

Operation manual, greenness loggers.

Keep them in:

- **OUTDOOR SUNLIGHT** (so they stay charged by their solar cells) or
- **TOTAL DARKNESS** (so they power-down into sleep mode).

Avoid keeping them in indirect light.

In low light they log useless data, and slowly waste their batteries.

Contact information:

Developer: Rick Shory

Company: rickshory.com

Address: 8828 NE Schuyler St, Portland OR 97220

Email: rickshory@gmail.com


Phone: 503-253-8100

Document version 4.1 **DRAFT**

22 June 2017

Overview:

The Greenlogger is a fist-sized instrument that logs infrared and visible reflectance, and so records ground-based Normalized Difference Vegetation Index (NDVI) trends. NDVI is a "greenness index" that can be a surrogate for plant phenology. For more details, see Appendix A.



Phenology data can be very expensive to collect in remote sites because it requires re-visits as frequent as twice-weekly, for human observers to accurately catch the desired plant growth stages.

The Greenlogger runs unattended, and automatically records spectral reflectance data from the vegetation it is set over.

The Greenlogger is not a phenology camera, which looks horizontally at a landscape. Such cameras have their use, but this system looks down vertically. Down-looking sensors are more comparable with satellite imagery, and are especially important in arid lands with significant bare soil between plants.




In a typical installation, a Greenlogger has a roughly circular field of view about one meter across. At this scale, a season-long recording can capture the greenness trend of an individual shrub or patch of grass. Comparing trends between instruments can show the effects of treatments such as supplemental watering or early removal of snow, or variations due to site parameters like aspect and slope.


Field of View, details

The Greenlogger has a cone-shaped field of view. The cone angle is about 45 degrees out in every direction from normal. So the total cone is about 90 degrees across.

The diameter of the circle is about double the height above ground, so 50 cm above ground would give a circle roughly 1 m across. Mounted higher, above taller vegetation, would give a correspondingly larger view circle.



You can conceive of what the sensor "sees" by looking at the collimator, the rhombic black plastic component that shows through the bottom of the instrument case.



The irradiance sensing chip is down in the hole of the collimator. By tilting the Greenlogger back and forth, you can see where the view cuts off. From angles where you can see the chip, light from that angle would be detected.


The cutoff angle is not exact because the chip has a nonzero horizontal extent. Although it is small, it is not a geometric point. The chip contains both the infrared and the broadband photodiodes, so the view cone is much the same for both bands.

Modes of operation:

A Greenlogger records data all the time. You do not need to turn it on or off. In between readings, it goes into sleep mode to save power.

Normally, in Day mode, it wakes up every ten seconds and takes readings. This corresponds to the blink of the yellow pilot light you see when a logger is sitting out in the sun (there may be one blink or two). After taking readings, the logger goes back to sleep for ten seconds.

Every two minutes, on the even numbered minute, the logger writes data to the SD card. This is not data averaged over the past two minutes, but the instantaneous readings at the time point. This is more useful for NDVI.



When ambient light levels drop below a certain threshold, which is currently the equivalent of late dusk, the logger goes into Night mode. (In future versions, this threshold will be user adjustable.) In Night mode, the logger only wakes up and takes readings once per hour, on the hour. In the morning, when one of these hourly wakeups finds light levels above the threshold, the logger goes back into Day mode.

You can use Night mode to keep the loggers on low-power, for example for shipping. Just put them inside a lightproof covering. Once they are taken out into the light, on the next hour turnover they will resume Day mode.

On their hourly Night mode wakeup, loggers always record their data (as long as there is sufficient battery power). This provides a once-per-hour baseline of conditions at night or in storage.

Rousing a logger to communicate

You control the loggers and get their data out by a Bluetooth radio link. You do not need to open the case. All data are backed up on the internal SD card, for emergency recovery in case of accident; but in routine operation you do not need to open anything.

A logger normally keeps its Bluetooth module turned off, to save power. To "rouse" a logger to communicate, you give a sharp tap on the front of the case, or the front-bottom. You have to hit quite hard. This is intentional so it will not happen by accident. You may barely be able to tap hard enough with your fingernail, more likely with a small tool like the handle of a screwdriver. It's not the force of the blow, but the sharpness; technically the "acceleration".

You actually need to give two taps, with a short interval between. The first tap brings the logger to half-roused mode, where the yellow pilot light blinks continually (once per second). The next tap, after the logger is half-roused, rouses it fully. The red pilot light starts blinking, which means the logger has activated its Bluetooth module and is waiting for a connection. If you leave the logger alone, in two minutes it will go back out of Rouse mode. The red light will go off, and the yellow light will resume blinking only every ten seconds.

You can also use Rouse mode to bring the logger out of Night mode. Rouse mode automatically puts the logger in Day mode. If the logger leaves Rouse mode and there is enough ambient light, the logger stays in Day mode. This way, you can wake up your loggers when you take them out of storage, without waiting for the normal hour turnover.

Bluetooth

Researchers often deploy a lot of recording instruments, twenty or more. With cables, you would go around to each one, open the case, plug in, and download your data. This risks letting moisture into the case, and you have to remember to bring the right cables and/or tools to unscrew things. But at least you would know exactly which instrument you are talking to.

Now, with wireless Bluetooth, you will create a "virtual cable" to each logger. There will be no physical cables to drag around. Instead, they will exist electronically on your smartphone or laptop. However, there is a certain amount of procedure to keep track of which logger you are talking to.



We first describe setting up connection to an Android device such as smartphone, then an alternative using a computer such as a laptop. The Android connection is much simpler.




Android device: Connecting to loggers by Bluetooth

Begin by installing two apps on your smartphone, BlueTerm and AirDroid. Both of these are free.

BlueTerm is one of many Bluetooth terminal apps. It is the one we have found most satisfactory.

AirDroid makes file transfer to your computer easy, after capturing logged data to your smartphone.

You will usually be using your smartphone to collect data from loggers, using the "D" (dump) command, so these instructions concentrate on that, but you can use any of the commands (covered later).



Before connecting, familiarize yourself with the Bluetooth module in the logger. This will enable you to connect to the correct logger if you have several.

The physical Bluetooth module shows through the transparent bottom of the logger case.

The last four characters of the serial number on the Bluetooth module will be the "device" your Android will see.

In this example, the characters are "B364". On some BT modules, the characters are right before "FLY".




On newer devices, the characters are at the end of the MAC address.



You may wish to write this code on the logger mount as a reminder. Grease pencil, also called "china marker", is the most sunlight resistant.

If you are not sure of the code, the following steps will make it clear which one is correct.



Rouse the logger by tapping it. Tap twice so the red LED starts blinking.

On your Android, start BlueTerm.

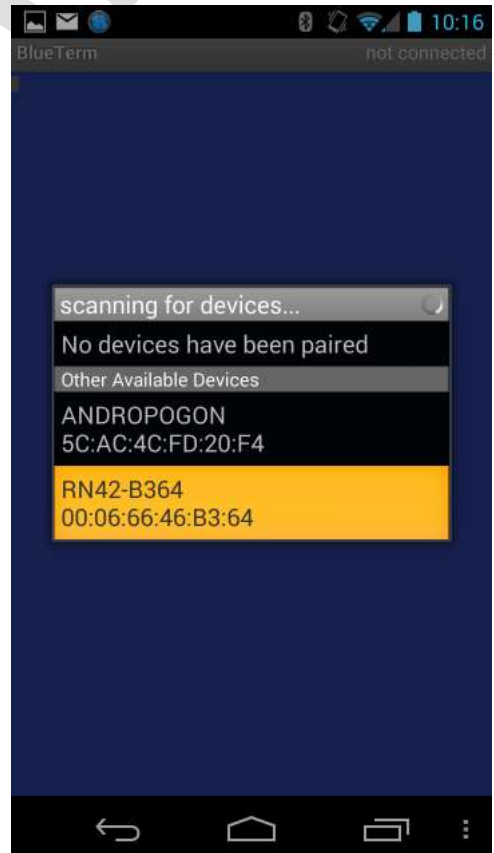
Tap the three-dots menu area to bring up the options, then tap "Connect device". Occasionally, you may have to try connections two or three times.



If you have not connected to any loggers yet, only the "Scan for devices" option will be available.

When you tap this, your Android will find the logger and show it as a device starting with "RN42-". The last four characters will be from the serial number.

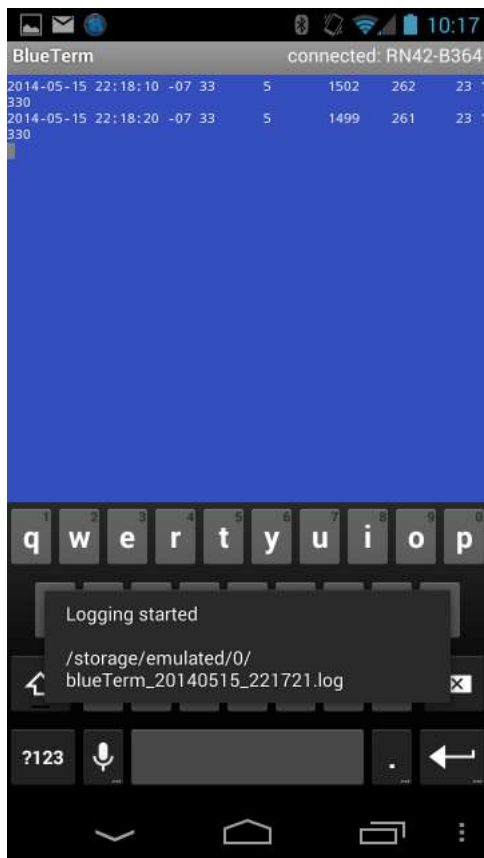
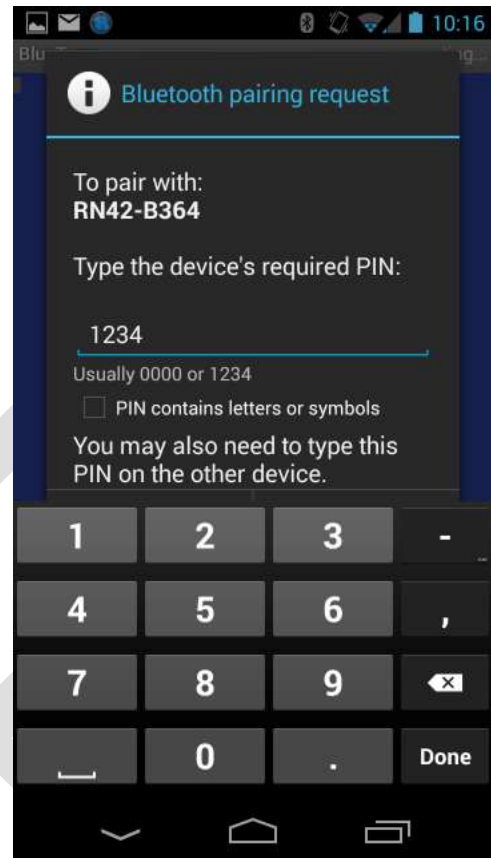
Next time, your Android will show it in a list as soon as you tap "Connect device".



The first time, complete the pairing by entering the PIN. For the RN42, this is always "1234".

Some newer modules will not ask you to input a PIN. Instead, the pairing will just ask you to verify the device is displaying a certain code. Since the module has no display, just verify it. If the logger is the only active, unpaired Bluetooth device nearby, this has to be the correct device.

If you have difficulty pairing with the logger through BlueTerm, it may be easier, the first time, by going directly through your phone's Bluetooth settings. Then return to BlueTerm, do "Connect Device", and the module will show up in the list.



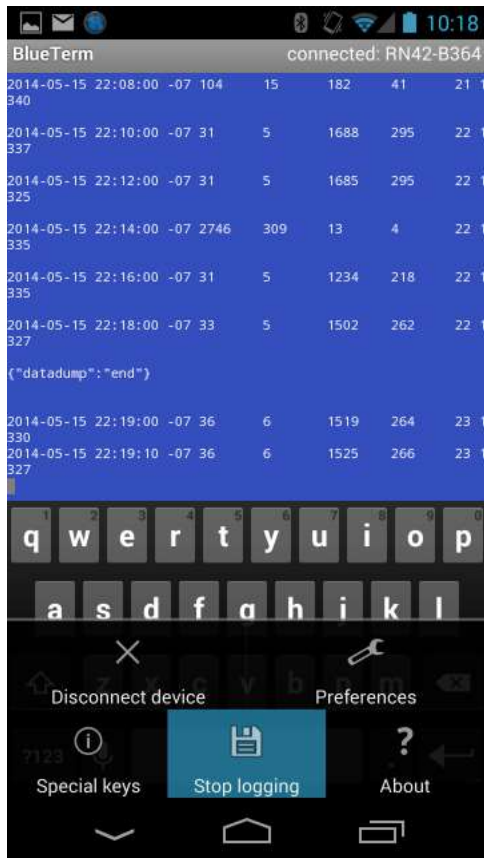
When the connection completes, you will see the logger's data stream. On the logger, the red LED goes out and the green one comes on.

Sometimes, you may have to try the connection two or three times.

To begin logging, tap the menu area again, and this time choose "Start logging"

A message briefly appears saying where data will be logged. The location is standardized, and the file name is based on the date and timestamp.

Each logger tags its data with identifying metadata. This will go into the log file, so the name of the log file is not important.



Logger commands work the same on a smartphone and on a laptop.

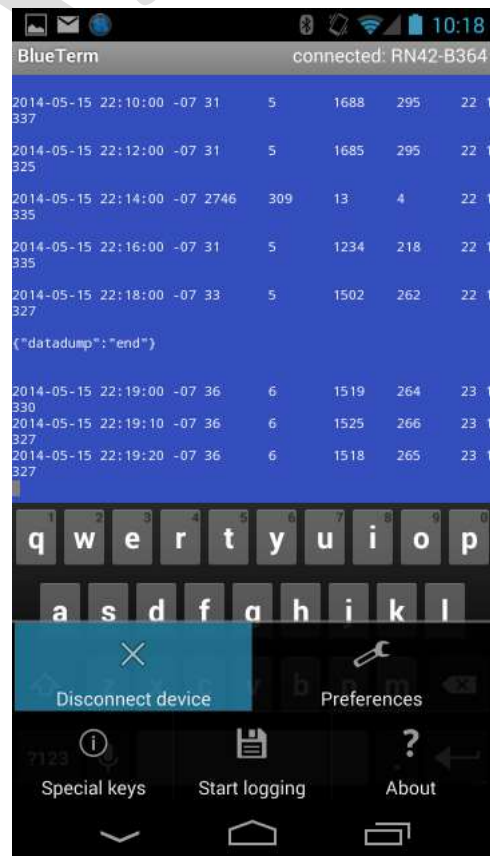
In this example, the "d" (data dump) command causes the logger to send stored data. The Android captures it to the log file.

When done, do "Stop logging" to assure the file is closed.

Finally, do "Disconnect device" to break the connection. The red LED on the logger starts blinking again.

You could keep logging while you connect to one logger after another, but this could yield a huge file. It is usually less confusing to stop logging and re-start when you connect to the next logger.

Metadata that specifies which data came from which logger is within the log file, so the file name is unimportant.



To copy the log files from your smartphone, sign up with AirDroid online, at airdroid.com.

Using AirDroid, you can get the files any time both your phone and your computer have network connectivity.

You could manually copy the files by cable, viewing your phone as a media device, but AirDroid avoids the confusion of where your files have "gone" in the sometimes-cryptic Android file system.






In some Android versions, newly created items, such as these log files, will not show up in the file system until the device is rebooted. They do show up in AirDroid.

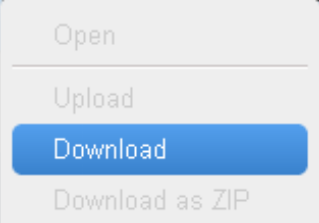
In your browser, AirDroid displays a console with many options.

Open "Files".



Your log files from BlueTerm will appear at the "bottom" of the default file location, listed below all the folders.

Name	Date Modified	Type	Size
 TerminalLog	05/13/2014 15:30	Folder	
 Toolbox	01/14/2014 19:52	Folder	
 Tumblr	09/07/2013 09:50	Folder	
 Video	12/21/2013 11:32	Folder	
 blueTerm_20140515_221721.log	05/15/2014 22:18	log	20.7 KB



A context menu is open over the log file, showing options: Open, Upload, Download (highlighted in blue), and Download as ZIP.

To copy to your computer, right click and choose "Download".

After downloading, you can also use AirDroid to delete these now-redundant files from your phone.

Laptop: Setting up a Bluetooth connection -- First one, first time

Setting up a Bluetooth connection to a laptop the first time is somewhat involved, especially the first one. After creating the connection, however, reconnecting is easy. Also, for additional loggers, you can copy much of what you have done on the first one.

These Bluetooth connections are "virtual serial ports" on your laptop, so you will need a program that talks to serial ports. We recommend "PuTTY".



<http://www.putty.org/>

Download and install PuTTY on your computer before proceeding with these instructions.

Read through the following procedure. Some of these steps need to happen in a fairly tight time sequence. The loggers are forgiving, but some computer operating systems don't give you very long to respond. You can always start again and re-do it. However it may take some untangling, to understand what you've already done, and what remains to do.

You wake up a logger by tapping it, as described above under "Rousing a logger to communicate". When roused, the red LED will blink (along with the yellow) to indicate the logger is available for a Bluetooth connection. You have two minutes to make the connection before the logger goes back to sleep.

The first time, you need to "create" the connection. If you are not able to complete the procedure in two minutes, you can rouse the logger again to finish.

From your computer's standpoint, you are going to "create a device". This is like an internal doorway on your computer that opens into the logger. To keep various loggers straight, you need to know which "device" is which.

The following screenshots are from Windows 7. The appearance and exact procedure will vary in other operating system versions, or with different view settings.

In Windows 7, look in Control Panel > Devices and Printers, to see which, if any, COM ports are already there.



In this screenshot, there are none yet. The "Bluetooth Module" shown is the part of your computer that *controls* Bluetooth. This is different from a Bluetooth serial port "device". If there were any such devices, they would (depending on view settings) probably show up with an icon as below.



If there are any on your computer, take note of what they are named so as to distinguish the new one we are going to add.

To add a new Bluetooth device, right-click on the Bluetooth icon in the tray at the bottom of the screen.

This will pop up the Bluetooth menu.

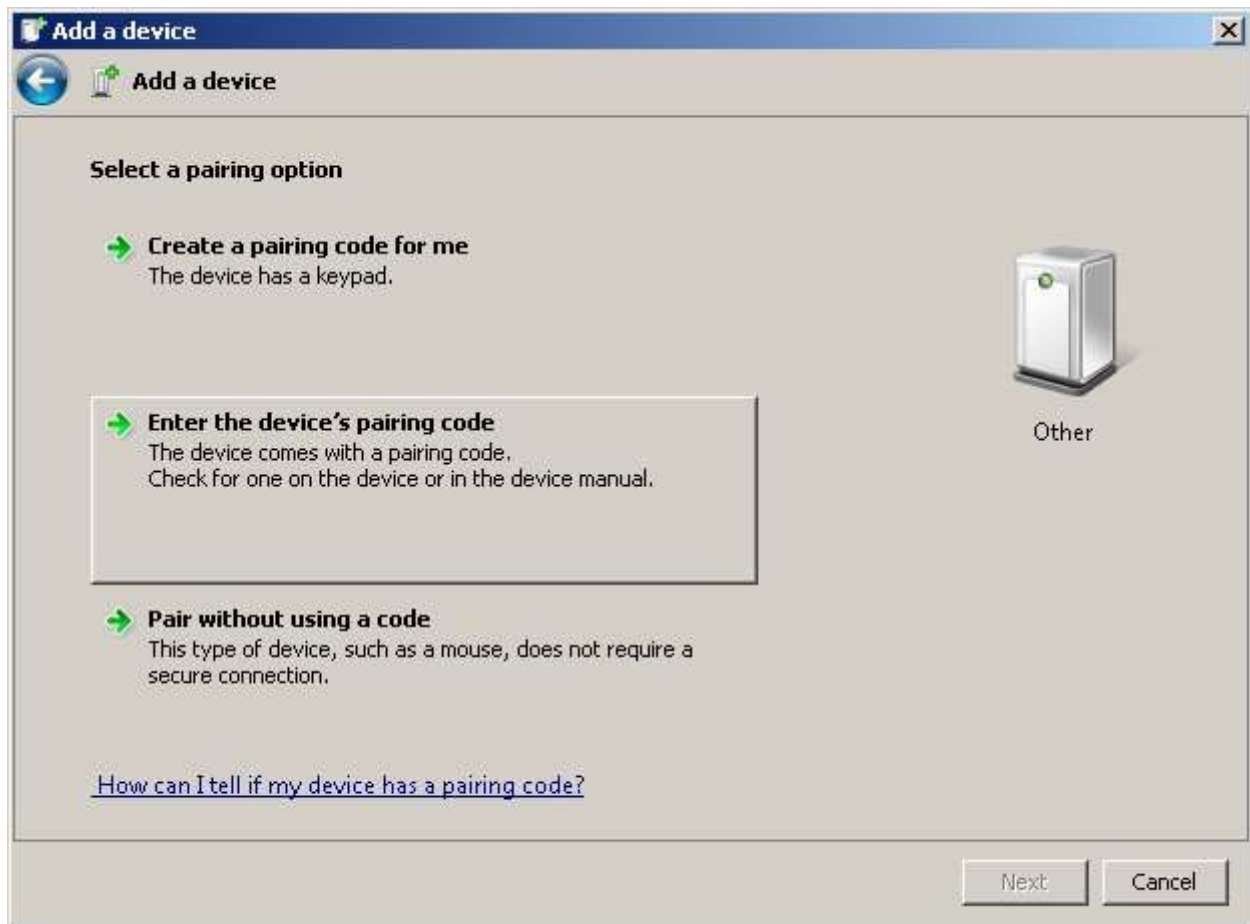
Choose "Add a Device".

When the logger's red light is flashing, its internal Bluetooth module will be sending a signal that your computer's Bluetooth controller can find.



Within a short time, this dialog should show a Bluetooth device. It may have a name such as "RN42-xxxx" where "x" is a numerical digit or a letter A to F. The device may only show as "Other", as in this example. This is the Bluetooth module in the logger that has its red light flashing.

Click on the device to select it, then click "Next".



The Bluetooth modules in the loggers have a pairing code. Click the "Enter the device's pairing code" button.

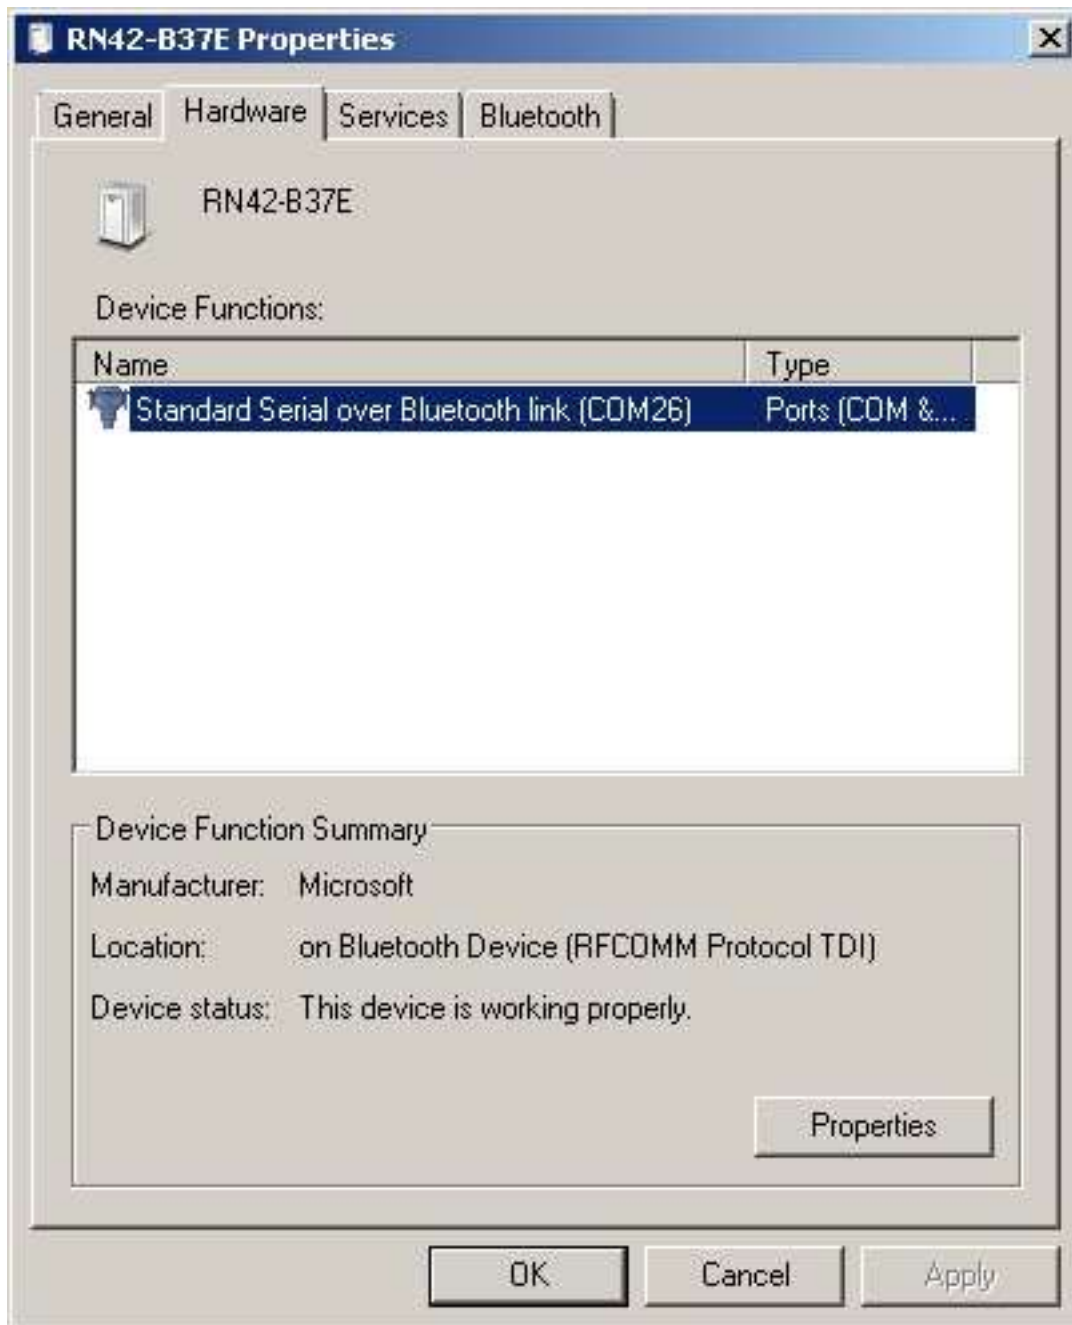


The pairing code is always "1234". Enter these digits and click "Next".



Now the control panel should show the new device. The name here is "RN42-B37E". This is the identifier, similar to a serial number, of the particular Bluetooth module in the logger you are connecting. "RN42" will be the same for all the loggers (unless we use a different part number in future versions), but the last 4 characters will be unique.

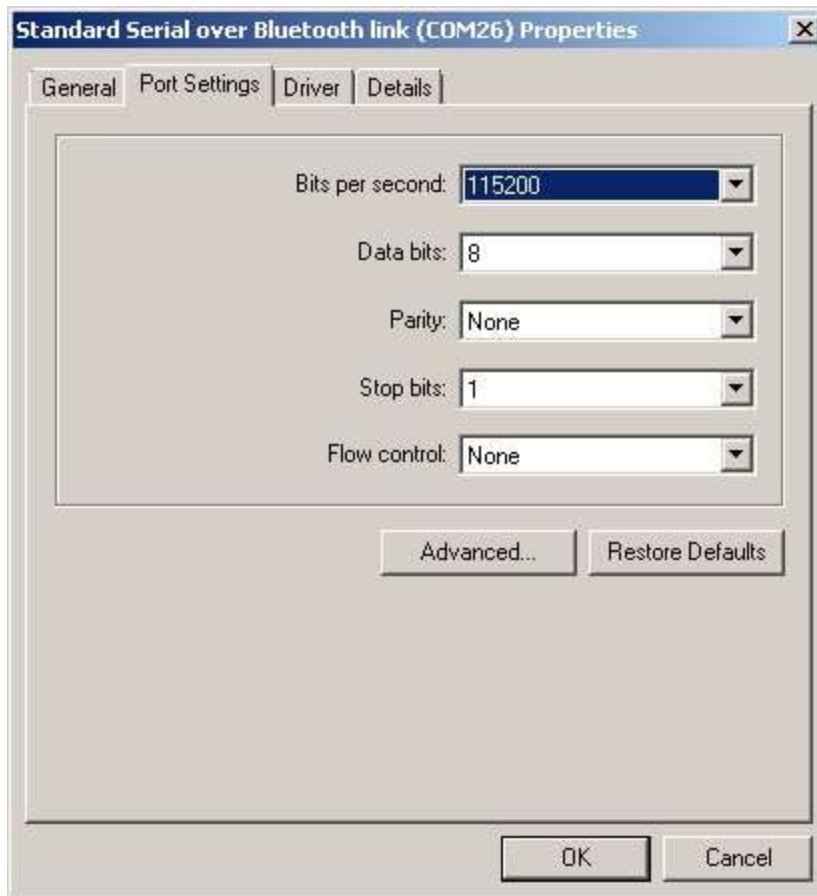
Your computer sees this as a "Device", but PuTTY will see it as a "serial port". You need to know what serial port it is. Right click on the new device. In the pop-up menu choose "Properties".



Look on the "Hardware" tab.

From your computer's standpoint, this is a serial port device, with a COM number. In this example the computer has assigned it as "COM26". Write down this COM number.

The properties of the port are probably set correctly, except for "Bits per second". Right click on the device name and from the pop-up menu choose "Properties".



Click on the Port Settings tab.

The settings should be:

Bits per second: change this to 115,200 (the default is probably 9600)

Data bits: 8

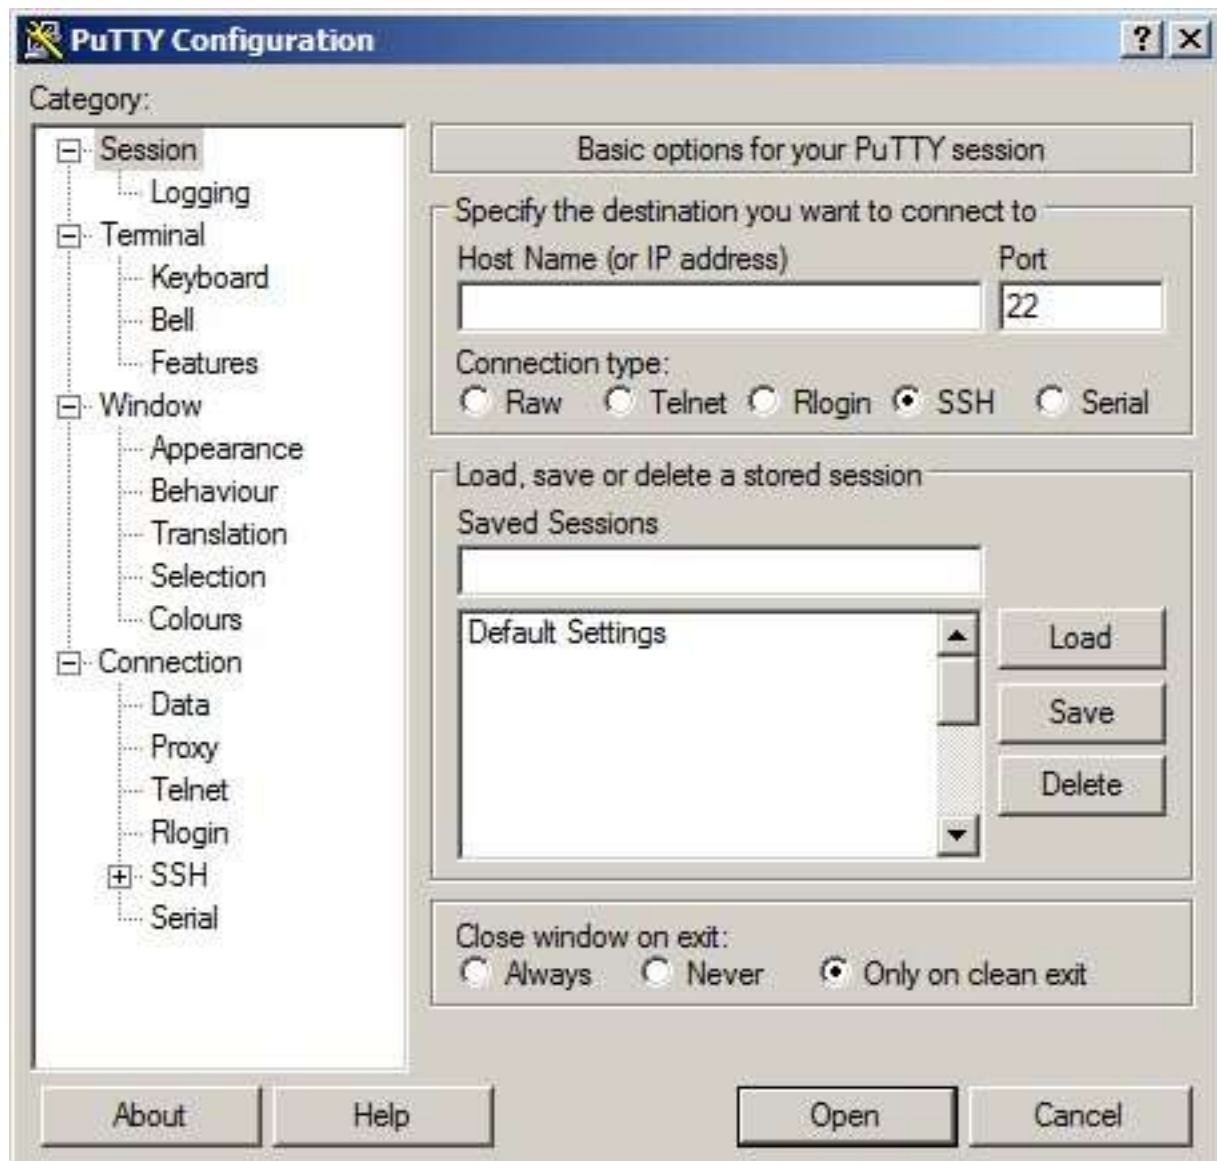
Parity: None

Stop bits: 1

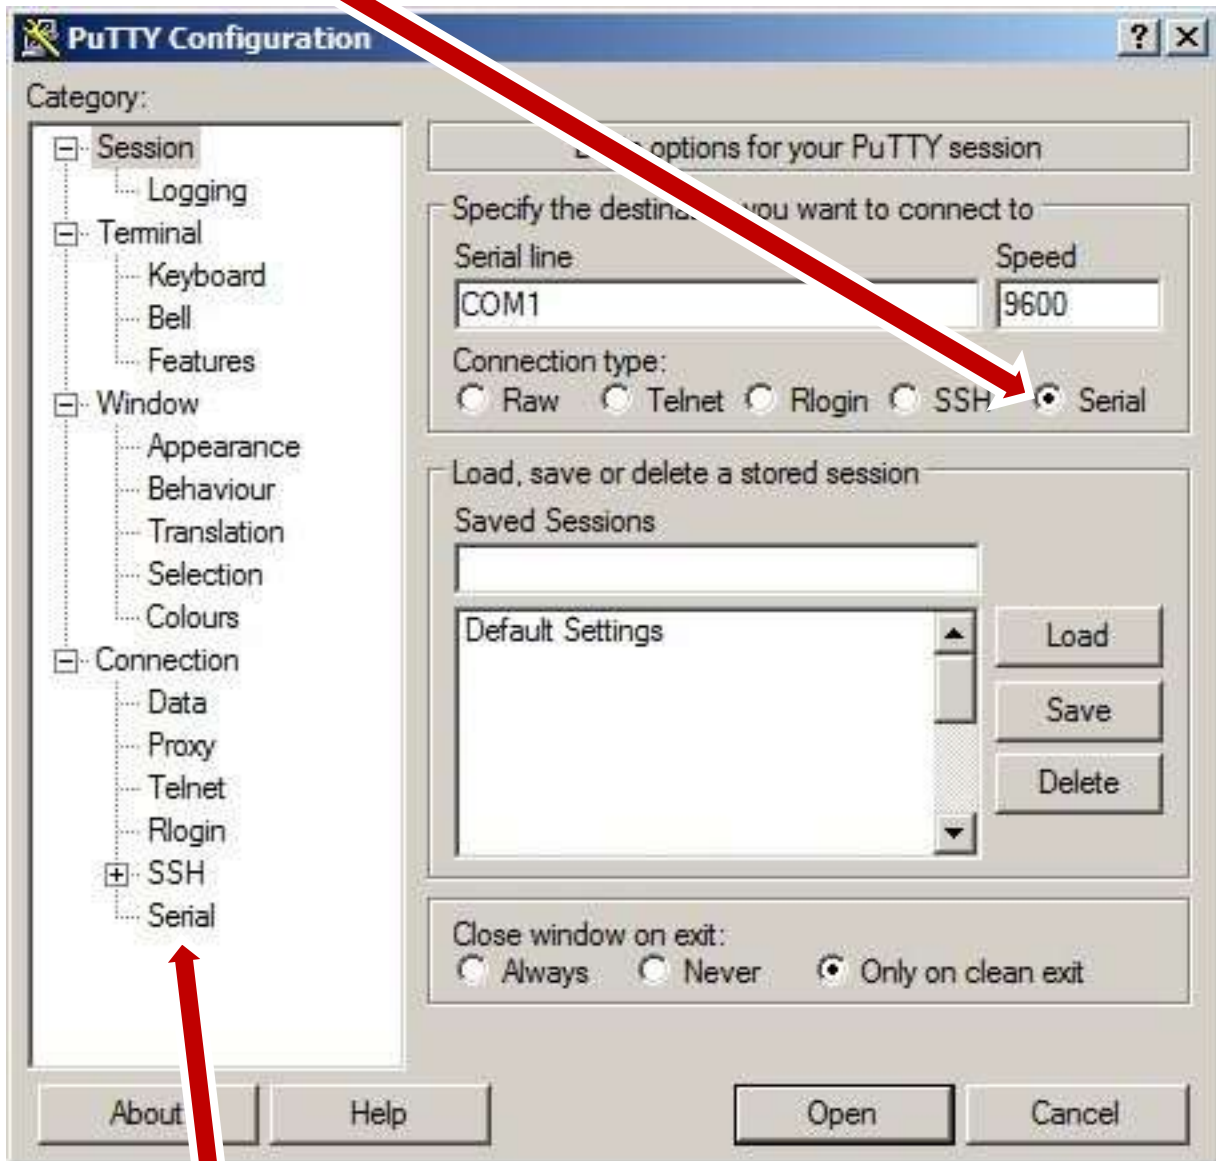
Flow control: None

Click "OK" in this and the other Properties dialog. The "device" is now set up on your computer.

Now we store this in PuTTY to automate the connection.
Start PuTTY. It comes up in the Configuration screen.

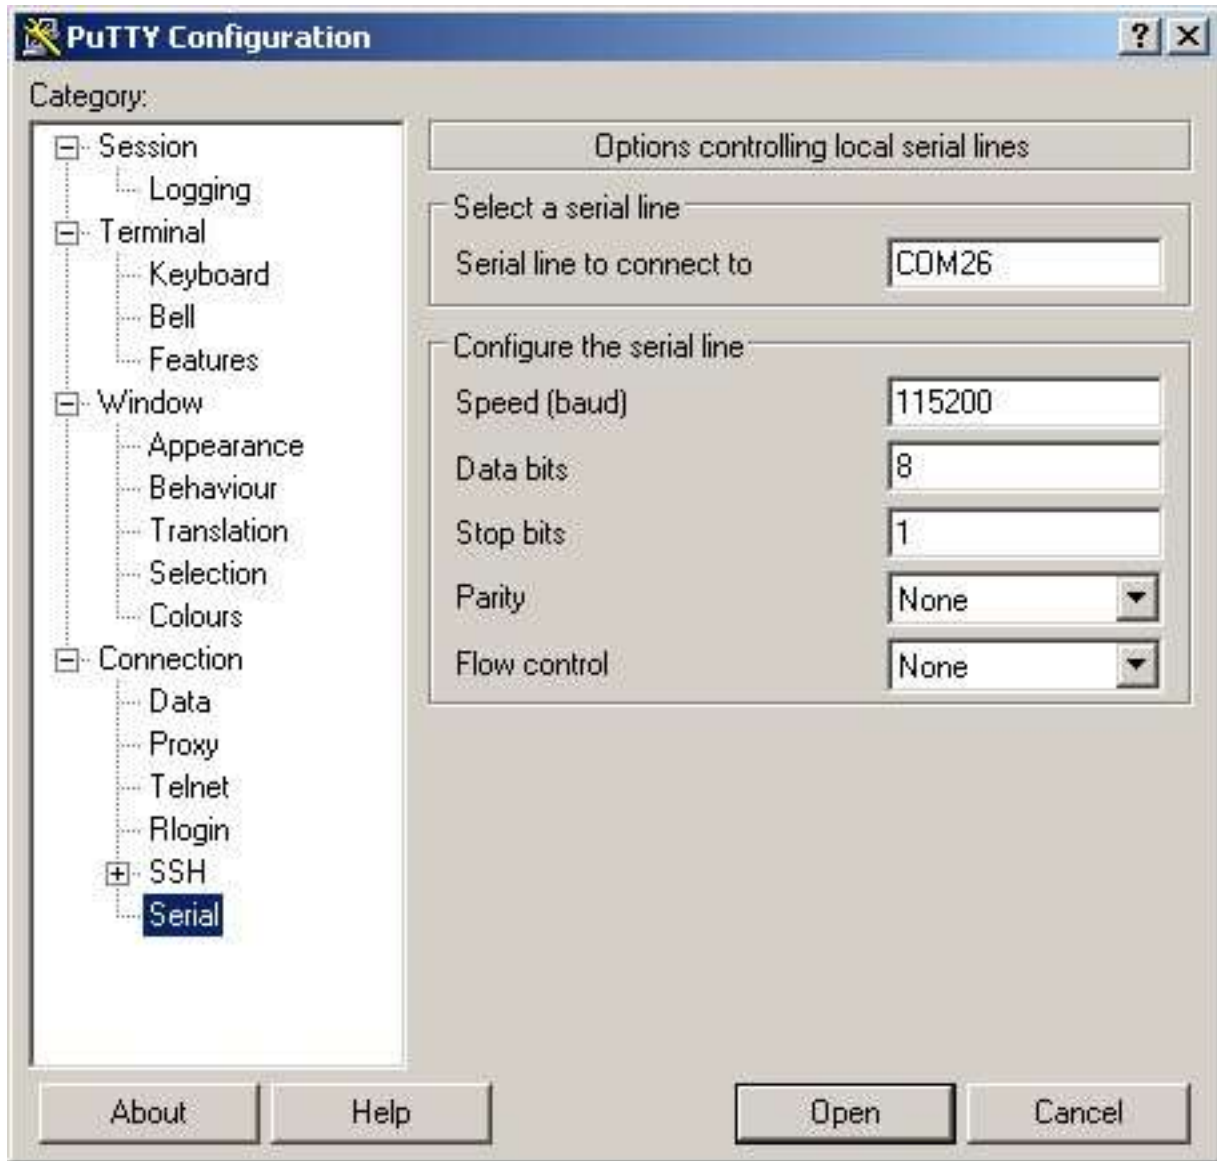


The default settings are for a different kind of connection than serial. To change this, click on Serial.



Then click on Serial in the left-hand column.

Make the settings the same as for the Device you set up earlier in Control Panel.



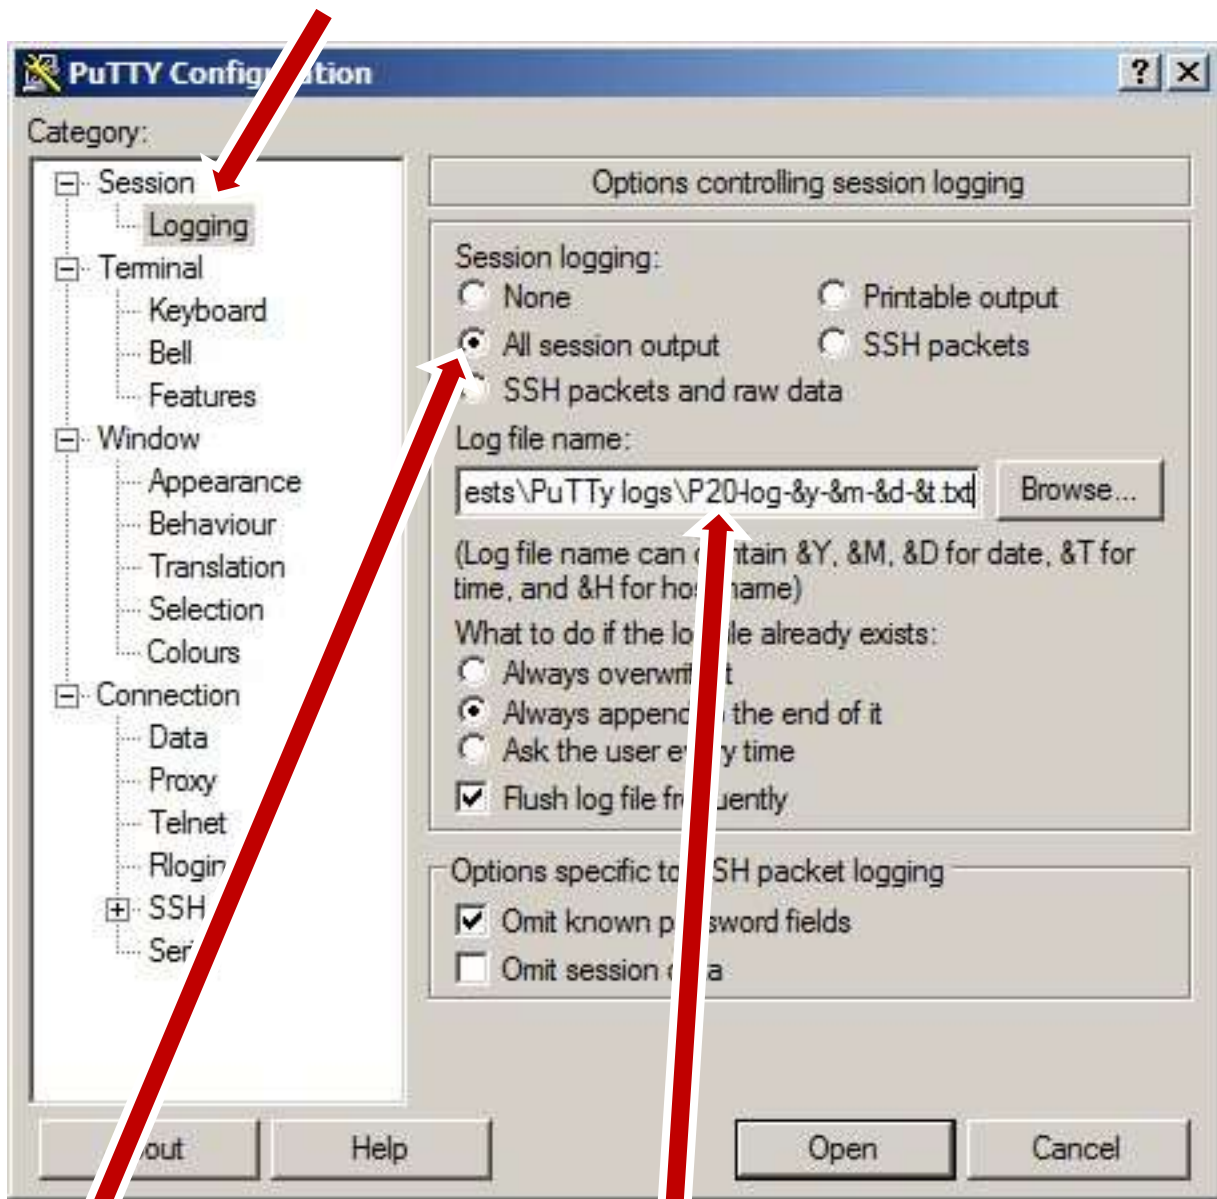
We are going to save this "session", and we want it easy to know which logger it is for.

On the logger, there is a label saying which logger it is. In this example, "P20" (which means prototype 20).



The label is in this position in circuit Version 12. It may be in a different place on other versions, but there should still be an identifier that will be easy to see when you want to connect to any particular logger.

Now, in PuTTY, click on Logging.



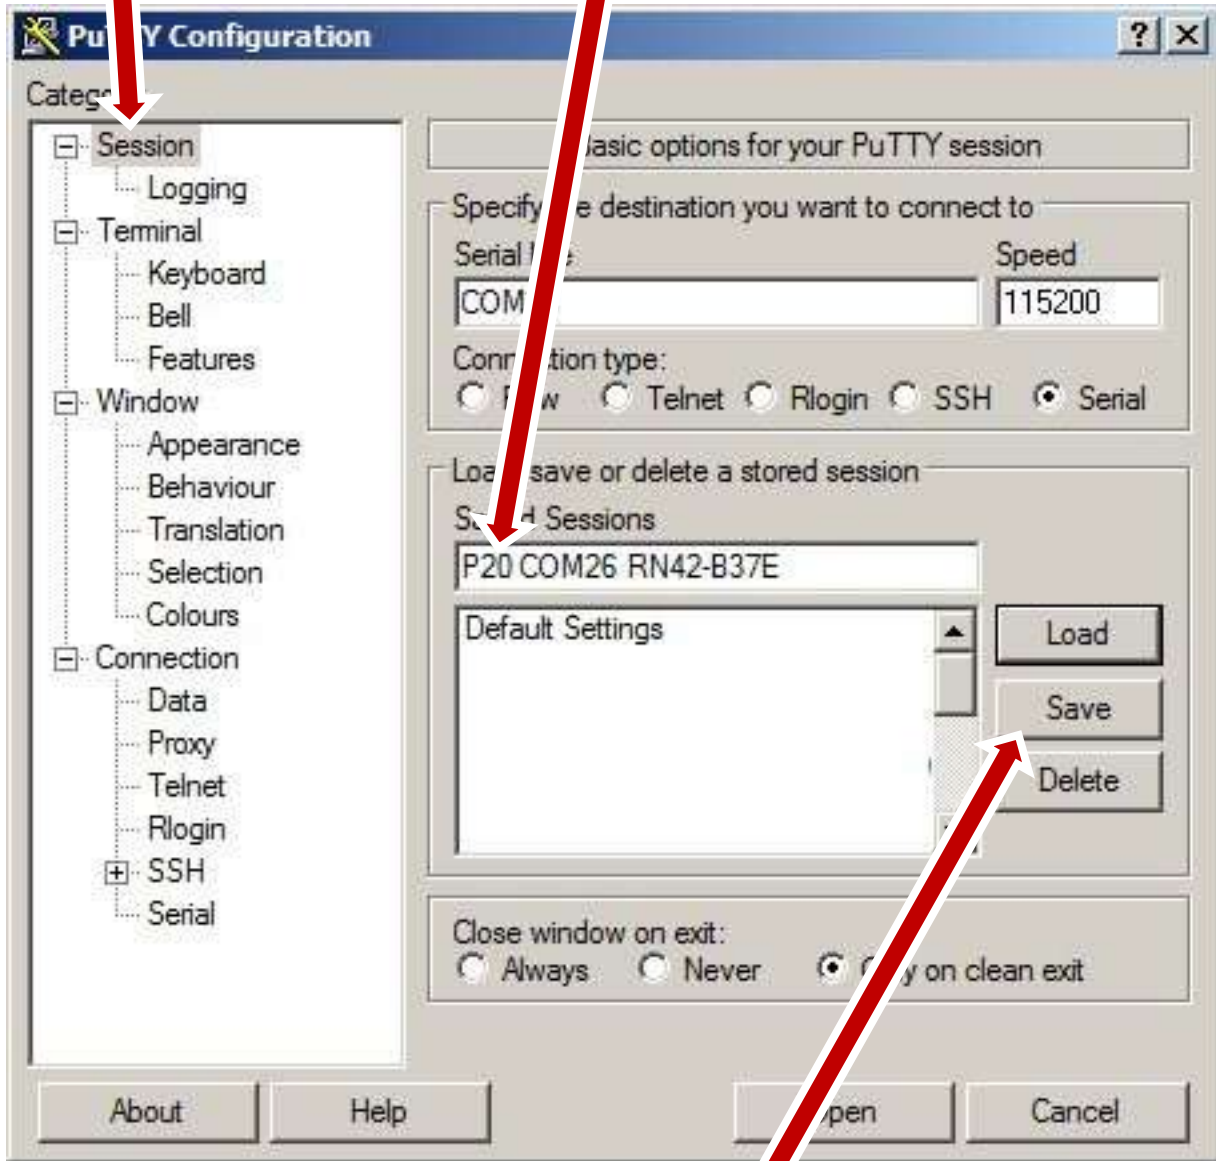
Pick "All session output".

Under "Log file name:", choose a folder. In the file name, put in the identifier for the logger.

The rest of the filename, "-log-&y-&m-&d-&t.txt", will create a timestamped file each time you connect. This will help with debugging, and is necessary to capture logged data.

Now, click on Session again.

Type in a name for the session. The logger identifier (e.g. "P20") is the most important part, so when you look at a logger, you know which session goes with it. However, putting all this other information here will make debugging easier.



Click "Save". This will save all these settings in one tidy package as a "session".
Close PuTTY by clicking "Cancel" (trust me).

Connecting to a logger after you have a session set up in PuTTY

Get a logger you previously set up a PuTTY "session" for. Tap the logger so its red LED starts flashing. If you just set up PuTTY for a new logger, the logger's Bluetooth connection has probably timed out and the logger has gone back to sleep. Tap it again to wake it up.

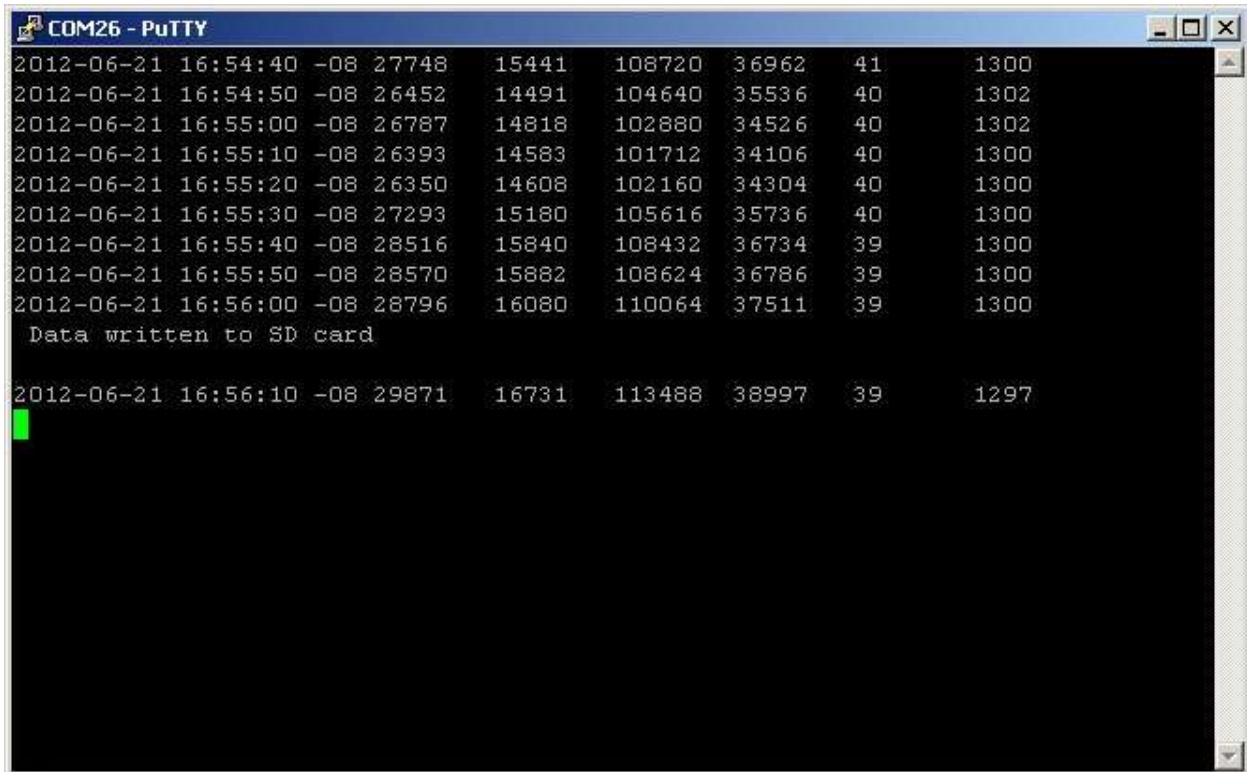
Open PuTTY. You should see the session you previously set up and saved.



Double click on it.

This is all you'll have to do to connect to that logger in the future.

Occasionally it fails and you have to try two or three times, but a window should open that shows the data stream from the logger.



```
COM26 - PuTTY
2012-06-21 16:54:40 -08 27748 15441 108720 36962 41 1300
2012-06-21 16:54:50 -08 26452 14491 104640 35536 40 1302
2012-06-21 16:55:00 -08 26787 14818 102880 34526 40 1302
2012-06-21 16:55:10 -08 26393 14583 101712 34106 40 1300
2012-06-21 16:55:20 -08 26350 14608 102160 34304 40 1300
2012-06-21 16:55:30 -08 27293 15180 105616 35736 40 1300
2012-06-21 16:55:40 -08 28516 15840 108432 36734 39 1300
2012-06-21 16:55:50 -08 28570 15882 108624 36786 39 1300
2012-06-21 16:56:00 -08 28796 16080 110064 37511 39 1300
Data written to SD card
2012-06-21 16:56:10 -08 29871 16731 113488 38997 39 1297
```

The numbers, of course, will be different from this example.

The red light on the logger stops flashing, and a green steady LED comes on instead. The data stream you see on screen is automatically being saved to the file you specified for "Logging".

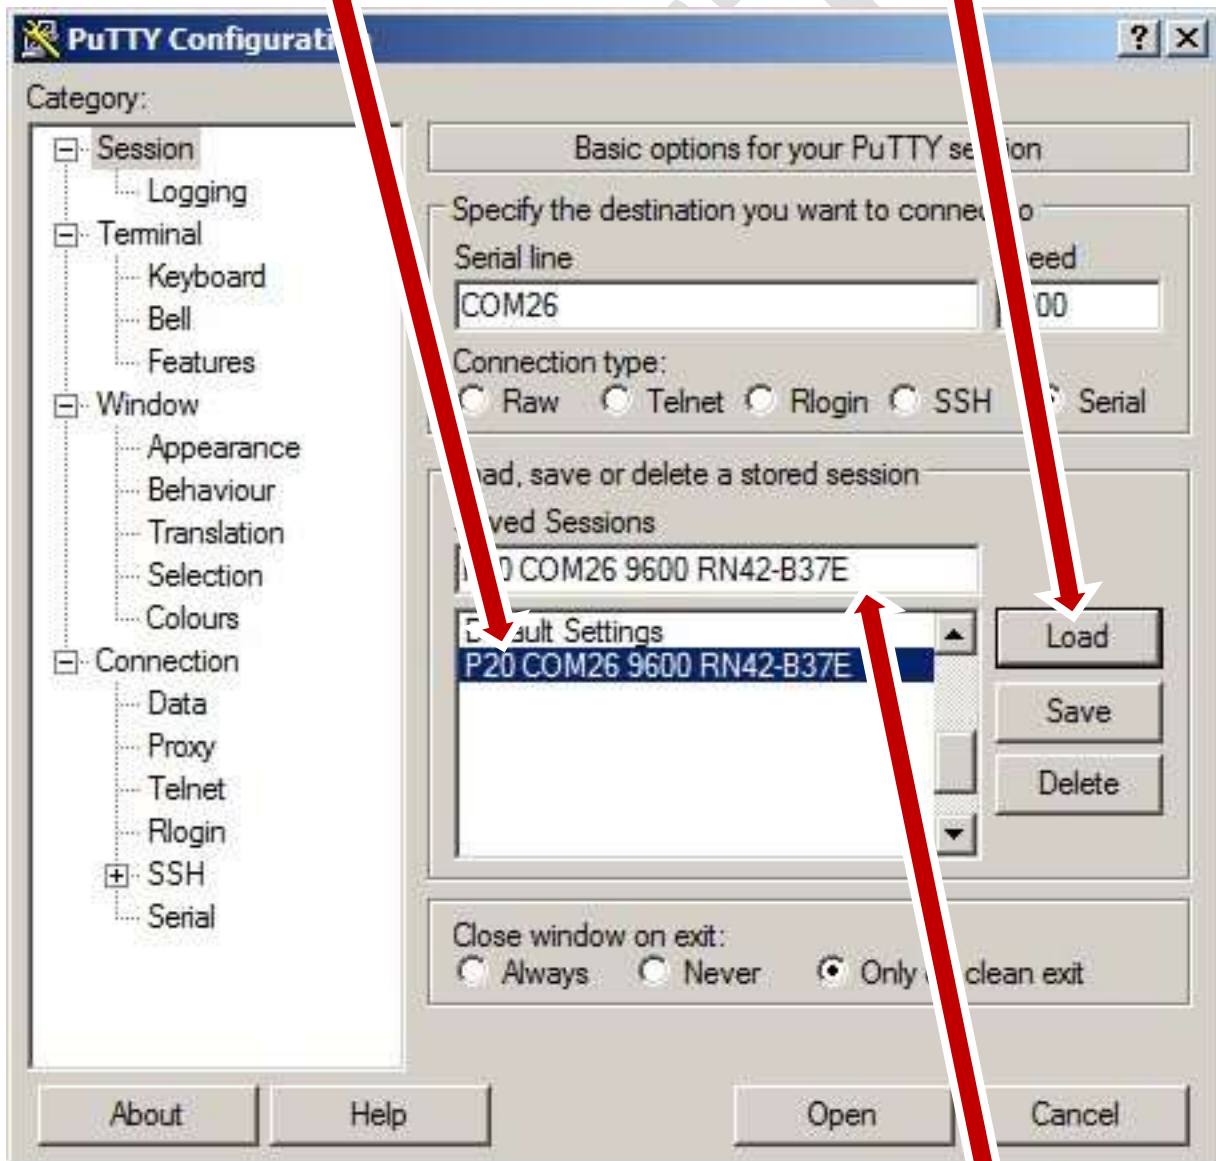
When you are done, close this window. The red light on the logger starts flashing again. You can reconnect, or let the logger time out and go to sleep.

Creating a PuTTY session for a new logger, after you already have a session for another logger set up.

You can copy most of the parameters for an existing session to create a new session, without having to start from scratch.

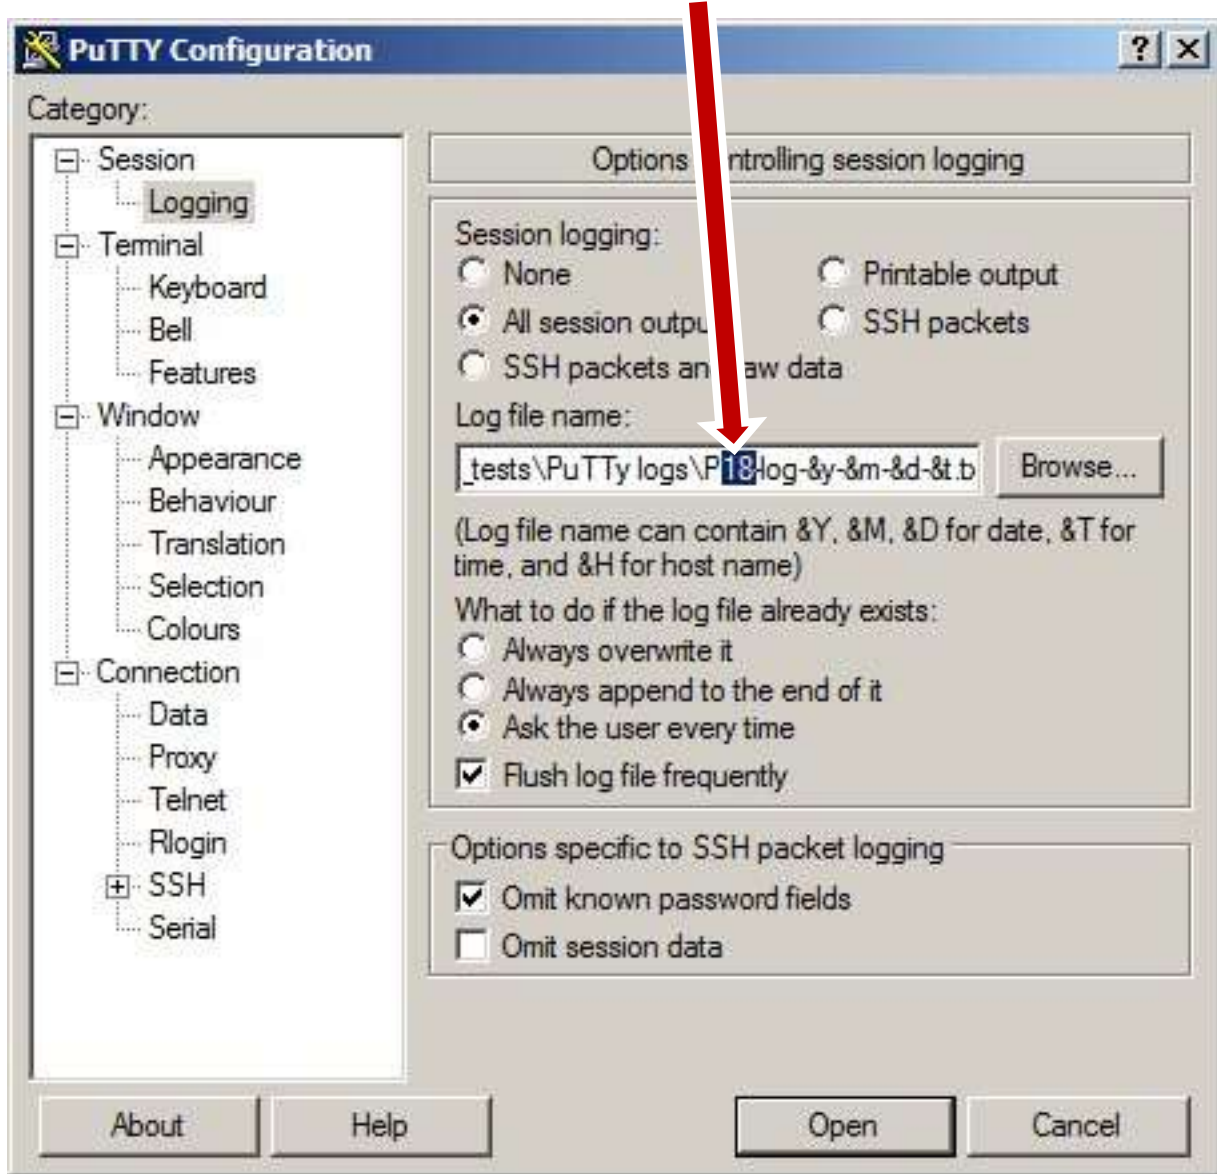
Set up the logger as a Bluetooth device, as previously explained. Take note of the COM number your computer assigns, and the identifier label on the logger.

Open PuTTY and click *once* on the saved session in the list. Then click "Load".

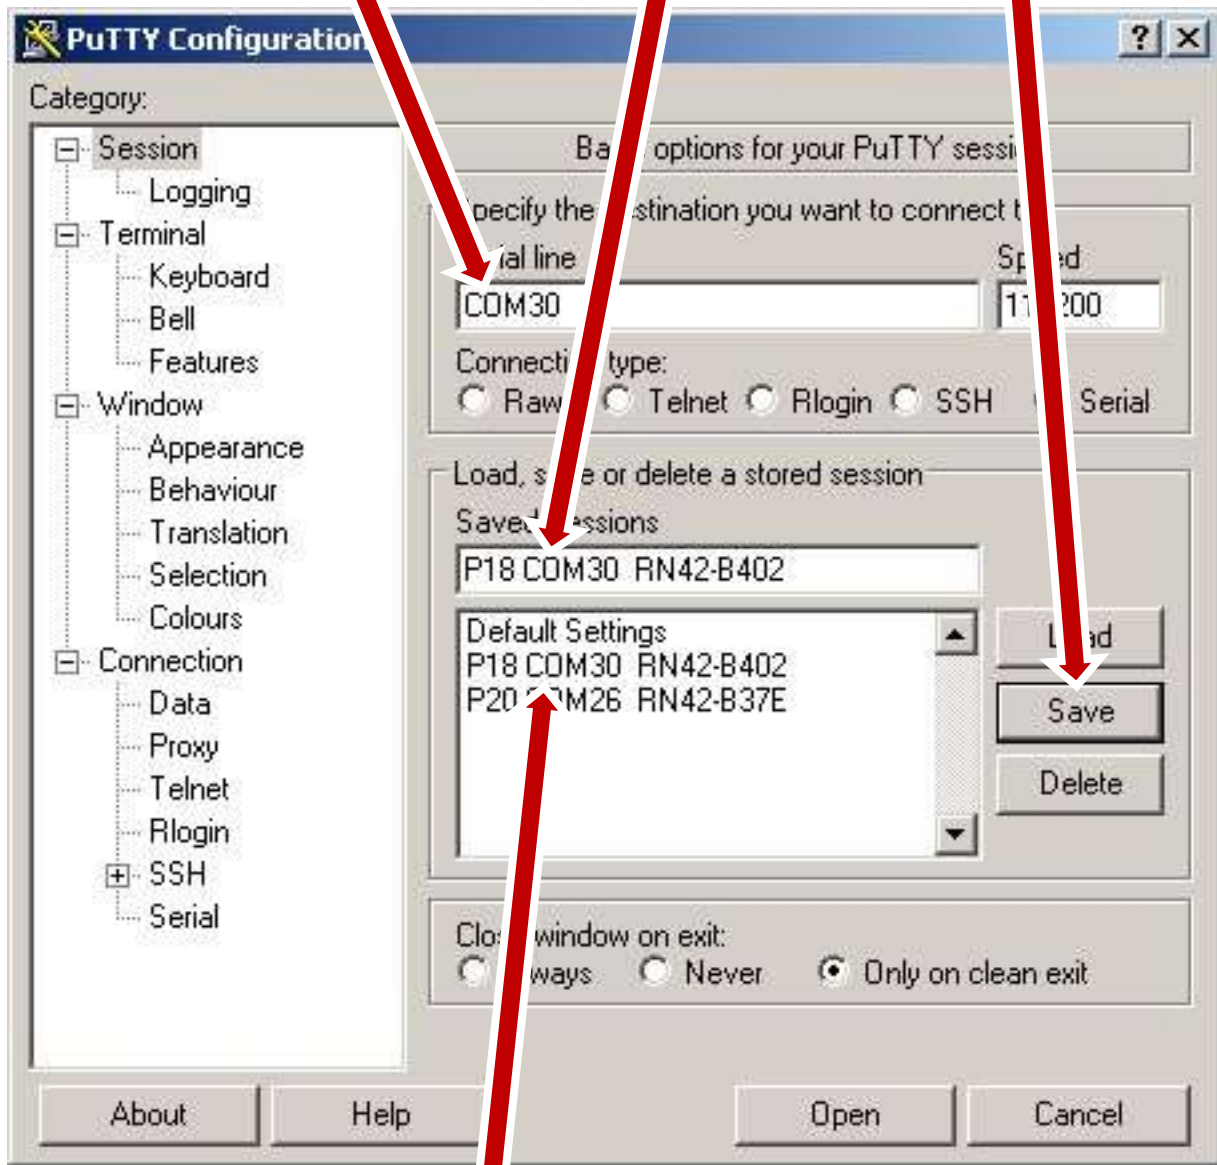


The name appears under "Saved Sessions", and all parameters are loaded.

Under "Logging", adjust the filename for the new logger identifier.



Adjust for the new COM number, and change the session name. Then click Save



The new session is added to the list.
You can now double-click it to connect to that logger.

You can open connections to multiple loggers at the same time, as long as you use the right session for each one.

What the numbers mean

Here is an example of a logger's data stream when you monitor it using BlueTerm on a smartphone, or PuTTY on a laptop. This is much the same format as stored on the logger's internal SD card (except with certain diagnostic messages inserted, and things like the "Data written to SD card" lines left out).

Timestamp	BBDn	IRDn	BBUp	IRUp	T(C)	Vbatt(mV)
2016-06-21 16:55:00 -08	26787	14818	102880	34526	40	1302
2016-06-21 16:55:10 -08	26393	14583	101712	34106	40	1300
2016-06-21 16:55:20 -08	26350	14608	102160	34304	40	1300
2016-06-21 16:55:30 -08	27293	15180	105616	35736	40	1300
2016-06-21 16:55:40 -08	28516	15840	108432	36734	39	1300
2016-06-21 16:55:50 -08	28570	15882	108624	36786	39	1300
2016-06-21 16:56:00 -08	28796	16080	110064	37511	39	1300
Data written to SD card						
2016-06-21 16:56:10 -08	29871	16731	113488	38997	39	1297
2016-06-21 16:56:20 -08	31940	17925	119840	41283	39	1300

Note that the write to the SD card is on the even number minute rollover.

This is what the column headings mean:

Timestamp: The time, from the system's internal clock, in Unix format. The clock is rated to 20ppm accuracy (roughly a minute per month) and has performed better than this in testing. The final part (e.g. "-08") is the timezone offset from Universal time (Greenwich Mean).

Examples are:

- 08 Alaska Daylight Time
- 07 Pacific Daylight Time
- 06 Mountain Daylight Time

You can check the time from <http://www.time.gov/>

The loggers have an internal GPS unit so they can get their time (and location) automatically. More on this later.

The following four columns are data from the irradiance sensors (see Appendix A). The unit of measurement is relative, but NDVI is normalized and so is itself dimensionless.

The first 2 columns are irradiance reflected up from below. These are from the down-pointing sensor that looks through the collimator.

BBDn Broadband, down-looking.

IRDn Infrared, down-looking.

The next 2 are reference irradiance, coming from the sky. These are from the sensor on the top of the logger, covered by the Teflon diffuser

BBUp Broadband, up-looking.

IRUp Infrared, up-looking.

T(C) Temperature, degrees Centigrade, from a sensor inside the instrument case. This allows compensating for the temperature response of the irradiance sensors, which is specified by the manufacturer. Also, if the circuit fails, this can give an indication whether it happened from high (or possibly low) temperature.

Vbatt(mV) The voltage, in millivolts, of the power storage, which is a single rechargeable NiMH cell, size AA. This data column allows tracking how well the solar cells maintain power under different conditions. The system monitors this voltage and disallows functions if too low (e.g. the SD card, which could scramble if there were not enough power to complete a write cycle). Below 1170 is low. The range of 1200+ is pretty good. The range of 1300+ is very good. The range of 1400+ is about as good as it gets.

Controlling loggers through Bluetooth

The serial data screen (colored blue in BlueTerm, black in PuTTY) where you see a logger's data stream is the main interface for controlling a logger. You do this by one-letter commands, sometimes followed by other information. All the letter commands can be either upper- or lower case.

D	data	dump data
G	get time	get time from the GPS
L	leveling	read the accelerometer
T	time	set the time manually
V	version	report the firmware version

As soon as you start typing anything, the normal data stream stops and you get an arrow prompt showing what you typed. To quit out of this, send <Enter>. If the data stream stops and you get an arrow prompt, the system interpreted something as input, so just send <Enter> to cancel it.

The "D" ("Data Dump") command is for getting data out of the logger.

This command causes the logger to output stored data, which will be captured to a log file, either using BlueTerm on a smartphone or PuTTY on a laptop.

In the future, we plan to implement data transfer using a wireless file server, if power requirements can be kept low enough. Presently, you get the data by the old method of simply causing the data to dump, and capturing that stream to a file. The data stream is tab delimited text, with JSON messages (explained below) to signal special conditions.

Data are always dumped for entire days. Each day takes about 5 seconds, for a day that has 12 hours of daylight. Spring and summer days from high latitudes have longer periods of daylight, and therefore more data records, and so their dumps will take longer.

"D" alone will dump the data since the last dump, which is usually what you want. For example, if you dumped data yesterday, and then did "D" today, you would get the data for yesterday (because data always goes by whole days) to complete that day, and then all today's data up to the present. You would have to filter out the redundant records, but this is preferable to missing some data.

"D" followed by a valid date will give you the data for only that date. For example:

```
d2016-05-02
```

...will dump the data for only May 2, 2016. This sets the last-dump date to May 2nd, so the next "D" alone will start from May 2nd. Dumping a one-day sample first avoids a long tie-up if you chose the wrong date.

If there are no readings for the chosen day, you will get a message saying this. To see a list of the dates that have data, use the command "DD" ("Data Dates").

If you ever want all the data in the logger, use the command "DA" ("Data All"). This can take quite a long time. Remember the rule of thumb of about 5 seconds per day's data. A year's data will take the better part of an hour to dump. The instrument continues logging data during one of these long dumps.

If the Bluetooth connection depletes the system power below what's needed to read the SD card, the dump will stop. It will remember the last date dumped, so "D" will resume from that point after the solar cells have recharged the system power enough.

All stored data are recorded as ASCII text files on the SD card, so they can be retrieved manually. If you want to do this, please contact the developer before opening the instrument case.

"JSON" means "JavaScript Object Notation". It is a coded format that the logger uses to record messages that are not part of the normal data stream. For example, when a logger updates its time setting, the logger writes into the current data file a JSON message like this:

```
{"timechange":{"from":"2017-04-07 03:06:35 +00","to":"2017-04-07 03:06:33 +00","by":"GPS"}}
```

It is both visually obvious when you are viewing the data, and easy to detect if you are automating data processing and want to either find such messages or ignore them.

The "D" command itself also inserts JSON messages, to mark the beginning and end of a dump.

```
{"datadump":"begin"} ... {"datadump":"end"}
```

The "G" ("Get from Gps") command gets time from the GPS.

The loggers' programming has heuristics to periodically acquire an accurate date and time by using its internal GPS, but you can force the logger to get the time using the "G" command.

The logger performs a sequence that activates the GPS and acquires the time and location from it. On the logger upper board, a set of LEDs flash to track progress. The procedure typically takes about 30 seconds.

Since GPS data comes from satellites, you typically do this outdoors, at a field site. This GPS acquire may work inside a building that has one level of roof overhead, such as wood frame house; but it is likely to fail inside a large structure such as a laboratory building. In this case, the procedure tries for 3 minutes, then stops.

The sequence is this:

Green LED	One or more brief flashes, to power on the GPS.
Yellow LED	Very brief flash to activate the GPS.
Orange (or red) LED	Continual blinking while acquiring and parsing GPS data, working to extract valid time and location information.
Yellow LED	Very brief flash to deactivate the GPS. Information successfully acquired, or procedure timed out.
Green LED	One or more flashes, to power down the GPS. Then all lights go off.

Error code	Very brief flash on power-up, means the upper board is not properly connected.
------------	--

If successful, the GPS sets the system time by sending the same format of "time" command as for setting the time manually (see below). The timezone is always "+00" which is standard Universal Time, UT, (Greenwich Mean Time). This is convenient as it is standardized for anywhere on earth, but you should understand that the hour will be different from your local time, unless you happen to be in the UT timezone.

The logger tries to acquire the GPS time periodically, so you may see the indicator LED sequence occur spontaneously. This happens every few days in normal operation, to keep any time drift to less than 10 seconds. The logger also attempts to acquire GPS time in other cases, such as recovering from power failure, or on being taken out of long storage.

In normal use, the logger tracks its cell charge, and tries to run the GPS when most solar power is available. If it has been moved a long distance, or been in storage a long time, it may take several days to re-optimize.

The "L" ("Leveling") command is for leveling the logger.

This will output the X, Y, and Z axis numbers of the logger, showing any tilt. This verifies the internal accelerometer is working.

You can use the Leveling command to orient the logger's view normal to the ground. This command puts the logger into Leveling mode, where the data stream shows the X, Y, and Z axis numbers continually. When a logger is level, the X and Y should be close to zero, while Z will be in the negative 200 range (unimportant).

```
X = 0, Y = 0, Z = -260
X = -1, Y = 0, Z = -260
X = 0, Y = 0, Z = -260
X = -2, Y = 0, Z = -261
X = -3, Y = 0, Z = -259
X = -2, Y = 0, Z = -260
X = 0, Y = 4, Z = -261
X = -1, Y = 1, Z = -260
X = 0, Y = 0, Z = -260
X = 0, Y = 0, Z = -260
X = 3, Y = 0, Z = -259
X = 0, Y = 0, Z = -260
X = 4, Y = 0, Z = -261
```

In some early prototypes, X and Y may be a few tens of points off level, due to quality control in soldering the accelerometer chip in place. However, returning the X and Y to zero after moving a logger will still restore it to its previous consistent orientation.

The logger will go out of Leveling mode in a few minutes, once X, Y, and Z are no longer changing. Also, any serial input will cancel Leveling mode, so you can just send <Enter> when you want it to stop.

The logger writes the axis numbers into each day's data, as a check that the device was in a reasonable orientation on that date. This would allow you to eliminate invalid data, such as if a logger were knocked off its mount during a long field deployment. You could also see on what date this happened.

Normal data logging continues while the logger is in Leveling mode.

The "T" ("Time") command is for manually setting the time.

You can manually set the time in the logger. You seldom need to do this, as the logger automatically acquires the time from its internal GPS. You might use this command if you are testing a logger indoors where it cannot get GPS satellite signals.

To manually set the time, enter a series of characters in this exact format (this is an example):

```
T2016-05-02 13:56:30 -07
```

The first character is a "T" (lower- or uppercase). Then, the 4-digit year, 2-digit month, and 2-digit day separated by dashes. Then a space and the 2-digit hour, minute, and second separated by colons. Then a space, and the timezone offset preceded by a plus or minus sign and zero-padded out to two digits.

A quirk of the BlueTerm app is that it translates input colon (":") into semicolon (";"). To accommodate this, the set-time command accepts semicolons as well as colons.

You can get the correct time from

<http://www.time.gov/>

On this site you can choose your location, and it will tell you the numeric timezone offset. Numeric is important to avoid confusion over daylight saving time. For example Pacific Standard Time and Alaska Daylight Time are both -8. "Pacific Time" and "Alaska Time" are ambiguous, as the hour offset would be different by time of year, but -8 is unambiguous.

The format for the "T" command is very strict. If you get it wrong, the command will fail and the old time setting will be retained. Backspace to correct input errors does not work yet. The time will be set at the moment you send <Enter>, typically within a fraction of a second.

When the logger next acquires a GPS time, it will overwrite any manually set time, on the assumption that GPS time is by its nature very accurate. A GPS time will always have a timezone offset of zero.

The "V" ("Version") command tells the current firmware version.

The current firmware version, as of this writing, is version 3.00.


There may be other obscure diagnostic commands, not documented yet.

DRAFT

Troubleshooting

There is minimal troubleshooting you can do without opening the instrument case. Please contact the developer before opening the case.

If the logger does not appear to be behaving correctly, you can try resetting it. A failure requiring a reset is inherently undefined, but often the logger will recover, even retaining its clock time.



You reset a logger by holding a strong magnet by the lower left front edge, as shown in the picture. A type of magnet that works well is a 3/8 inch diameter rare earth disk magnet, with poles on the faces. The dark rodlike component on the circuit board inside the case is a magnetic reed switch. The sensitivity is greatest in the position shown, with the magnetic field lines going the length of the switch. This causes an internal contact, and produces a system reset.

When you take the magnet away, the logger comes out of reset. If it has sufficient battery power, the logger will come out of reset searching for a Bluetooth connection. The red light will be flashing. It will have an especially long timeout of 3 minutes in case you need to set up a new connection.

If you previously set up a PuTTY session for this logger, or paired it with your Android, use the same session or pair. Resetting the logger does not invalidate it. However, you will need to close any PuTTY window for that logger and reconnect.

DRAFT

Moisture

The Greenlogger has a desiccant pack inside to absorb any moisture. This is one of the reasons you should not open the case. The desiccant pack would hydrate from ambient air and stop working.

If condensation does appear inside the instrument case, please contact the developer.



Notice that the central interior of the case, which contains the electronics, is the only area that will stay dry. In wet environments, moisture will normally get into the mounting channels, which run vertically up the corners of the case. If you look carefully, you will see that these do not connect with the case interior.

Appendix A: Irradiance sensors and NDVI

The rationale for NDVI

Normalized Difference Vegetation Index ("NDVI") is a "greenness index". In simple terms, plants absorb visible light for photosynthesis, and they reflect infrared because it is no use to them. So, higher infrared reflectance relative to visible means greater "greenness".

NDVI from satellite bands can give NDVI maps of whole continents. At the other end of the scale, clip-on instruments can judge plant health from individual leaves. The Greenlogger works at in intermediate range, tracking "greenness" for a small piece of an ecosystem, on the order of one meter in size.

The NDVI formula

To look at how much infrared reflectance there is compared to visible, the simplest math would be subtraction:

Infrared - Visible

The more infrared you have, and the less visible, the higher the number. However, natural light levels constantly change, such as by time of day and weather. So this number would vary by overall intensity. To compensate for that, the difference is normalized by dividing by the sum.

$(\text{Infrared} - \text{Visible}) / (\text{Infrared} + \text{Visible})$

This is the essence of "Normalized Difference Vegetation Index".

Spectral bands

Detectors to directly deliver "visible" and "infrared" would be complex and expensive, utilizing spectrophotometry, or specialized narrow-pass filters. The Greenlogger uses a common technique of numerically processing the inherent spectral response of simple sensors, photodiodes. Other electronic sensors such as phototransistors or cadmium sulfide photocells are available, but photodiodes are preferred for their linearity and stability.

Hardware components

If you are familiar with measuring irradiance using photodiodes, you understand the issues involved. Photon flux produces a tiny current, typically on the order of microamps, which must be carefully shielded enroute to any amplifier/digitizer. The usual approach is to convert the current to a voltage using a transimpedance amplifier, then do analog-to-digital conversion on that voltage. The transimpedance parameters must be rigorously specified, or you will have to custom-calibrate each circuit you build. Your circuit must of course deal with the photodiode's frequency response, biasing, and junction capacitance. For NDVI, you have two photodetectors, the "visible" and the "infrared", so you have to assure they "see" the same thing.

The Greenlogger bypasses all these headaches by using an integrated electronic component, the TSL 2561 light-to-digital converter. It gives direct digital output, and the Greenlogger programming auto-ranges the numbers so the logged output is linear over several orders of magnitude (moonlight, to direct full sun). The two photodiodes are together in the same tiny package, so they inherently see the same view. The logger reads them by switching "channels".

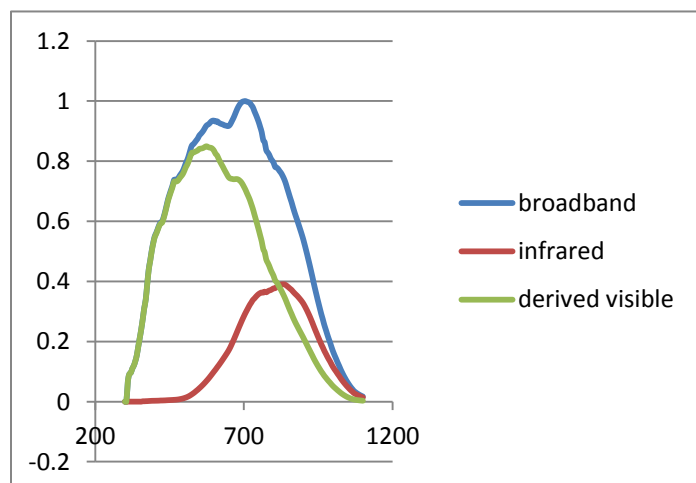
Numerical processing

The spectral response of commonly available photodiodes is not inherently ideal for "visible" and "infrared" bands. What follows is a common technique to better approximate these bands. The derivation presented here is the very simplest. Other more complex methods are possible, analogous to the Lux formula in the TSL 2562 data sheet.

The "Channel zero" ("Ch0") photodiode in the TSL 2562 responds to wavelengths throughout the visible and into the infrared, with a peak around 700 nm. This is shown in the chart as a line labeled "broadband". This peak is near the red limit of the visible spectrum, and of itself is not a very satisfactory "visible" band.

The "Channel one" ("Ch1") photodiode has attenuated response in the visible, with a peak around 815 nm. This is well into the infrared, and is so labeled in the chart.

The simplest derivation of "visible" is to subtract the "infrared" from the




"broadband". This shifts the peak response to about 580 nm, well into the visible. A more elaborate method might use weighting factors, but let's see what we get from these basic calculations.

In this rudimentary derivation, the "infrared" response is less than the "visible", however it is useful to look at some example data.

Below is a Greenlogger recording from a near worst-case day. The weather was rainy and cloudy, with a few sunbreaks.

"BB..." is shorthand for "broadband", "IR..." for "infrared", and "VS..." for the derived visible. The "...Dn" notation means "down looking", which means the reflectance from looking down at the vegetation. Later we will include "...Up", ("up looking"), which is the direct irradiance as looking up to the sky.

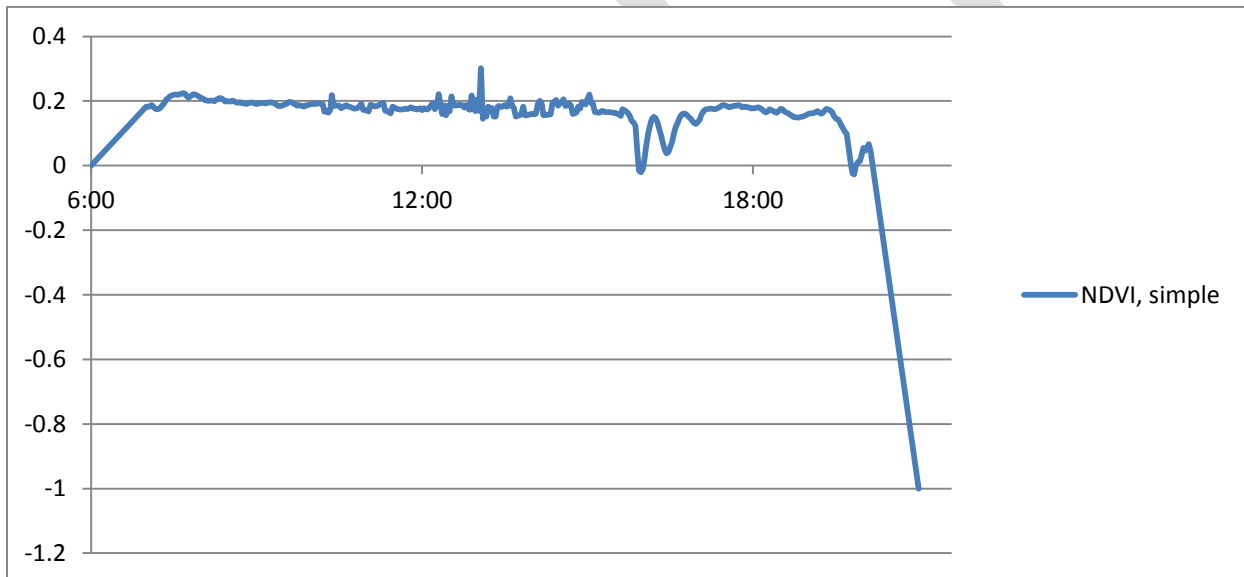


Data like these are a mess. Spectral quality changes slightly during sunbreaks. Typically, there is significant reflectance off bright cloud edges. You can think of this as dropping random copies of the sun around the sky, over the blue-sky background, and making the ambient daylight "whiter" or "redder". The cloudy times can be spectrally anomalous in various ways, not to mention the reddening at dawn and dusk. Wet vegetation reflects somewhat more like plain

water, varying depending whether the moisture is frosty or dewy or a continuous shiny layer. Still, with all these drawbacks, let's see what we get.

Notice that, even though the sensor *response* to infrared is lower than visible, the infrared *signal* is higher. This illustrates that most of natural daylight is infrared and gives some confidence that this simple derivation corresponds to reality.


Charting the NDVI formula, even this simple derivation gives a fairly stable trace. As expected, the numbers go off the rails at the very limits of dusk and dawn, when sensor readings approach zero. The major mid-day excursions probably correspond to sunbreaks, but the traces during early morning and late evening hours are quite level. A common protocol in ground-based NDVI is to ignore all but mid-day readings, and further filter by heuristics that reject too-high, too-low, or too rapidly changing readings.



Spectral compensation

Each Greenlogger contains two TSL 2561 components, one facing "down" to read reflectance from the ground, and one "up" to read irradiance from the sky. Both these provide a Ch0 (broadband) and Ch1 (infrared) signal.

The "up" sensor is intended to compensate for spectral variations in ambient daylight. As you see in the chart below, it also does a lot to normalize the intensity. The "BB..." trace below corresponds to the down-looking broadband channel, except now divided by the broadband channel from the up-looking sensor. The "IR..." trace is adjusted the same way, from the up-looking infrared channel. The "VS..." (derived visible) trace is their difference, after adjusting BB and IR.



Using these adjusted numbers for the NDVI calculations gives an even more stable trace.

