A Guide to PHD Guiding

Introduction

"PHD Guiding" from Stark Labs has to be the most popular guiding software in use today. That's partly because of the price (free), but it also happens to be a darn good guiding tool. But even if it came with excellent documentation (it does not), there would still be a need to explain many of the functions and settings of PHD, especially for the beginner. This "guide to guiding" was created to fill that void. Portions of this text are lifted directly and without shame (but with permission) from notes created by Neil Heacock. It is based on PHD Guiding version 1.13. Some features may be different for other versions.

First Steps

When PHD Guiding is launched it does not automatically connect to your camera or mount. Make sure your guide camera is connected and powered on, then click on the camera icon in the lower left corner to bring up a list of supported cameras and select your guide camera from the list. Depending on the type of camera you might be prompted to select some options. To connect to your mount, first select it from the <MOUNT> menu. The choices are ASCOM, GPUSB, GPINT (at various port addresses) or "on-camera" (meaning a guide port built into the guide camera, if available). ASCOM is used to connect to a mount through a software driver that is compliant with the ASCOM standard. GPUSB and GPINT are hardware devices (accessed via a USB port or parallel port, respectively) that connect on the other side to an ST-4 compatible guide port on the mount. While parallel ports are rarely found on modern computers, the GPUSB device is a very popular way to connect a computer to a telescope mount. The "on-camera" selection relies on the driver for the guide camera to provide a guide port interface. After selecting the mount type, click on the telescope icon to connect. As with the camera, you may have to make additional selections depending on the mount and interface.

The status bar at the bottom of the PHD window shows whether the camera and mount are connected (PHD refers to the mount as "scope"). It will show "no cam", "no scope" and "no cal" (no calibration) initially. This status bar provides other useful information that is often overlooked, so get familiar with it.

The <TOOLS> menu also has some important features, but we'll get to those later. The other menus are not important (unless you don't read this document, in which case the <HELP> menu will be essential).

Once the camera is selected you can click on the "loop" icon (3rd from left) to start displaying images from the guide camera. The first thing you want to do is select the exposure time from the pop-up menu to the right of the "stop" icon. To some degree the exposure may be defined by the brightness of available stars and the sensitivity of the camera. However, exposure duration is also important in that too short an exposure will yield images that are affected by atmospheric disturbances (seeing). Unless seeing conditions are excellent you should use exposures of at least 2 to 3 seconds. This will average out the variations in star position caused by atmospheric disturbances, thus allowing PHD to more accurately calculate the true centroid of the guide star. Once you have selected the exposure time (and camera gain, if available), click the "Take Dark" button and cap the guide scope as directed. This will capture a reference dark frame that will be subtracted from subsequent frames from the guide camera. This is important to improve the accuracy of centroid calculation.

Next to the exposure setting is a slider that is supposed to adjust the appearance of the displayed images. I've never found it to be useful and it doesn't affect guiding, just how the image is displayed to you. However, you should play with it to see what position looks best in your system.

In the lower right corner is the "camera dialog" button. For some cameras you can click on this button to change settings in the camera driver. Webcams typically provide such an interface. If your camera does not support it the button will be grayed out.

The "Brain"

Now we come to the all important and widely misunderstood "brain" icon. Although it is described as "Advanced setup", it is a near certainty that you will need to change something in this menu at some point. And even if you don't, it's very helpful to understand what the options are. Each is described below:

1) RA aggressiveness: A setting of 120 means that you are going to apply 120% of the calculated correction to the RA movement. This is probably too much. Anywhere from 70 to 100 should likely be what you want, sometimes less, but never over 100.

2) RA Hysteresis: Determines the degree to which previous corrections affect the current calculation. The idea here is that a quick change in the measured error in guide star position is probably due to a bad measurement (perhaps due to atmospheric disturbance), so the correction is "diluted" by mixing in some of the recent correction trend. In addition to sensitivity to seeing conditions, this setting is influenced by the characteristics of the mount. A good mount will not have sudden large changes in the RA position, so a large hysteresis setting is appropriate to filter out bad measurements.

3) Max RA Duration (milliseconds): Defines how large of a correction PHD is allowed to make (in milliseconds). As with other settings, the idea here is to limit the damage caused by bad measurements of the guide star position. Allowing a very large correction (in one cycle) might require a large correction in the other direction later. If there really is a large error the exposure is probably ruined anyway, but in any case it isn't going to be much worse by breaking up the correction into multiple cycles, which is what happens when the max duration limit is hit. Unlike parameters specified in pixels, the proper setting for such timing

parameters depends on the guide rate set in the mount. With the typical guide rate of 0.5X (50%), the max RA duration should be set at something like 300 milliseconds, which corresponds to a movement of about 2.5 arc-seconds. Another consideration is that in most mounts (especially using an ST-4 guide port) the guiding process is delayed for the duration of the correction pulse, so allowing a large correction in one axis may delay a necessary correction in the other. Theoretically, guiding through a serial link to the mount (usually via an ASCOM driver) could make the correction more quickly, but this is generally not done. The next guide exposure cannot begin until the mount movement is done.

4) Search Region (pixels): This defines the size of the area (centered on the previous position of the guide star) to search for the star to find its new position. The default of 15 pixels (i.e., a 30 x 30 area) is generally quite adequate. Larger areas take more time to process and if it gets that far off you've got a serious problem that is probably not going to recover anyway.

5) Min Motion (pixels): This is the minimal amount that the star is allowed to move *without* sending a correction. It applies to both RA and DEC. If it were set to .25, the star would be allowed to "float around" a quarter pixel without PHD sending corrections to the mount. This is like a non-linear version of RA aggressiveness: The idea is to eliminate corrections that are probably erroneous, but in this case it is done through a threshold rather than a linear modification of the correction value. A setting of 0.15 is typically good, but it will depend to some degree on the relative magnification of the guide system and imaging system as well as on the seeing conditions and the quality of the mount.

6) Calibration step (milliseconds): This is the length of a pulse in milliseconds that PHD will send to your mount during the calibration process. Unfortunately, no single recommendation for this parameter will suffice because it depends on what part of the sky you are imaging, what the guide scope magnification is and what guide rate is set in the mount. Start with the default setting of 500ms. During calibration, look at the status bar to see how much the selected star has moved from the initial position (the goal in pixels is shown in parentheses - this number is calculated by PHD based on the characteristics of your camera). If it is moving just a few pixels per iteration you could speed things up by using a larger setting for calibration step. On the other hand, if it exceeds the goal in just a couple of iterations the calibration may not be accurate. However, a high degree of accuracy in calibration is not necessary. If it takes at least 4 iterations it's probably close enough.

Note that this setting also defines the duration of pulses used to manually guide the mount (see <TOOLS> menu). After calibration is done you can safely change this setting to the duration you want for manual guiding, but remember to change it back before you calibrate again.

7) Time lapse (milliseconds): A delay after a guide movement before the next guide camera exposure. In some cases the mount movement may not be complete or may not have settled when the "pulse" is done. Setting some delay here will provide time for the mount to settle so that the next exposure is done entirely at the new position.

8) LE port and LE read delay: These settings apply only to guide cameras that use a parallel or serial port on the computer to control the "long exposure" mode of a modified webcam. Beyond that I can't give you any information on how to use them.

9) DEC guide mode (Off/Auto/North/south): Normally you would leave this set at "Auto". If your polar alignment is excellent you might prefer "Off" to avoid spurious corrections. If DEC is drifting you can note which way it's drifting and enable corrections only in the opposite direction. This can be advantageous because it avoids sending DEC corrections that aren't really necessary.

10) DEC Algorithm (Lowpass filter/Resist switch): The default selection is "resist switch", which means that PHD will make corrections normally if they are in the same direction as recent corrections, but will ignore small errors in the other direction until they become large. The lowpass filter option effectively averages the corrections so that most spurious movements are avoided, but small steps to correct for drift are applied. The latter strikes me as the better choice, but it does not seem to work in the current version of PHD. The corrections executed are much too small. If DEC guide mode is set to "North" or "South" a setting here of "Resist switching" has no affect, so lowpass filtering should be the better choice – if only it would work!

11) DEC slope weight: The author of PHD, Craig Stark, doesn't explain this well and I certainly have no insight, but suggest taking him at his word and leave it at "5".

12) Max DEC Duration (milliseconds): Same as for Max RA Duration.

Note: It has been suggested that setting max DEC duration to some very high value such as 2000 milliseconds will correct problems with certain mounts, such as the Celestron CGEM. The problem that is addressed by this setting is that the DEC motor sometimes won't respond to a small correction due to mechanical issues. It is true that allowing a large DEC correction will eventually overcome this fault in the mount, but the exposure will already be ruined. Other than fixing the mount, the only solution to such a defect is to get the polar alignment good enough that DEC guiding can simply be turned off.

13 Star Mass Tolerance: To ensure that it is guiding on the correct star, PHD checks the "star mass" (basically, how bright it is) on each exposure. If it changes by a factor greater than the specified tolerance, a warning will sound. A value of 0.5 is reasonable, but you may sometimes get false warnings due to various conditions that affect the apparent brightness. In this case, you can avoid the warnings by setting a higher tolerance. A setting of 1.0 disables the feature.

14) Noise reduction (none/2x2/3x3): Normally set at "none", but if you're working with a dim guide star and can't increase the exposure time, you might want to try one of the other settings. What it does is to trade off resolution for sensitivity (like binning, but done after the capture rather than in the camera). This can be especially useful for cameras that have very small pixels. Such cameras may provide more resolution than is useful, but suffer from low sensitivity and noise that degrades the accuracy with which PHD can calculate the centroid of

the guide star.

15) Camera gain(%): This setting is available only with some types of cameras and the recommended setting depends on the camera model. In general, you want to use the setting that provides the best signal to noise ratio.

16) Force calibration (check box): When checked, clicking the PHD "target" icon will start a calibration cycle, then start guiding. When unchecked, PHD will use the previous calibration data and begin guiding immediately when the target icon is clicked. PHD automatically “checks” this item when launched and “unchecks” it when a calibration is successfully completed. However, when you move to a different part of the sky you will have to check it to do a new calibration.

17) Log info (check box): If checked, PHD will write all calibration and guiding iteration information into a log file. I always leave this checked so that I can later check how well the guiding performed. Not that I always do such a check, but it's good to know that you can investigate later if necessary. PHD creates a new file for each date in which logging is started. This is much better than simply overwriting the old file (which older versions did), but can still be confusing if you start logging after midnight on one night and then start it again before midnight the next night.

18) Disable guide output (check box): When checked, PHD does everything the same, except that no commands are issued to the mount. This is useful because it allows you to study the behavior of the mount without guiding. In particular, by disabling the output and enabling logging you create a log file that can be imported into a program to analyze the mount's periodic error and generate a periodic error correction (PEC) file.

19) RA-only dither (check box): Dithering is a feature that requires using both PHD Guiding and Nebulosity. For details on dithering please refer to the Nebuosity documentation. Normally, dithering is applied to both RA and DEC, but checking this box forces it to apply dithering only to RA. This is a good idea if you have problems guiding in DEC because it might take too long for DEC to settle after a dither move.

20) Use subframes (check box): Some cameras are capable of reading pixels from only a portion of the captured image. If this selection is checked and your camera supports the feature, PHD will use it to decrease the amount of time it takes to download an image by reading only the area around the guide star (the search region).

Calibration

Once the advanced settings are made we are ready to begin calibration. As noted before, calibration is the process of measuring how the stars move (apparent motion) in response to guide commands. PHD measures both the speed and direction of movement. Note that it is theoretically possible to guide with the camera at any given angle, but you will find it least confusing to orient the camera such that RA and DEC movements correspond to horizontal and vertical movement on the screen. It's not too important which axis is horizontal or vertical, but you should keep it the same so you are familiar with how it moves. If the guide camera has non-square pixels it's a good idea to turn it such that RA movement corresponds to the smaller dimension because RA guiding is more important (indeed, with proper polar alignment, DEC guiding should be unnecessary). For example, with my Meade DSI Pro camera the pixels are 9.6 x 7.5 microns (horizontal x vertical), so I have it turned 90 degrees from what would appear to be normal to get RA on the more accurate vertical axis.

Calibration must be done on a star that is reasonably close to your imaging target. Within a few degrees should be quite adequate (except near the pole, where guiding gets quite tricky, but is also relatively unimportant). Note that when doing a meridian flip the directions for guiding change. PHD provides a selection under the <TOOLS> menu to flip the calibration data.

After slewing to the target area, click the loop icon to see what stars are available. You could do the framing of your target first and then calibrate at the final position, but I like to get calibration out of the way first. While PHD is looping (or after you stop looping) you can click on any star and see the metrics for it in the status bar. Primarily, you want the highest SN (signal to noise ratio). It's fairly obvious that you shouldn't use a dim star for guiding or calibration, but it's less obvious that you also don't want one that's too bright. Specifically, if any pixel in the star image saturates, the calculation of the centroid will be degraded. PHD doesn't give a very clear indication of how bright a star is or whether it is saturated. Fortunately, the SN ratio will be poor for a saturated star, so you can just go by that and pick the star with the best reading. Be aware, however, that the SN ratio may change over time. PHD has a function to automatically select a star, but it doesn't work well with cameras I've seen, so I don't recommend it. You should also try to avoid stars that have other stars of similar brightness nearby because PHD could select the wrong one. Decreasing the size of the search region might help in such a case.

Once a star has been selected click the "target" icon to start calibration. PHD will attempt to move the mount in RA and you should see the star move. You can see the numeric measurement of the movement (in pixels) in the status bar. If there is no significant movement after a few iterations, click "stop" and increase the calibration step size before trying again. PHD will measure the movement in one direction, then reverse direction to get back to the starting point. If DEC guiding is enabled it will repeat the process for the DEC axis. The amount of time needed to move a given amount in either axis varies with the RA and DEC position. In particular, when imaging near the pole RA will appear to move very slowly. Unfortunately, you can't specify different step sizes for RA and DEC, so patience is your best option. PHD usually does not go all the way back to the starting point before it stops calibrating and starts guiding, so there will be a slight shift in your framing. I prefer to calibrate before doing the final framing and camera rotation, then selecting a new star to guide on (without doing the calibration over again).

When calibration is done PHD will automatically begin guiding on the selected star. At this point you may want to stop it as you adjust the scope to get the target framed and focused because PHD will likely lose the guide star during this process and complain about it. Once the target is framed and focused, go back to PHD and repeat the process to select a guide star. When you did the calibration PHD automatically unchecked the "force calibration" option, so you can just click a star and then click the "target" icon to start guiding.

You should wait a bit before starting to image. If DEC Guide Mode is set to "auto" and “resist switch” PHD will watch the DEC drift and then favor corrections for the dominant drift direction, although it will still make corrections in the other direction if the error gets large enough. To use the "north" or "south" settings YOU are the computer and need to watch which way it is drifting and select north or south accordingly.

Tool Time

Now that basic guiding is in operation you can start the image capture process. But this is also a good time to look into the <TOOLS> menu:

1) Manual Guide: Brings up a window with 4 buttons labeled North, South, East and West. Each time you click on one of these buttons PHD will move the mount in the selected direction for a period defined by the "calibration step" parameter. Obviously, you wouldn't normally do this while auto-guiding. It is sometimes useful to make a small correction in position, either to better frame the image you are capturing or to get the guide star away from the edge of the frame, but it's generally too slow to make very much of a move. It is also a convenient way to make sure that PHD is actually talking to the mount and to determine which way stars move on the screen in response to the guide commands. Finally, if for some reason you can't or don't want to let PHD auto-guide, you can watch the star position and guide it manually using this tool. I don't recommend it.

2) Erase Dark Frame: I can't think of any reason you would want to erase the dark frame (as opposed to replacing it with a new one), but here's the option, just in case.

3) Auto-select star: This feature was discussed previously. I don't recommend using it.

4) Enter calibration data: This could have been a useful feature if it supported saving and loading calibration data to/from a file. Instead, it simply prompts you to manually enter the numbers. The numbers can be found in the log file for an earlier calibration. This would be useful if you were returning to the same target a second night, but finding and entering the data is so awkward that you might as well just re-calibrate. Another potential use of this feature is to increase or decrease the “gain” for each axis. If calibration is already done PHD will show the measured rate and angle. To increase the DEC gain, for example, reduce the DEC rate number. PHD will then use a larger correction for a given amount of positional error. This works the same as RA aggressiveness, but can be done to either axis.

5) Flip calibration data: Use this button when you do a meridian flip (and continue imaging the same target). PHD will invert the previous calibration so that it is correct for after the flip.

6) Overlay selection (No overlay/Bullseye/Fine Grid/Coarse Grid): When one of the overlay patterns is selected PHD will display the pattern on top of the image. The bullseye is useful if you want to center a star for some reason. I've never used the others.

7) Enable logging: Same as the "Log info" check box in the advanced setup.

8) Enable star image logging: Also logs a small image of the guide star for each cycle. This would be useful for a really detailed analysis of guiding issues. Way beyond what almost anyone would want to do.

9) Enable server: This opens a link between PHD and Nebulosity (another Stark Labs product) to support "dithering". This is a useful feature, but beyond the scope of this document.

10) Enable Debug logging: As with "star image logging", this is only for very serious analysis of guiding issues and is mostly used by the author.

11) Enable Graph: This is what you want! As soon as you start guiding, enable this graph and move the window it brings up such that you can see the guide star, the PHD status bar and the graph. Some of the advanced setup parameters can be adjusted in real time within the graph window. This is very useful. This is also where you will find the display of RMS error, the best measure of how your guiding is performing. RMS is calculated over the same period as the graph, which can be adjusted to 50, 100, 250 or 500 iterations (personally, I wish that the X axis of the graph were true time rather than just iterations, but so be it). Remember that the RMS error value shown is only for the RA axis.

12) Enable Star profile: This selection brings up a window that shows you a detailed graph of the most recent acquisition of the guide star image. This can help you determine whether you have a good guide star.

13) Enable manual lock position: This option allows you to place the guide star at a specific location in X/Y coordinates of the guide camera. This is rarely useful, but if you are particularly anal, you could specify the center of the image to make everything very neat and tidy. Or you could specify something like (100, 100) so that the guide errors are easy to evaluate arithmetically. PHD will prompt you for the coordinates you want and when you start guiding it will attempt to move the selected star to the specified coordinates.

The Graph

Enabling (and watching) the graph is highly recommended. The buttons to the left of the graph are described below, starting with the top button:

1) Horizontal scale: Click to toggle between 50, 100, 250 or 500 points. Note that these are guiding cycles, which have a somewhat variable time component, so there is no specific time scale to the plot.

2) Vertical data type: Click to toggle between RA/DEC and dx/dy modes. The former is the default and shows the errors in terms of the RA and DEC axes of the mount. "dx/dy" mode instead shows the errors in terms of the image captured from the camera. How the two

modes correspond depends on the rotation of your guide camera.

You can change the colors of the plot lines by doing a "shift click" (for RA/dx) or "control click" (for DEC/dy) on this button.

3) Hide: Click to hide the graph. Data will still be collected and when you enable the graph again you can still see up to 500 "old" points.

4) Clear: Click to reset the graph and clear the oscillation and RMS error calculations. This is useful when you have made a change in the setup and want to see the results without influence from previous settings.

Below the graph are a few selectors for some of the "brain" settings. They work the same as under the "brain" menu, but can be changed here on the fly. The other advanced settings can only be changed in the "brain" menu, which requires you to stop guiding. The settable parameters are RA aggressiveness, RA hysteresis, Minimum motion, Max RA duration, Max DEC duration and DEC guide mode.

General Guiding Concepts

That covers just about everything in PHD Guiding, but there are some general comments on guiding that should be understood:

Differential Flexure

First and most importantly, guiding uses one camera to control the pointing angle of another. It seems obvious that for this to work there must be a fixed, solid connection between the two imaging systems. What is not so obvious is that a very small amount of difference between them (well below unaided human perception) can destroy the usefulness of guiding. This problem can be greatly reduced by using either an off-axis guider system or a dual-chip camera (essentially, two cameras in one), but such systems have other issues, so they are not always the right choice.

Polar Alignment

Also very important is polar alignment and the idea that you should strive to eliminate the need for guiding. In particular, DEC guiding should, for the most part, be totally unnecessary. If polar alignment is absolutely perfect, virtually all DEC guiding corrections come from factors that are short-term aberrations that should not be corrected for because they will go away by themselves. The only real exception is that atmospheric refraction changes slightly with the thickness of the atmosphere and therefore changes the apparent position of the stars (albeit very slowly) as you track the sky. When polar alignment is not perfect you get "DEC drift" - apparent motion of stars in declination. The direction and magnitude of this drift depends on both the direction and magnitude of the alignment error and what part of the sky you're looking at. However, for a given mis-alignment and imaging target, DEC drift will always be in the same direction (north or south). If you disable guiding outputs and watch the PHD window you can see which way it drifts (eventually you'll see the star move away from the crosshair, but it's quicker to just watch the status bar and see which direction PHD is predominantly trying to move the mount, paying attention only to north/south corrections). If polar alignment is good the average north/south correction will be near zero. If there are a lot of corrections in both directions it is probably caused by atmospheric disturbance ("seeing"). You should not try to correct for seeing because the system cannot respond quickly enough to do any good and the average will be near zero anyway. But if there is a "trend" toward either north or south you can enable correction for just the opposite direction. Other factors might still cause errors in the other direction, but the constant drift will soon compensate for it. The amount and direction of the DEC drift can also be used to correct your polar alignment. This is called "drift alignment", but I won't go into the details here.

DEC Guiding

As indicated above, you should not try to correct for errors caused by seeing. Another reason to minimize corrections in DEC is that, unlike the RA axis, if corrections are allowed in both directions of DEC the motor has to change direction and you are then dealing with backlash, which can really complicate things. If you allow DEC corrections in just one direction the gears will always be engaged in the same way (unless some other force, such as wind, pushes it the other way), producing much more consistent results. This is not an issue (normally) with RA guiding because the motor is always moving west - guiding simply changes the speed. In most mounts you can set the "guide rate". There are several different ways of specifying guide rate, but they all relate to the normal tracking rate (sidereal rate). In almost all cases a "west" correction cannot exceed two times sidereal rate and an "east" correction at most stops all RA motion. That is, it never reverses direction in RA for guiding. While there are some mounts that are capable of reversing direction for RA guiding, you should definitely avoid such settings. Determining the best guide rate is discussed in more detail below.

RA Balance

Because the RA motor is always going forward, we can select a "bias" in the balance to improve tracking consistency. That is, by shifting the counterweight slightly, gravity keeps the gears engaged on the same sides of the teeth at all times. It is preferable to bias the balance toward the east so that the motor is always pushing up on the weight rather than letting it down. However, this means that when you move from one side of the meridian to the other you have to change the balance. You don't want to loosen the clutches to check the balance in the middle of an imaging session, so it's a good idea to experiment with balance beforehand and mark positions on the counterweight shaft for the two possible positions of the weights.

Guide System Magnification

The relationship between the magnifications of the guide system and imaging system is somewhat flexible, but cannot be totally ignored. Magnification is also referred to as "image scale" and is expressed as arc-seconds per pixel. The guide camera can have a higher magnification (lower arc-seconds per pixel) than the imaging system, but if the guiding system magnification is very high you may have trouble finding a good guide star, particularly since high magnification usually means a slower focal ratio. On the other hand, a guiding system can have lower magnification than the imager because PHD (and other guiding software) can find the center of the guide star to much finer resolution than an individual pixel. But if you take this too far guiding performance will suffer. I would recommend that the image scale of the guide system be no more than two times that of the imager (half the magnification).

Guide Star Quality

There is also some flexibility in the quality of the guide star image. For example, an inexpensive achromatic refractor will usually work very well, although you might get very slightly better performance from an ED doublet or true apochromatic scope. It has been said (I know because I've said it myself) that guiding can actually be better with a slightly out of focus image. This is theoretically true because if a star is so tightly focused that all its energy lands on just one pixel the software cannot determine the star's centroid to anything smaller than that one pixel. Normally, the energy is spread over many pixels and the brightness at each location can be used to project where the center is to a much higher precision. However, in the real world it almost never happens that a star image is that well rendered on the image sensor, especially if the scope is not of top quality. In fact, I recommend that most people should focus their guide scope as well as possible because this will yield the greatest brightness and signal-to-noise ratio. Since guiding is done with relatively short exposures, sensitivity of the system is often the most critical feature and a fuzzy image greatly reduces the ability of the camera to cleanly detect the guide star.

Pixels and Durations

PHD does not know anything about your magnification/image scale (some guiding programs do). It also doesn't know your guide rate or what part of the sky you are imaging. It doesn't really need to know these things to do its job, but it is helpful for you to understand the relationships between these parameters.

First, calculate your image scale with this formula:

image_scale = arcsin(pixel_size / focal_length) * 3600

“pixel_size” and “focal_length” must be in the same units. Since pixel sizes are typically stated in microns and focal lengths in millimeters, divide the pixel size by 1000 to convert it to millimeters. The arcsin function returns the image scale in degrees per pixel. Multiplying by 3600 (60 minutes / degree * 60 seconds / minute) gives you arc-seconds per pixel. You can use this formula for both the guide system and the imaging system. Remember that pixels may have different dimensions for height and width, in which case the image scale is different for the two axes. Use the effective focal length of the scope after accounting for any focal reducer or extender.

Sidereal rate is 15 degrees per hour, which (very conveniently) converts to 15 arc-minutes per minute or 15 arc-seconds per second. So without guiding, an equatorial mount is moving west (in RA) at approximately 15 arc-seconds per second and does not move at all in DEC. This rate is the same as the apparent motion of the stars, so we can think of it as "not moving". Guide rates are specified as a fraction of sidereal rate, but there are several ways of expressing it: If the available guide rates are all greater than "1" (usually up to "2"), the number is the multiplier applied to sidereal rate for a WEST correction. If the available rates are "1" or less, the number represents the fraction of sidereal rate that is added to (WEST) or subtracted from (EAST) the normal motion rate. That is, the number represents the apparent rate of motion – what we would see looking through an eyepiece. If the guide rate is specified as a percentage this is the same as the second example, except that the fraction is expressed as a percentage. So a typical guide rate might be given as "1.5X", "0.5X" or "50%". These all mean the same thing. (Personally, I think the first method, used by Orion, is confusing and should be avoided. The second method is much better, but is too easily confused with the first. So the third method is the easiest and most clear. Naturally, it is probably the least commonly used.)

A guiding system controls the motion of the mount by controlling the duration of a pulse. When an ST-4 guide port is used the computer does the timing directly. For most ASCOM interfaces the computer calculates the duration and sends this information to the mount, but the actual timing is handled by the processor in the mount. In either case, there is a desirable range for the pulse duration: If it is too long you could be missing changes in the tracking while it is executing the move. If it is too short the accuracy of movements suffers because the processor cannot measure the time period that precisely (especially if it is doing other things at the same time). So the guide rate should be selected such that typical guiding commands are in the range of 10 to 300 milliseconds. The image scale of the guide system will affect this choice. In most cases the image scale will be 2 to 5 arc-seconds per pixel and a guide rate of 50% is recommended. If the image scale is greater than 5 (such as when using a very short guide scope) a higher guide rate is probably appropriate – as much as 100%. If the image scale is less than 2 you might want to go down to 25%. This would often be the case when using an off-axis guider through a long 'scope such as an SCT. In any of these cases, the characteristics of the mount might also influence the choice of guide rate. A very well behaved mount should use a slower rate to get the best precision in pulse timing, but one that moves erratically might need a faster guide rate to keep the pulses reasonably short.

Note that the term "sidereal rate" is really only appropriate for RA motion, but mounts use the equivalent motor speed to define DEC guide rates.

When setting the calibration step size you can use image scale to determine how many pixels it will move per step. For example, if the image scale is 2.5 arc-seconds per pixel and you set the step size to 1000 milliseconds (1 second) with a guide rate of 50% you have:

((15 arc-seconds / second) * 1 second * 50%) / (2.5 arc-seconds / pixel) = 3 pixels

This may be a reasonable value, but watch the status bar during calibration for the rightmost number (in parentheses). It is the “target” for the total movement, as determined by PHD. You want to reach that goal in 5 to 10 steps, so divide the number by 5 to 10 to determine your goal for pixels per step.

However, the above is really only true when DEC is near zero. As DEC approaches 90 (or -90) degrees the amount of movement in RA decreases proportionately. For example, at 45 degrees DEC you would want to double the step size. You might notice that this suggests that apparent motion becomes zero at the pole. That is correct and it means that you really can't guide at the pole. All of the motion is rotation around the pole. When imaging even 30 degrees away from the pole (DEC > 60) you will notice that you have to significantly increase the calibration step size and within 10 degrees (DEC > 80) it just gets ridiculous.

Sensitivity and Resolution

There is a very real trade-off between sensitivity and resolution. For a given size of sensor chip, to get higher resolution you make the pixels smaller, but that decreases the sensitivity. So a guide camera should be selected based on the needs of both resolution (image scale) and sensitivity. A very short guide scope will need higher resolution to get a suitable image scale, but that will then require either longer exposures (thus requiring a better quality mount) or a fast focal ratio in the guide scope. Most cameras used for guiding do not support binning. That's unfortunate because a choice of binning ratios would allow you to adjust the trade-off between sensitivity and resolution to match the equipment and situation.

Because sensitivity and low noise are important to guiding, a color camera should pretty much never be used. The color filters substantially reduce the amount of light that reaches the sensor. Furthermore, differences in brightness between red, green and blue filters will cause errors in calculating the location of the guide star's centroid.

Understanding RMS

“RMS” stands for “root mean square” and it is a common way of specifying the average (mean) magnitude of a varying signal, especially one that goes both above and below the zero point, because it accounts for the magnitude without regard for the direction (positive or negative). For example, AC voltages are usually expressed as an RMS value. For a sine wave the RMS value is pretty easy to understand – it's basically the average over a full cycle where the negative portion is inverted (RMS values are, therefore, always positive). But for a more complex signal, such as the error in the position of a guide star, short term errors have relatively little effect on the RMS. This is apparent when you look at the PHD guiding graph and see spots where the error is one or two pixels, but the RMS value shown is much less than one pixel. Indeed, if there is any “bias” in the error (if RA tends to stay a bit above or below the line), the RMS will probably be dominated by this bias. And if the bias stays the same through a whole sub-exposure the resulting image quality may be better than the RMS value would suggest (because what really matters is that the position stays constant, not that it is exactly where it's supposed to be), but that is not usually what happens.

RMS tells you the range over which most of the energy from a star will be distributed. That is not to say that an RMS of 0.2 pixels means that the star image (or even most of the star image) will be just 0.2 pixels in diameter. The nominal star size will be determined by the optical system and then the error in guiding will "enlarge" the star image. Unless something is wrong with the system, most of the error will be in the RA axis, so the star will be elongated by an amount proportional to the RMS value (remember that PHD shows RMS only for the RA axis).

Although RMS is a great metric of guiding performance, it doesn't tell the whole story in terms of image quality. For example, say a star is bright enough that it is twice the saturation level of the sensor. Even though most of the energy is confined to the area defined by the optics and RMS error, pixels at half the peak level will record just as brightly as the peak and so the error will appear larger. A similar effect happens when we stretch an image to bring out faint details. Stars that were OK before the stretch become saturated and more elongated.

I don't know of any easy way to quantify the image quality that will result from a given guide graph. The best that we can do is to make adjustments to get the RMS as low as possible and watch out for larger, short term errors.

Connections

In general, there are two ways to connect the computer to the mount: An ST-4 compatible guide port or a proprietary interface (typically an RS-232 serial interface). The latter type is proprietary to each mount manufacturer because there is no standard interface definition. There is, however, a standard software abstraction of the interface, the ASCOM standard. Not so many years ago, computers had parallel printer ports that could be easily set up to drive an ST-4 interface. Now that parallel ports are pretty much nonexistent, a device such as the GPUSB from Shoestring Astronomy is typically used instead. It plugs into a USB port on the computer on one side and the mount's ST-4 guide port on the other. (Note that mount manufacturers usually don't use the ST-4 designation, as that name comes from an SBIG product, a stand-alone guide system.) Most mounts also have some way to connect directly to a computer - the proprietary interface. It is usually an RS-232 serial interface and, again, many modern computers do not have such ports on them and a USB-to-serial adapter is required. These are simple and inexpensive devices that generally work pretty well, but often have the annoying behavior that the virtual COM port assigned to the device changes each time you plug it in to a different USB port. It's a good idea to always use the same USB port for these adapters so you don't have to change settings in your software. Once the physical connection is made you must select the interface (i.e., select which driver to use) in PHD and other application software.

Dealing with Seeing

Atmospheric disturbances cause the apparent position of the stars to vary slightly, but quickly, and this can be problematic for guiding (the degree of such disturbance is called "seeing"). With an adaptive optics (AO) system it is actually possible to "guide out" these disturbances and produce very sharp images. However, this requires a lot of speed in all the components:

The guide camera must produce a clean guide star image, the computer must process this image and the AO system must respond to guiding commands all within a small fraction of a second. Perhaps the hardest part of this is getting a good guide star image with such a short exposure. AO units are generally fast enough to handle moderate seeing conditions, although this capability does not come cheaply. The other way to deal with seeing is to average the guide star image over a long enough period that the effects of seeing are mostly eliminated. This usually requires a period of at least 2 or 3 seconds. The easiest way to do it, of course, is to simply use an exposure time of at least 2 or 3 seconds and let the camera integrate the star image. Many astro-photographers try to set the exposure time as low as possible so that corrections can be made to the mount position before they become too large. However, unless seeing conditions are excellent the "corrections" sent to the mount will be erroneous and guiding will be worse rather than better.

Conclusion

In fact, it is wise to adopt the attitude that measuring the position of a guide star is always imprecise except when you average many measurements over a significant period of time. Therefore, the positional errors detected in a single cycle should be treated not as "facts" but as "suggestions" of what might be a good way to move the mount. Parameters such as RA aggressiveness, RA hysteresis and min/max durations are the tools to implement this attitude.

When you combine the above attitude with the importance of polar alignment and minimal flexure, the approach to optimal guiding can be summed up as “The best guided mount is the least guided mount”. That is, you should strive to minimize the frequency and magnitude of guiding corrections.