate electrical isolation to avoid problems in readings when multiple electrodes are put in the same solution. This is a substantial improvement over other low-cost boards where using multiple probes can cause heavy electrical noise and interference. In order to build this project you will require the following materials:

Note, some of the links below are amazon affiliate links. This means that I get a small commission if you purchase through these links at absolutely no extra cost to you. The links to other websites are not affiliate links.

1. [Arduino MKR Wifi 1010](#)
2. [uFire pH probe](#)
3. [uFire EC probe](#)
4. [A rugged pH probe with a VNC connector](#)
5. [An rugged EC probe with a VNC connector](#)
6. [Two Qwiic-to-Qwiic connectors](#)

7. [One Qwiic-to-male connector](#)
8. A project box to put everything inside (optional)
9. [A micro USB cable](#)

The code for the project is shown below:

```
#include <uFire_EC.h>
#include <uFire_pH.h>
#include <WiFiNINA.h>
#include <ArduinoMqttClient.h>

#define SECRET_SSID "ENTER WIFI SSID HERE"
#define SECRET_PASS "ENTER WIFI PASSWORD HERE"

//calibration solutions used
#define PH_HIGH_SOLUTION_PH 7.0
#define PH_LOW_SOLUTION_PH 4.0
#define EC_HIGH_SOLUTION_EC 10.0
#define EC_LOW_SOLUTION_EC 1.0
#define CALIBRATION_TEMP 20.0

// topics for the mqtt sensors
// Make sure all stations have different topics
#define EC_TOPIC "EC1"
#define PH_TOPIC "PH1"
#define CALIB_TOPIC "CALIB1"
#define MQTT_BROKER "ENTER MQTT SERVER IP HERE"
#define MQTT_PORT 1883

int status = WL_IDLE_STATUS; // the Wifi radio's status
String message;

uFire_pH ph;
uFire_EC ec;
WiFiClient wifiClient;
MqttClient mqttClient(wifiClient);

void check_connection()
{
  if (!mqttClient.connected()) {
    WiFi.end();
    status = WiFi.begin(SECRET_SSID, SECRET_PASS);
```

```

    delay(10000);
    if (!mqttClient.connect(MQTT_BROKER, MQTT_PORT)) {
        Serial.print("MQTT connection failed! Error code = ");
        Serial.println(mqttClient.connectError());
        delay(100);
    }
    mqttClient.subscribe(CALIB_TOPIC);
}

void setup()
{
    Serial.begin(9600);
    while (!Serial);

    // connect to wifi and mqtt broker
    check_connection();
    // coorectly initialize the uFire sensors
    // note the Wire.begin() statement is critical
    Wire.begin();
    ec.begin();
    ph.begin();
}

void loop()
{
    // mqtt keep alive
    mqttClient.poll();

    // read messages
    message = "";
    while (mqttClient.available()) {
        message += (char)mqttClient.read();
    }

    // execute calibration if requested
    Serial.println(message);
        if (message == "EC1_HIGH")
ec.calibrateProbeHigh(EC_HIGH_SOLUTION_EC, CALIBRATION_TEMP);
        if (message == "EC1_LOW")
ec.calibrateProbeLow(EC_LOW_SOLUTION_EC, CALIBRATION_TEMP);
}

```

```

        if (message == "PH1_HIGH")
ph.calibrateProbeHigh(PH_HIGH_SOLUTION_PH);
        if (message == "PH1_LOW")
ph.calibrateProbeLow(PH_LOW_SOLUTION_PH);

// Measure EC
ec.measureEC();
Serial.println((String) "mS/cm: " + ec.mS);

// Measure pH
ph.measurepH();
Serial.println((String) "pH: " + ph.pH);

// Ensure the wifi and mqtt connections are alive
check_connection();

// post EC to MQTT server
mqttClient.beginMessage(EC_TOPIC);
mqttClient.print(ec.mS);
mqttClient.endMessage();

// post pH to MQTT server
mqttClient.beginMessage(PH_TOPIC);
mqttClient.print(ph.pH);
mqttClient.endMessage();

// ensure sensors are not probed too frequently
delay(1000);
}

```

Once you get all the materials you should first assemble the components. Connect the pH and EC board together using the Qwiic-to-Qwiic connector, then use the Qwiic-to-male connector to hook up one of these boards to the Arduino (doesn't matter which one). Connect the black cable to ground, red cable to 5V, blue cable to SDA, and yellow cable to SCL. Set up your board according to the instructions in the [Arduino MKR wifi 1010 getting started page](#), modify the code above to properly include information about your wifi network, calibration solutions, and MQTT server, then upload the code. The Arduino

will connect to your Wifi and MQTT servers and automatically reconnect when there are connection issues.

The above code will also post the readings of the pH and EC sensors to topics PH1 and EC1 respectively if you add an input in MyCodo to capture these readings you should be able to store them and take control actions using the MyCodo interface. Additionally, the Arduino code will respond to calibration requests published to the topic "CALIB1". For example, if you want to calibrate your EC sensor with a two-point calibration method with a standard solution with an EC of 10mS/cm, you would put the electrode in the calibration solution, then send the message "EC1_HIGH" to the CALIB1 topic and the Arduino will perform the task as requested. The code assumes you will want to do 2 point calibrations for both EC and pH, with the calibration events triggered by EC1_HIGH, EC1_LOW, PH1_HIGH, and PH1_LOW. Note that the definition of the EC and pH values of the calibration solutions should be changed to the solutions you will be using within the code. The high/low values in the code, as is, are 10mS/cm|1mS/cm for EC and 7|4 for pH.

A simple cheatsheet for macro nutrient additions in hydroponics

In hydroponic growing, we are often faced with the need to adjust the nutrient concentrations of a fertilizer reservoir or foliar spray directly, in order to increase the quantity of some nutrient by a specific amount. Although you can use a program like [HydroBuddy](#) in order to quickly calculate these

values, it is often the case that these calculations need to be done in the field or in a growing environment, and a computer to calculate things is not at hand. For this reason, I have created a small “cheat sheet” that you can use in order to figure out the amounts of salts that you would need to add to a solution to increase any of the macronutrients by 10 ppm.

Salt Name	ppm	Element	ppm	Element	g/L	g/gal
Calcium nitrate (ag grade)	10	N (NO3-)	13.19	Ca	0.0694	0.2629
MAP	10	N (NH4+)	22.1	P	0.0821	0.3108
Ammonium Sulfate	10	N (NH4+)	11.4	S	0.0472	0.1785
Gypsum	10	Ca	7.99	S	0.0430	0.1626
Calcium Chloride	10	Ca	17.69	Cl	0.0277	0.1048
Magnesium Nitrate Hexahydrate	10	N (NO3-)	8.67	Mg	0.0915	0.3463
Epsom Salt	10	Mg	13.19	S	0.1014	0.3839
Magnesium Chloride	10	Mg	29.16	Cl	0.0392	0.1483
AgSil 16H	10	Si	10.9	K	0.0411	0.1554
MKP	10	P	12.62	K	0.0439	0.1663
Potassium Nitrate	10	N (NO3-)	27.87	K	0.0730	0.2763
Potassium Sulfate	10	K	4.10	S	0.0223	0.0844
Potassium Chloride	10	K	9.067	Cl	0.0191	0.0722

Cheatsheet for macronutrient additions in hydroponics

With the above cheatsheet, you can quickly evaluate some of the most common options you would have to increase all the different macronutrients in a hydroponic or foliar solution by 10 ppm and which secondary elemental contributions you would get from these additions. For example, if you add 0.0694g/L of Calcium Nitrate, this would add 10ppm of Nitrogen as nitrate plus 13.19ppm of Calcium. Careful consideration of secondary contributions need to be taken into account, especially when using salts that contain elements that can be toxic, such as

chlorides.


How to make your own stabilized mono-silicic acid for use in hydroponics

Please follow [this link](#), for an updated an easier process for the synthesis of mono/ortho-silicic acid.

During the past several years, there have been a lot of “mono-silicic acid” products being marketed for their use in hydroponic culture. These differ from the traditional potassium silicate based products in that they are very acidic in their concentrated form and are stable in solution for longer periods of time at the pH values used in hydroponics. While a hydroponic nutrient solution that has potassium silicate added to it and the pH adjusted to 5.5-6.5 will generally see extensive polymerization of the silicon-containing molecules within 24 hours, these stabilized mono-silicic acid products will be stable for far longer periods of time. They are therefore ideal for use in recirculating systems, where potassium silicate additions can be less effective.

If you watched my [youtube video](#) on silicon in hydroponics, you’ll know that the most common way to produce these stabilized products is quite complicated and involves the use of silicon chloride and very specific reaction conditions. These are unavailable to hydroponic growers, reason why it is not easy to produce a home-made version of choline stabilized ortho-silicic acid (ch-OSA). However, choline is not the only reagent that can be used to stabilize silicic acid in

solution, and research in industry has shown us that it is actually possible to form stabilized silicic acid products starting from potassium silicate itself.



Model representation of orthosilicic acid, also called mono-silicic acid.

[This patent](#) describes how to prepare mono-silicic acid solutions that are stabilized by carnitine and several other additives, in the region from 0.01 to 8% silicic acid by weight. The great thing about this process is that we can start from potassium silicate, which is readily available. The concentration of Si obtained will be significantly lower than what is possible when generating ch-OSA – where solutions can reach 40% mono-silicic acid – but the fact that we can prepare it from readily available materials might compensate for this to some extent. **It is also worth noting that this process comes from an unexpired patent, so it should not be used commercially without licensing the technology from the owner of the intellectual property.**

Extrapolating from the contents of the patent and the examples given, we can come up with a process to brew our own mono-

silicic acid at an 8% concentration. Here are the things you will need:

Note the amazon links below are affiliate links. This means that, if you choose to purchase through these links, I get a commission for your purchase, at no extra cost to you.

1. [Potassium silicate \(at least 32% K as K₂O\)](#)
2. [Carnitine hydrochloride](#)
3. [Phosphoric acid \(85%\)](#)
4. [Propylene glycol](#)
5. Distilled water
6. [Scale to weight the materials \(precision of at least +/- 0.1g, max at least 500g\)](#)

To prepare around 425g of stabilized mono-silicic acid, you could follow this process.

Before continuing please make sure you understand what you're doing. Wear adequate eye and body protection, carry this out in a place with enough ventilation and make sure you read the material safety data sheet (MSDS) of all the materials used. These instructions are provided for educational purposes only, follow them at *your own risk*.

1. Add 10g of carnitine hydrochloride to a clean 1000mL beaker
2. Add 65g of distilled water to the mix.
3. Stir until all the carnitine hydrochloride dissolves
4. Add 10g of propylene glycol.
5. Add 240g of 85% phosphoric acid.
6. Put the mixture on an ice bath with ample ice.
7. Wait for 15 minutes, so that the mixture cools down.
8. During the course of an hour, *slowly* add 125g of potassium silicate to the mixture with constant stirring. Add more ice to the ice bath if needed to keep the solution cool. Note that predissolving the silicate in 150mL of distilled water and adding it as a liquid

makes this process easier, although KOH additions might be required to complete its dissolution.

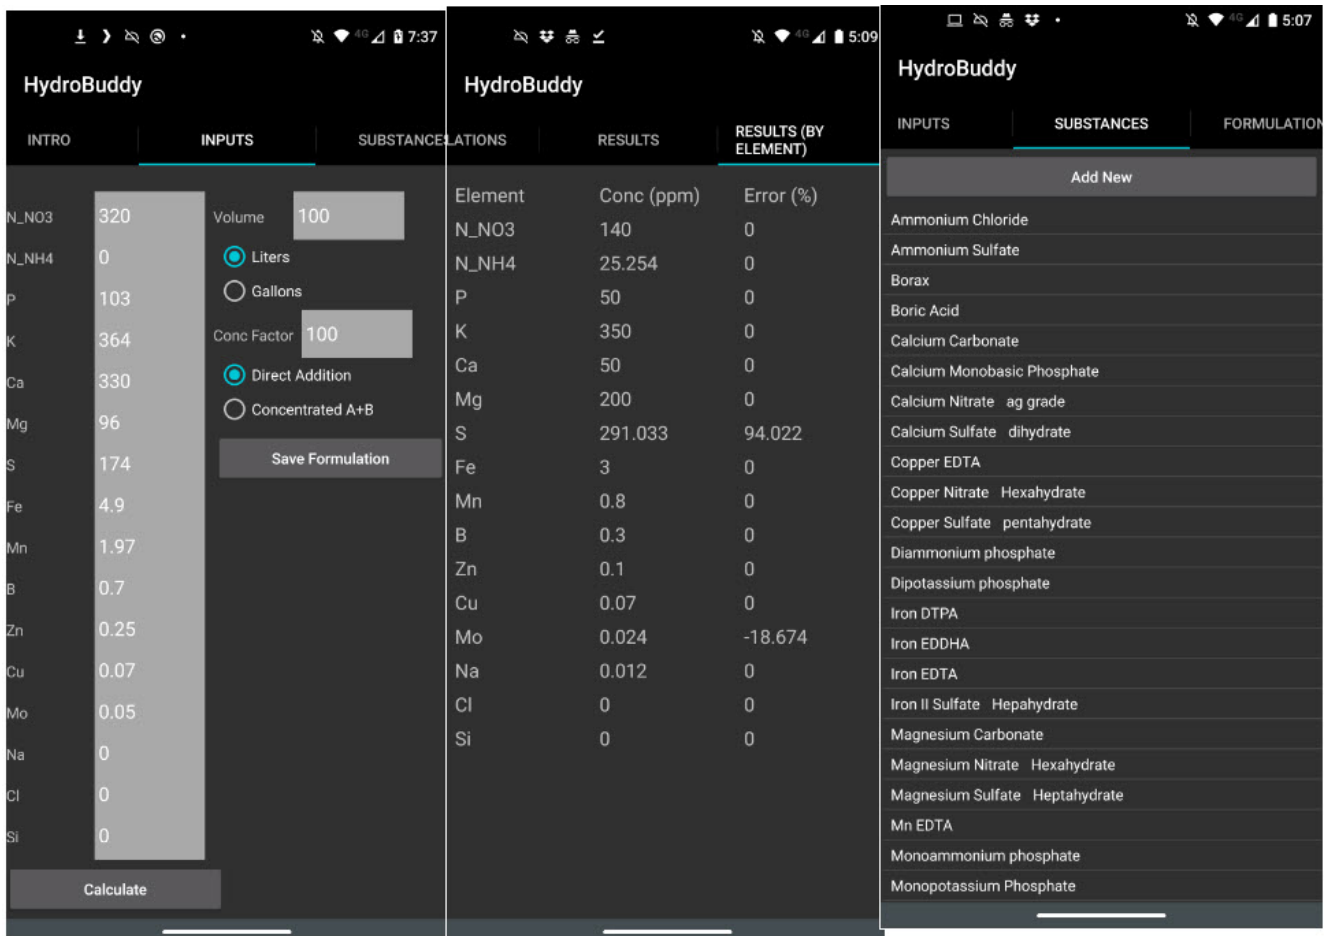
After this, you should be left with an acidic, completely translucent, carnitine and propylene glycol stabilized mono-silicic acid solution that should be around 7-8% w/w of Si as elemental Si. If there's any precipitate in the solution then the stabilization process did not go well and the silicic acid formed polymerized into silica. **This solution should then be used at around 1g/gal, which will provide ~18-20ppm of Si as elemental Si in your hydroponic solution.** When using this solution, the silicon present at the pH used in hydroponics should be much more stable than when derived from direct addition of potassium silicate.

If you go through the above process, leave a comment and let us know how it went.

HydroBuddy coming to Android, free and open source!

The [Hydrobuddy open source hydroponic nutrient calculator](#) – which has been used extensively by both professional and amateur growers for the past 11 years – is finally making the leap to the Android platform. This is a big leap, as many growers – especially in developing countries – lack access to a PC but have easy access to Android phones. Thanks to the effort of the [LAMW development team](#) – who made the development of Android apps using Lazarus possible – I have been able to recode and port HydroBuddy to mobile and will release it for free and with NO ads within the next couple of weeks within the Google Playstore. The [source code](#) is already available on github and will be updated as the application development

process continues.



Some screenshots of the current testing release of the HydroBuddy Android app.

The HydroBuddy Android application will include most of the features present in the desktop application. It will include all the same substances and formulations that are included in the desktop application and will allow you to perform the same types of calculations you do on the regular application. It is however true that some functionality will be missing initially. The “Copy Commercial Nutrient Formulation” section has not been implemented yet, neither have the “Water Quality Parameters”, EC estimations, or calculation of instrumental errors. The “Concentration from weights” and “set weight” calculations are also missing at this point in time.

Right now, the HydroBuddy application is in its internal testing stage, which means that only people who I add to the testing group are able to download and use the application. **If**

you want to participate in this testing phase, you can download HydroBuddy in your Android device through [this link](#). If you participate in this process, please share meaningful feedback about the application with me. *This feedback can be left as a comment in this post as well.*

The port to Android is expected to be a stepping stone towards HydroBuddy v2.0, which will include a complete recoding of several portions of the calculator. This will be done in order to implement a more flexible database structure that allows for more effective saving and sharing of the inputs and outputs of the program. The aim of this is to allow growers of all origins: hobbyists, industry professionals, or researchers, to exchange complete calculation sets and allow for a much more profound use of the calculator as a community-building tool.

Sadly there is no iOS version planned, since I do not own, or plan to own, any Apple devices, and therefore cannot program/compile/test code in this platform. However, the licensing terms of both the desktop and mobile versions of the application do allow anyone to port and publish the application in iOS, provided it is non-commercial and released for free, with no ads and under the same licensing terms (GPLv2).

Calibrating a capacitive moisture/water content sensor for hydroponics

As I've mentioned multiple times in my blog, moisture sensing is one of the most important measurements in a hydroponic crop

that uses a significant amount of media. It allows you to properly time irrigations and avoid over or under watering your plants. Capacitive sensors are the lowest cost initial approach to adequate moisture monitoring of your media. In this post, we are going to learn how to use an Arduino with a gravity shield and a low-cost capacitive moisture sensor in order to accurately monitor the water saturation of our media.



A capacitive sensor plugged into an LCD shield and an Arduino UNO. The measurements are very easy to track with this setup.

An analogue capacitive moisture sensor [like this](#) one does not expose any metallic parts to the media and can be used for the monitoring of moisture content. This sensor is powered by 3.5-5V and gives you a voltage signal that is proportional to the dielectric constant of the media it is in. As the dielectric constant of media changes with the presence of water, so does the signal of the sensor. However, no specific voltage corresponds to a specific water content measurement by

definition, so we need to calibrate the sensor in order to interpret the voltage values we read from it. In order to get readings from this sensor, I use an Arduino UN0, the above-linked capacitive sensor, and an [LCD shield](#) from dfrobot that I can use to easily get the device readings. The arduino code used for this device is also shared below.

```
#include <U8glib.h>

#define XCOL_SET 0
#define XCOL_SET_UNITS 50

U8GLIB_NHD_C12864 u8g(13, 11, 10, 9, 8);
float capacitance;

void draw() {

    u8g.setFont(u8g_font_helvB10);
    u8g.drawStr(0,21,"M_SENSOR:");
    u8g.setPrintPos(XCOL_SET,51);
    u8g.print(capacitance);
    u8g.drawStr( XCOL_SET_UNITS,51,"mV" );
}

void setup() {
    pinMode(LED_BUILTIN, OUTPUT);
    pinMode(4, OUTPUT);
    Serial.begin(115200);
    analogReference(DEFAULT);
    Serial.println("MOISTURE");

    u8g.setContrast(0);
    u8g.setRot180();
}

void loop() {

    draw();
    capacitance = analogRead(1);
    Serial.println(capacitance);
```

```
u8g.firstPage();
  do {
    draw();
  }
  while( u8g.nextPage() );

  delay(1000);
}
```


The calibration of a moisture sensor usually requires the creation of what soil scientists call a “water retention” curve, which is a curve where you plot the sensor’s signal as a function of a fixed volume or weight of water added to the media. However, this approach involves the use of many different containers and the addition of water to the media followed by oven drying steps in order to accurately determine how much moisture was actually in the media for every measurement carried out.

In order to do this procedure in an easier manner, losing as little accuracy as possible, I have created a calibration procedure that makes use of natural drying and only requires one single oven drying step. The procedure is as follows:

1. Heat the media that will be used for an experiment in an oven at 110°C (this drives out all water).
2. Wait for the media to cool to room temperature.
3. Fill the container that will be used for the test (this can plastic but has to have holes at the bottom). Put the sensor in the media, make sure the sensor is driven into the media until the top line is reached.
4. Take a reading, this is the “completely dry” media reading.
5. Disconnect the sensor from the Arduino.
6. Weigh the container+media+sensor. This will be the “dry weight”.
7. Take the sensor out.
8. Add water to the media until there is ample runoff coming out of the bottom.

9. Wait until no more runoff comes off.
10. Put the sensor in the media, making sure you drive it in until the top line is reached. The sensor won't be taken out for the remainder of the experiment.
11. Connect the sensor and take a reading. Take note of this value.
12. Disconnect the sensor.
13. Weigh the container+media+sensor. Take note of this value. The first reading you take is the "max saturation" weight.
14. Repeat steps 11 to 13 every 1-6 hours (time is not very important as long as you gather enough data points) until you reach close to the dry weight. This can take several days depending on the media used. The more points you measure, the more accurate your curve will be.

After carrying out this procedure, you can create a curve like the one shown below. You can use the difference between each weight and the dry weight divided by the difference between the weight at max saturation and the dry weight in order to calculate the water saturation percentage. As you can see, the curves for these sensors are fairly linear in the 40-100% moisture range for the media I tested, while below 40% the regime changes and the measurement increases exponentially until it reached the "dry weight" sensor value. The entire curve can be described with a power-law equation. This behavior will not be the same for all media tested, reason why it's very important for you to calibrate the sensor for the specific media you want to use.





Calibration curve for a capacitive moisture sensor. In this case, the media was a mixture of 50% river sand and 50% rice husks.

Once you have your sensor calibrated you can know what a measurement in mV implies in terms of water saturation for a given media type. This allows you to time your irrigations at a given water saturation % effectively. Which water saturation percentage might be better, depends on the properties of your media and how the water potential changes as a function of the water saturation. However, this allows you to experiment with irrigations at different water saturation % values and figure out exactly where you need to water so that your plants are not under or overwatered.

It is also worth noting that the above sensor is probably not rugged enough for use in a hydroponic setup without a bit more hardening. In order to use these sensors in practical applications, you should apply conformal coating to the electronics at the top of the sensor and then use shrink tubing in order to fully protect these electronics from the elements.

Properly positioning temperature and humidity sensors in a hydroponic growing environment


Temperature and humidity are two key variables you need to measure. They are important because they determine how your plants will transport water, and transpiration controls a lot of processes, with nutrient transport being one of the most important ones. However, the value of these variables in a growing environment – being that a greenhouse or a grow room – can change substantially depending on where they are measured. It is therefore critical to know where to place sensors and how to interpret their readings based on their location. In this post, we are going to discuss where it is best to measure these variables and what consequences it could have if these values are not measured properly.



Sensor placement relative to a control source (AC, humidifier/dehumidifier) for one or two sensors. Note that this setup assumes good circulation throughout the room, including middle of canopy.

Let's start with the worst possible case, you only have one set of sensors and you need to control your environment with it. *In this case, place your temperature and humidity sensors at canopy height, as far away as possible from both the AC ducts and the humidifier/dehumidifier, make sure the sensors are hanging in the air and not stuck to a wall or tubing.* Then, make sure you use a hot wire anemometer to verify that your air movement speed is at least 0.3m/s across the entire room. This setup ensures that the worst controlled part of the room is at the correct value and it also attempts to minimize the gradient created from the control sources to the sensor using a good amount of air circulation. It is not perfect though and significantly different "climate zones" will be created close to and far away from the climate control devices.

The above setup can be used effectively in small growing environments, but can be problematic as both the number of plants and the size of the growing environment increases. At this point, using a single set of sensors is not an option if adequate climate control is desired. In these cases, multiple sensors need to be placed to ensure that climate control is being done properly. When using multiple sensors, place the second sensor at the place with the lowest air circulation inside the room, at canopy height, which is usually in the middle of the room, then place subsequent sensors as far away from either this or the first sensor in sequence. When doing climate control, the system needs to ensure all of the sensors remain within a "safety band", making sure no sensor becomes too cold/hot, humid/dry, during control cycles.



Sensor placement for multiple control sources. Sensors are placed in order trying to always be as far away from sensors as possible but within the plant canopy.

When you implement a sensor system like this, you will realize pretty quickly that climate becomes very hard to control in a larger room when there is only one source of control (one AC, one humidifier, and one dehumidifier) because gradients become too big for effective control, so it takes too long for the AC

to be able to properly control the room while ensuring all sensors remain within proper boundaries. In this case, it becomes necessary to add multiple sources of control, so that the extent of gradients within the room can be minimized. This means adding multiple ducts for the output of an AC, multiple humidifier/dehumidifiers, etc.


In these cases, sources of control are placed outside of the plant canopy to avoid plants being exposed to the flows from these control sources (which would expose them to very cold/hot/dry/humid air). Then the sensors need to be placed within the plant canopy, starting from as far away as possible from all sources of control – usually the middle of the canopy – and then to the corners of the growing environment.

Note that the control algorithm needs to ensure all of these sensors are within the proper control band and not attempt to control the average reading of these sensors. If you try to use the average of sensors to control a room, you might be left with a room where two extremes are present and the control system believes everything is ok while these extremes are maintained. The median is a better way to control a room, but it only becomes useful when 5 or more sensors are used. If only 2 or 3 sensors are used, ensuring all of these sensors are within adequate bands is fundamental to ensure that the room will have a lower chance of having humidity/temperature microclimates that will be detrimental to plant growth.

Making the most out of your hydroponic setup's logged

sensor and control data


If you have a hydroponic crop with a data logging and automated control solution, you probably have a lot of sensor and control data recorded that could be useful to take your crop's results to the next level. In this post, I am going to talk about some things that you could be doing with these stored data. You will see how the usage of this data opens up many possibilities and that even implementing the most basic of these suggestions could lead to important improvements in your understanding of your crop and its results.



Use of different moving average to smooth out sensor readings. The lowest hanging fruit to take advantage of logged data is to be able to download the data and put it into a database structure that you can properly query and search. Most data logging solutions record the data in either very simple structures, like csv files, or non-relational databases – like MongoDB – which are rather limited and do not allow for the degree of versatility that a true relational database engine offers. Having the data in a properly built database will allow you to start using it in a creative way. For example,

with the data in a proper database, it becomes possible to create a custom data visualization that can help you understand what's going on inside your growing environment.

The images in this post show you some examples of this. The first one shows a simple example where a rather noisy humidity sensor is smoothed out using different moving averages, these averages can then be used to implement more effective control algorithms. The second image shows a detailed map of the temperature and humidity values experienced in a room, colored by the hour of the day. We can use this plot to easily locate where problematic times and VPD conditions might be, just by looking at when extreme readings happen. This behavior would be harder to observe and diagnose on a regular VPD Vs Time plot. Regular data logging web interfaces and platforms will not allow you to create plots of this sort, which is why putting your data into a proper DB and manipulating it to create custom visualizations can be very powerful.



Relative Humidity Vs Temperature map colored as a function of the hour of the day for a growing room being constantly monitored

The most powerful uses of the data come into play when you actually piece together your control and sensor data. Say you have an AC system coupled with a temperature sensor but you have a lot of other temperature and humidity readings and you also know the age of your plants at each point in time. Using this, you can create an advanced control algorithm where a system will use all of this additional information to know when to trigger AC systems and dehumidifiers to control the environment. Having a lot of logged data from a set point control system is a great starting point to train a reinforcement learning algorithm for climate control, since we know which control actions were taken at each point in time and we know the effect these had. Implementing such control mechanisms can lead to control systems that avoid spikes in

humidity and temperature across light on/off cycles, greatly smoothing out the environmental transitions for your crops.

Finally, there is also the potential to improve yields by gathering detailed mappings of yield data in a room and relating these yields with environmental sensors. If you have several different sensors in a room and you know the yield that you obtained on a per-plant basis, then you can create a map of all the yields in a room in order to see if there are important disparities in your yields because of differences in local humidity, temperature, light or air circulation levels. This can lead to important insights that can help better adjust climate conditions for the entire grow room. If multiple rooms are available, the information about environmental sensor data can be related to yields in order to stir all rooms towards more favorable conditions.

For example, after analyzing yield and temperature data from multiple growing cycles of one of my customers, we realized that the greenhouse with the lowest temperature standard deviations between sensors was giving the best yields, we then implemented better control algorithms on the other greenhouses to prevent this from happening, obtaining significantly better results across the board after that.

Data is a treasure. If you have been recording judicious sensor, climate control, and yield data through time, you're probably sitting on a gold mine that you haven't exploited yet. If you're interested in using my help to do so, please consider [booking an hour of consultation](#) time with me so that we can discuss your needs and how we could leverage your data to improve your growing results.

Commercial sensor and data logging solutions for hydroponics

On a previous post, I discussed a very interesting open-source sensor/data logging alternative for Hydroponics called [MyCodo](#), which offers a lot of features and flexibility for those growers with the time and skills necessary to implement their own sensor and data logging setup. However, many growers don't have the time to do this on their own – or the time and willingness to hire someone to do it for them – and all they want is a solution that “just works” out of the box and that fits most of their data logging needs. In this post I am going to talk about three commercial solutions – in no particular order – that I've had experience with along with some of the advantages and disadvantages that each one offers you. **Note that this post has not been sponsored by any of these brands.** *The statements below represent my opinions on the matter and the facts, to the best of my knowledge. I recommend you contact each company to ask specific questions pertinent to your needs.*



[Growtronix](#). This company offers a complete solution for monitoring and automation of hydroponic crops. Their sensors are hooked through cabled connections and they support a wide array of analogue sensors, both sold by them and by third parties. As long as a sensor can work on a 3.5-5V input and give an analogue reading, it can be installed in a growtronix setup. Their web interface is user-friendly, it allows you to

view sensor readings and create control schemes using simple if logic statements. They have also shared the source code of their web interface with some of my customers in the past, so if you would like to customize things beyond their base web application, I'm sure you could figure it out if you have the time and programming skills. Growtronix support – per the experience of the customers I have you have used it – has been stellar.

There are however some downsides to using growtronix. Since everything is cabled you will need to lay cables across your rooms if you want to hook up multiple sensors within them. The system lacks support for third party i2c sensors, meaning that you can only connect analogue sensors and will miss on some interesting third-party sensor offerings. The data is also stored in a non-relational mongoDB implementation, which means that querying data and doing complicated data analysis will not be easy with them. Their control algorithm technology is also rather simple, to the best of my knowledge they do not offer more advanced control mechanisms beyond the if logic statements they allow the users to program.



[Agrowtek](#). Similar to Growtronix, they also offer a complete monitoring and automation solution for hydroponic crops. However, they offer their own touchscreen computers to connect to their sensors, dosing pumps, and relay modules, so they do not have a dedicated web interface for their sensors that is hosted on any computer but you must purchase their own. Their “GrowControl” panels will hook with normal ethernet cables to any of the sensors they offer and you will be able to program all the behavior of the sensors and the relays from these stations. Their main advantage is easy setup, everything easily hooks up and you can then program things within the

GrowControl panels to fit whatever simple control needs you might have. You can probably setup 200 sensors/relays in a day to control an entire facility using this setup. Their custom computer also gives you more stability, meaning crashes of the system are rare (according to the customers I have who have used them). From the three companies discussed in this post, this is also the only one to offer nutrient injection systems in their offering.

However, one big limitation of this company is how closed the ecosystem is. You have absolutely no ability to hook up third-party sensors and sadly their offering lacks some important and basic sensors for a medium to large scale hydroponic setup, specifically water content and water potential sensors. You are also becoming reliant on the availability of support from them and – if the company went under – it would be very hard for you to be able to fix or find replacements for their sensors or their control panels. Their control algorithms are also fairly simple and are limited to basic if-logic, similar to the Growtronix system. Data is also not logged into any database but as basic csv files, which means substantial effort will be needed to perform advanced data analysis tasks.



[SmartBeeControllers](#). This company also offers a complete automation and monitoring solution for your hydroponic crop. Their main differentiating factor relative to the last two is that sensor stations connect wirelessly to your computer, allowing you to place sensors throughout your facility without having to set up cables through the entire place. Their sensor stations can hook up to a large number of sensors so, for example, you can use a water content station to hook up six of their capacitive water content sensors. They also require a computer server with the web software to communicate with – alike Growtronix – and their software has a focus on simplicity. In this case, control options are even more limited than in other cases, with basically only simple set-

point logic available to control relays (to the best of my knowledge).

The SmartBee ecosystem is also quite limited and offers no pH/EC/ORP sensors or water potential sensors (tensiometers). You have no ability to hook up third-party sensors as well, meaning you're stuck with this offering if you use them. Because of the wireless nature of communications, sensor readings and their stability can also be compromised due to excessive electromagnetic noise, which can be particularly problematic in a short room that has a lot of HPS ballasts. It is also true that in the past (2-3 years ago) their support seemed to have problems, with several complaints about their response time online. I do not know if their technical support has improved so I would advise you to seek recent opinions about it on social media if you're considering them for purchase. The people I know who used them didn't need to contact support, so I cannot comment on this aspect from my customers' experience.

The above are three commercially available data logging systems for hydroponics. All of them should be easy to hook up and should provide you with basic data logging and control capabilities for your grow. In my opinion, the most complete one is Growtronix, given the ability to add third-party sensors – even if only analogue ones – and the quality of their sensors and web application software. However, if controlling the nutrient injection process electronically is important for your situation, then Agrowtek might be a better solution. None of them however provide advanced control mechanisms – like reinforcement learning-based climate control – and none of them provide access to all sensors that would be desirable, so a custom DIY setup might be best if these features are very important to you.

Standard hydroponic formulations from the scientific literature

When researchers started looking into growing plants without soil, they started to look for mixtures of nutrients that could grow plants successfully so that these formulations could be used to study other aspects of plant physiology. If you have a mixture of nutrients that you know grows a plant without major issues, then you can use that as a base to study other things, for example how plants react to some exogenous agent or how changes to temperature or humidity affect the uptake of certain nutrients (see this paper for a view into the history of hydroponics and standard solutions). The establishment of these standard solutions was one of the great achievements of botanists during the twentieth century, which allowed thousands of detailed studies on plants to be carried out. In this post, we're going to be talking about these standard solutions and why they are a great place to start for anybody seeking to formulate their own nutrients.


ppm (mg/L)	1	2	3	4	5	6	7	8	9	10	11	12
K	132.93	187.28	241.24	312.79	236.15	237.33	89.54	157.57	261.57	302.23	430.08	312.79
Ca	136.27	36.07	149.09	163.52	200.39	160.31	161.11	120.23	184.76	172.34	220.43	160.31
Mg	19.69	18.71	37.19	49.34	48.61	24.31	55.90	48.61	49.10	50.55	36.46	34.03
N as NH ₄ ⁺	0.00	4.90	2.10	18.91	0.00	28.01	19.61	0.00	0.00	0.03	0.01	17.51
Na	0.00	0.23	1.15	0.46	0.00	0.46	0.00	2.07	0.46	0.69	8.74	0.69
Fe	36.86	2.79	4.02	0.00	1.44	1.12	1.12	5.03	1.34	1.90	7.10	0.84
Mn	0.00	0.62	1.23	0.00	0.50	0.11	0.14	0.40	0.62	1.98	2.40	0.55
Cu	0.00	0.06	0.01	0.00	0.02	0.03	0.00	0.02	0.01	0.10	0.04	0.04
Zn	0.00	0.01	0.01	0.00	0.05	0.13	0.13	0.05	0.11	0.10	0.12	0.03

N as NO3	123.82	77.46	161.50	226.63	210.10	196.09	112.75	112.05	167.80	201.28	241.62	224.11
P	103.45	42.74	64.74	40.89	30.97	61.95	71.24	61.95	30.66	59.78	69.69	38.72
S	25.97	27.90	54.51	65.09	64.13	32.07	96.84	64.13	111.59	67.98	87.22	44.89
Cl	0.00	0.00	0.00	0.00	0.64	1.77	0.00	0.53	0.00	0.00	13.47	0.00
B	0.00	0.28	1.19	0.00	0.46	0.27	0.10	0.40	0.43	0.30	0.34	0.27
Mo	0.00	0.41	0.00	0.00	0.01	0.05	0.00	0.03	0.05	0.19	0.06	0.34

Summary of standard nutrient formulations found in [this article](#) with the concentrations translated to ppm. The numbers in the list correspond to the following: 1. Knop, 2. Pennings-feld North Africa, 3. Pennings-Feld Carnations, 4. Gravel Culture Japan, 5. Arnon and Hoagland 1940, 6. Dennisch R. Hoagland USA, 7 Shive and Robbins 1942, 8. Hacsalyo 1961, 9. Steiner 1961, 10. Cooper 1979, 11 Research Centre Soil-less culture, 12. Naaldwijk cucumber.

One of the best places to find a comparison between these standard solutions is [this paper](#). In it, the authors explore the relationships between the different solutions and how they are similar or diverge. In the table above, you can see a summary of the elemental nutrient concentrations found in this paper for the 12 standard solutions they compare (the paper states them in mmol/L but I have changed them to ppm as these are more commonly used units in the field nowadays). As you can see, some of the older solutions miss some elements or contain much smaller amounts of them – as they were likely present in the media or other salts as impurities – while more recent standard solutions do contain all the elements we now understand are necessary for plant life.

NH₄⁺ proportion:
 2: 0.04
 3: 0.01
 4: 0.06
 6: 0.11
 7: 0.08
 12: 0.57




Standard solutions:

1. Knop
2. Penningsfeld N. Africa
3. Penningsfeld Carnations
4. Gravel culture, Japan
5. Arnon and Hoagland
6. Dennish R. Hoagland
7. Shive and Robbins
8. Hacskalyo
9. Steiner
10. Cooper
11. Res. Centre Soil. Cult.
12. Naaldwijk, cucumber

FIGURE 1. Cation composition of the standard solutions.

Figure showing the Ca/Mg/K ratio represented in a three axis plot. Taken from the paper mentioned above.



Standard solutions:

1. Knop
2. Penningsfeld N. Africa
3. Penningsfeld Carnations
4. Gravel culture, Japan
5. Arnon and Hoagland
6. Dennish R. Hoagland
7. Shive and Robbins
8. Hacskalyo
9. Steiner
10. Cooper
11. Res. Centre Soil. Cult.
12. Naaldwijk, cucumber

FIGURE 2. Anion composition of the standard solutions.

Figure showing the N/S/P ratio represented in a three axis plot. Taken from the paper mentioned above.

It is interesting to note that all of these solutions have been successfully used to grow plants, so their convergent aspects might show us some of the basic things that plants require for growth. As they highlight on the paper, the

K/Mg/Ca ratio for most of these solutions is rather similar, as well as the N/S/P ratios. This means that most of these authors figured out that plants needed pretty specific ratios of these nutrients and these ratios are sustained with minor variations through the 12 solutions, developed across a span of more than 100 years. All the solutions developed from the 1940s have similar final concentrations and their starting pH is almost always in the 4-5 range, due to the presence of acid phosphate salts like monopotassium phosphate.

Nonetheless, there are several things that improved in the solutions as a function of time. The first is the inclusion of higher concentrations of all micronutrients with time, as macronutrient salt quality increased, the media sources became more inert and the need to add them to avoid deficiencies became apparent. The need to chelate micronutrients also became clear with time, as solutions starting with Hoagland's solution in the 1940s started using EDTA to chelate iron, to alleviate the problem of iron phosphate precipitation in hydroponic solutions. This is clearly shown in the table below, where the authors show how the first three solutions had almost or all of their Fe precipitate out, while the newest solutions, like Cooper's developed in 1979, had less than 5.5% of its Fe precipitated.

Standard solution	% Fe precipitated as Fe ₂ (PO ₄) ₃	% Cu complexed with chelate	% Zn complexed with chelate	% Mn complexed with chelate
1. Knop 1865	100	-	-	-
2. Penningsfeld North Africa	99.9	-	-	-
3. Penningsfeld Carnations	99.9	-	-	-
4. Gravel culture Japan	-	-	-	-
5. Arnon and Hoagland 1940	87.8	40.3	6.4	0.3
6. Dennisch R. Hoagland	4.0	97.5	97.5	0.1
7. Shive and Robbins 1942	99.9	-	-	-
8. Hacsalyo 1961	4.0	99.3	42.4	0.2
9. Steiner 1961	4.9	99.5	48.8	0.2
10. Cooper 1979	5.5	98.3	22.4	0.1
11. Res. Centre Soil. Cultures	6.9	100	99.2	7.7
12. Naaldwijk cucumber	4.5	96.5	7.8	0

This table shows the precipitated Fe and chelated portions of the micro nutrients in all the standard solutions.

The natural question when reading about standard solutions is: which one is the best one to use? Sadly, I don't think there's a simple answer. There have been multiple studies comparing standard solutions (see [this one](#) for an example). What ends up happening most of the time is that, while most of the solutions manage to grow healthy crops, one of the solutions happens to be more fit to the idiosyncrasies of the study because its conditions are better aligned with those that the authors developed the solutions under. A study revealing a solution to be better than another to grow plants under a given set of conditions does not imply that this solution will be the best one for all plants under all conditions. For this reason, the optimization of nutrient solutions to particular conditions using tissue analysis is still pursued in order to maximize yields.

My advice would be to view the above solutions as well researched starting points for your hydroponic crops. These solutions, especially the ones developed after 1940, will do a good basic job growing your plants. If you're interested in

making your own solutions, starting with a solution like the Hoagland, Steiner, or Cooper solutions is a great way to begin making your own nutrients. Once you have a basic standard solution working for you, you can then tweak it to maximize your yield and improve your crop's quality.

The stability of metal chelates

When you get introduced to hydroponics and nutrient solution chemistry, one of the first concepts that you learn is chelation. A chelate is a molecule formed by a metallic ion and a chelating agent – which is also referred to as a ligand – where the metal ion is wrapped around very tightly by this ligand. The job of the chelating agent is to keep the heavy metal ion shielded from the environment, allowing it to exist in solution without forming potentially insoluble compounds that will take it out of the nutrient solution. However, these chelates can be unstable or too stable, both of which can hinder the availability of the nutrient to plants. In this post, we're going to talk about what determines the stability of a metal chelate and how you can know if a given chelate will be able to fulfill its job in a hydroponic environment.



$$K_b = \frac{|ML|}{(|M| \times |L|)}$$

A simplified view of the chemical equilibrium formed $|M|$ refers to the concentration of the free metallic ion, $|L|$ the ligand concentration and $|ML|$ the chelate concentration. Charges are omitted for simplicity.

Since chelates are formed by the reaction of a metallic ion – most commonly a cation – which a ligand, a chemical equilibrium is established between the free metallic ion, the ligand, and the chelate. Every second, there are lots of chelate molecules being formed from reactions between metallic ions and ligands, and free metallic ions and ligands are being formed from the disassembly of the chelate. The process is in equilibrium when the rates of assembly and disassembly are the same. The equilibrium constant – also known as the stability constant or K_b – tells us how displaced this equilibrium is towards the product (in this case the chelate). When the K_b value is large, the concentration of the chelate at equilibrium is very large, while when K_b is small, the opposite is true. Since these numbers are usually very large for chelates, we express them as pK_b which is $-\text{Log}(K_b)$. These constants depend on temperature, but their values are independent of other chemical reactions. However, things like pH can affect the concentration of ligand or metal cation, which can affect the concentration of chelate, since the equilibrium constant's value remains the same.

	Al(III)	Ba	Ca	Co(II)	Cu	Fe(II)	Fe(III)	Hg	Mg	Mn	Ni	Sr	Zn
Acetic acid		0.39	0.53	2.24				3.7d	0.51		0.74	0.43	1.03
Adenine													
Adipic acid		1.92	2.19		3.35								
ADP		2.36	2.82	3.68	5.9				3.11	3.54	4.5	2.5	4.28
Alanine		0.8	1.24	4.82	8.18					3.24	5.96	0.73	5.16
b-Alanine					7.13						4.63		4
Albumin			2.2										
Arginine						3.2				2			
Ascorbic acid			0.19									0.35	
Asparagine			0									0.43	
Aspartic acid		1.14	1.16	5.9	8.57				2.43	3.74	7.12	1.48	2.9
ATP		3.29	3.6	4.62	6.13				4	3.98	5.02	3.03	4.25
Benzoic acid					1.6						0.9		0.9
n-Butyric acid		0.31	0.51		2.14				0.53			0.36	1

Polyphosphate			3		3.5	3			3.2	5.5	3		2.5
Proline						4.07	10			3.34			
Propionic acid		0.34	0.5		2.2		3.45		0.54			0.43	1.01
Purine					6.9						4.88		
Pyrophosphate			5		6.7		22.2		5.7		5.8		8.7
Pyruvic acid			0.8		2.2								
Riboflavin				3.9	<6					3.4	4.1		<4
Salicylaldehyde				4.67	7.4	4.22	8.7		3.69	3.73	5.22		4.5
Salicylic acid	14.11			6.72	10.6	6.55	16.35		4.7	2.7	6.95		6.85
Sarcosine				4.34	7.83	3.52	9.7				5.41		
Serine			1.43			3.43	9.2				5.44		
Succinic acid		1.57	1.2	2.08	3.3		7.49		1.2	2.11	2.36	0.9	1.78
(+)-Tartaric acid		1.95	1.8		3.2		7.49		1.36		3.78	1.94	2.68
Tetrametaphosphate		4.9	5.2		3.18				5.17		4.95	2.8	
Threonine						3.3	8.6						
Trimetaphosphate			2.5		1.55				1.11	3.57	3.22	1.95	
Triphosphate		6.3	6.5		9.8				5.8			3.8	9.7
Tryptophan							9						
Uridine diphosphate									3.17				
Uridine triphosphate			3.71	4.55					4.02	4.78			
n-Valeric acid		0.2	0.3		2.12								
Valine					7.92	3.39	9.6			2.84	5.37		5
Xanthosine				2.8	3.4	<2					3		2.4

This table was originally present in a website that no longer exists. The data is taken from the [NIST reference of heavy metal complexes](#).

The table above shows you the pKb values for different metal ions and different ligands or chelating agents. Since the pKb scale is logarithmic, a difference of 1 indicates an order of magnitude higher stability. You can also find additional references to other stability constants in this link. These constants allow us to predict which chelates will be formed if different metallic cations and ligands are present. Let's say we have a solution that contains Ca²⁺ and Fe³⁺ and we add a small amount of sodium citrate, what will happen? Since the constant for Ca²⁺ is 3.5 but that of Fe³⁺ is 11.85, citrate will chelate around 1 billion Fe³⁺ ions for every Ca²⁺ ion it chelates. In practice, this means that all the Fe³⁺ that can be chelated will be, while Ca²⁺ will remain as a free metallic ion. However, if we have Fe²⁺ instead of Fe³⁺ then Fe²⁺ has a constant of only 3.2, which means that one molecule of Fe²⁺

will be chelated for every 3 of Ca^{2+} , meaning we will have around 25% of all the chelate formed as a chelate formed by Fe^{2+} and 75% as a chelate formed by Ca^{2+} .

We can see in this manner how chelating only one heavy metal can lead to problems. Imagine that you purchase Iron EDTA and add it to your nutrient solution, but you have added Manganese from Manganese sulfate. Upon addition, the FeEDTA chelate will disassemble to generate as much Fe^{2+} and free EDTA as dictated by the equilibrium constant and the free EDTA will then get into equilibria with all the other heavy metals, since the constant with Mn is 13.6 and that of Fe is 14.3 the ligand will redistribute itself so that it complies with all the chemical equilibria present. This means that for every 7 Fe^{2+} cations that are chelated we will have around 1 Mn^{2+} containing chelate, so you will lose around 14% of the chelated Fe in order to chelate free Manganese. That free Fe^{2+} will be unstable and precipitate out, which will shift the equilibrium and cause us to lose more of the Fe chelate. This is how competing equilibria can lead to the slow but sure depletion of available cations in solution.

With the above references and charts, you should now be able to look into any chelating agent you want to use and determine how good of a choice it is for your solution and what is likely to happen once you put that chelate in. The ligand will chelate different metals in order to comply with all the equilibrium constants, so it is up to you to add enough so that all heavy metals are satisfied or add ligands whose affinity for a given ion is so high that the others are just unable to compete for it, almost regardless of their concentration.