Analyzing PHD2 Guiding Results – A Basic Tutorial

Introduction

At some point, most imagers want to analyze the guiding performance they're getting. They may be trying to identify and fix particular problems or perhaps just trying to get the best performance they can from their set-up. Either way, the PHD2 guide logs are the best source of data, and using them is far better than trying to look at the guiding graph in real-time. The latter is useful as a quick check on how things are going, but any serious problem diagnosis or tuning requires a longer time window than what is available in the real-time graph. In this document, I'll be using Andy Galasso's PHDLogViewer application as a way to show common problems and guiding behaviors. This is the tool we most often use for helping PHD2 users, and it is likely to serve you well.

What I'll discuss in this document is simply the result of my own personal experience and research. I don't claim to be an expert on telescope mount mechanics, and I still encounter guiding behavior that I can't explain. But I've been using some version of PHD since 2006, and I've analyzed hundreds of guide logs in the ensuing years, both my own and those from other PHD2 users. Hopefully, what I've learned can help you understand your guiding results or can at least accelerate your learning curve a bit.

Basics – A Big Picture View

Summary Statistics

When looking at these graphs and statistics, you want to look at the measurements in units of arc-secs (a-s). If this isn't obvious to you, take a look at "Image Scale and Measurements in Arc-secs" in the Appendix. To get a big picture view of your guiding, start by looking at the RMS numbers that appear to the lower right of the graph:

When PHDLogViewer loads a guiding session, the "RMS" value is the standard deviation of all the star movements in that session In this example, about 68% of the movements were below 0.59 a-s, the remainder were larger. Just as a rough starting point, you probably want to see your total RMS value be 1 a-s or lower. How much lower will depend on many things, especially your seeing conditions and the quality of your mount. Now look at the individual RMS values for RA and Dec to see how they compare. With most mounts I've seen, it's common for the RA value to be somewhat higher than the Dec value, probably because the RA gear system is always running while the Dec gear system is not. There are also some seeing-related things that can tend to make the RA displacements a bit larger and more frequent. However, if the RA and Dec RMS values differ by a large amount, for example 2-3X, you are likely to see elongated stars in your images. So what happens if these two values are nearly equal but are also quite large? That's likely to lead to bloated stars in your images, and your resolution and sharpness will suffer. It's not enough for the stars to look round – you want them to be both small and round, where "small" is determined by your optics rather than by guiding errors. You'll sometimes see forum posts from people who say "I get perfectly round stars even with 30 minute exposures." Well, that's probably a good thing – but how large are those star images, and how do their sizes compare to what they get with short exposures of 10-20 seconds?

For many of us, guiding rarely runs all night without some sort of glitches. When looking at the overall statistics, you often need to filter out these events. You'll probably need to address them at some point, but it's good to get a feel for the overall guiding performance without having the numbers polluted by these unusual events. Consider this section of a guide log:
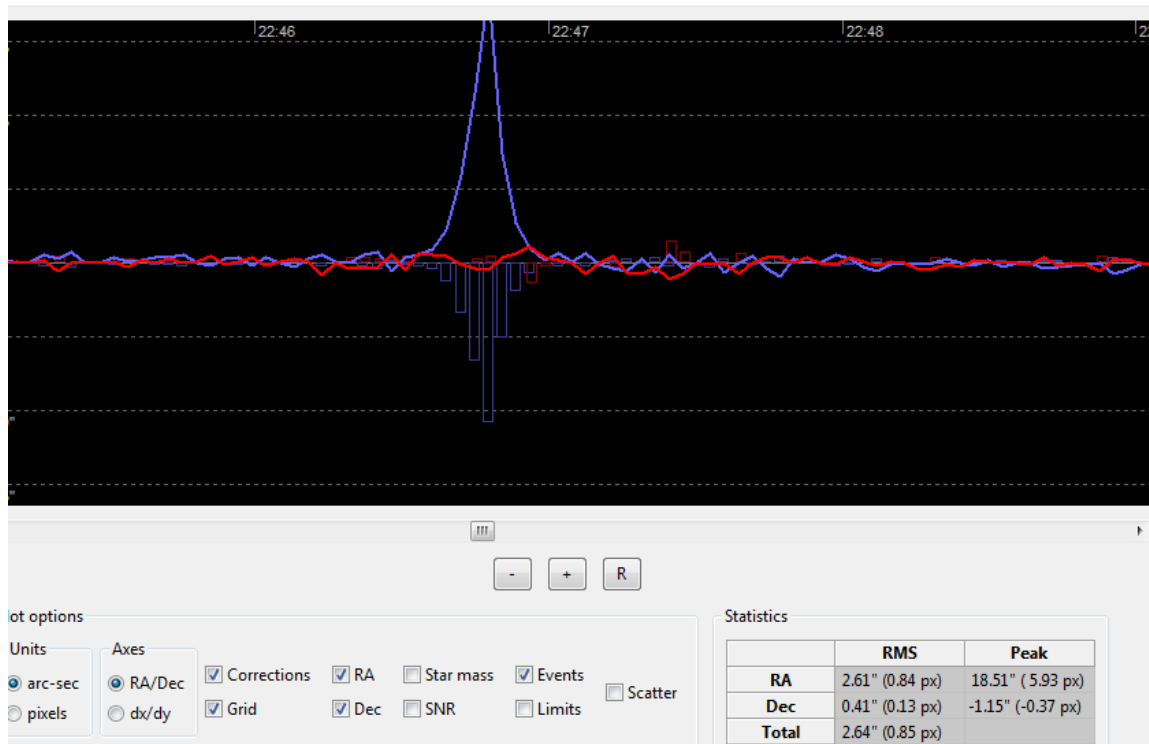

Something bad happened shortly before 22:47, a huge excursion in RA that caused the total RMS to reach 2.6 a-s for the time period shown. But the guiding before and after that event looks much better. You can use PHDLogViewer to isolate those regions to see how things were going. When I did that on this data, the statistics looked much better: a total RMS of 0.5 arc-sec before and after the unusual event. This tells you the overall guiding was going pretty well, but you'll need to figure out the source of this large RA excursion. Doing that will be discussed in the section on "Gremlins."

Seeing

When you look at a guiding graph, most of the rapid star motion you see is caused by seeing, something you can't do anything about. If this isn't a familiar subject for you, take a look at the section on "Astronomical Seeing and Guiding" in the Appendix. If the rapid movements are large, you should check the exposure time being used – longer exposure times can help to reduce the envelope of seeing deflections you see (see Appendix). These rapid excursions are reflected in the RMS statistics and may even be the primary contributor to them. It may help to run the Guiding Assistant periodically to get a sense of the typical seeing conditions at your site. Poor seeing can't be corrected by guiding and you'll simply have to do the best you can. Over time, you'll probably learn what to expect on most nights and will quickly recognize the effects of seeing when you look at the guiding graphs. However, if you're "chasing the seeing" because of poor choices in guiding parameters, you'll see that too many of these seeing-induced star movements trigger guide commands, and this will often result in a saw-tooth pattern in the graph. Identifying those sorts of problems will be covered later.

What's on the Graphs

When you look at the details shown in the guiding graph, you'll normally start by looking at two things: how far the guide star moved from one exposure to the next, and how PHD2 reacted to that move. Here's a typical part of a guiding graph zoomed way in to show the details:
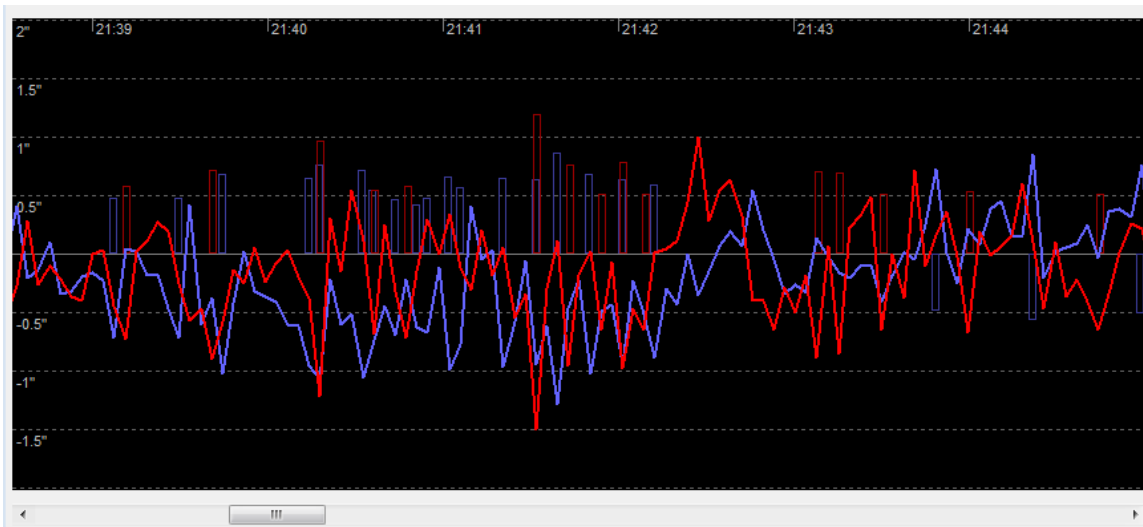
The star movement is shown by the connected lines, blue for RA and red for Dec. The rectangles show the guide commands that were generated by PHD2, and their heights indicate the relative sizes of the guide pulses. Don't over-analyze the up/down convention for displaying the rectangles. This orientation was chosen to reduce clutter and because it is more intuitive – you want the command to "push" the star in the direction opposite to its apparent movement. If you look closely, the rectangles always trail the star motion by a small amount because PHD2 is reacting to the move seen in the previous exposure. What you see here is pretty typical – sometimes it only takes one guide pulse to restore order while in other cases it may take multiple guide pulses in the same direction. You'll see that some star displacements don't trigger any guide pulses at all. That's usually because of the "min-move" setting for that axis but it can also be due to intentional damping in the guide algorithms. You also need to pay attention to the scale at the far left to keep things in context. Beginners often look at these graphs and think the guiding is horrible because there seems to be so much motion. But the motion you see here is generally falling within a range of +/- 1 arc-sec, and the total RMS error for this section is about

3

0.6 arc-sec. In this example, the guiding was about as good as the seeing conditions would allow.


Easy Problems

To get warmed up, we can start by looking at some easy problems to diagnose. One of the easiest is something that looks like this:

Wow! Absolutely perfect tracking, the guide star never moved! Um…… no. This is what happens when you're guiding on a hot pixel, which is the same as not guiding at all. This should almost never happen, but if it does you can try the following:
1. Let PHD2 auto-select the guide star (Alt-s). It can be hard to visually distinguish a hot pixel from a faint guide star when you're just peering at the display.
2. Be sure you're using either a dark library or a bad-pixel map.
3. Apply a 2x2 or even 3x3 noise reduction filter (brain dialog/camera tab).

Now let's look at another problem, one that's more likely to happen if you're using a guide cable attached to your camera and the mount's ST-4 guiding interface:
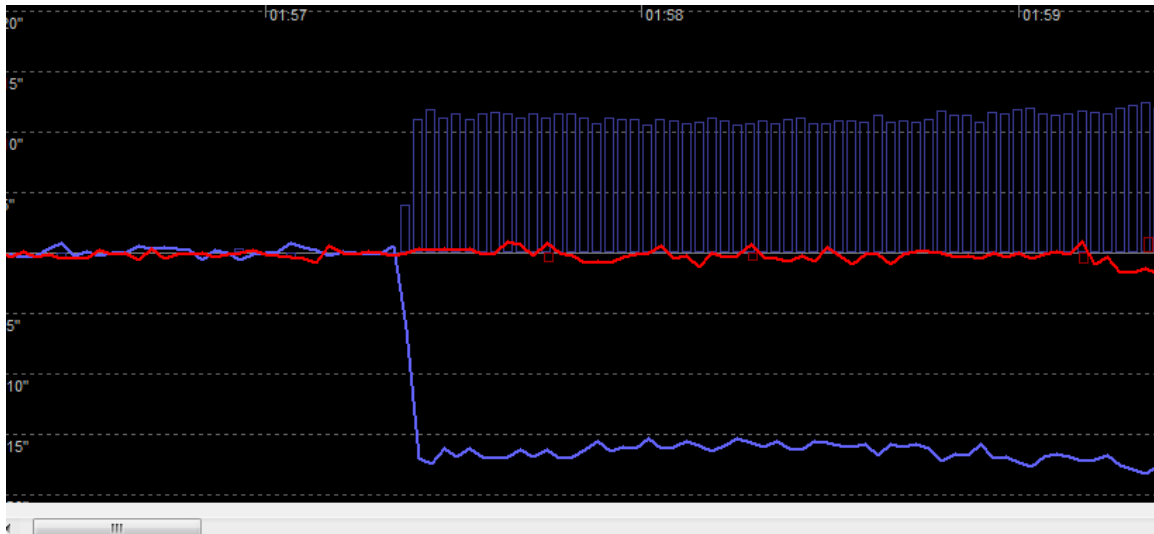
It looks like guiding got off to a reasonable start until the star took a big lurch to the west (down). But in this case, PHD2 was never able to recover. Look at the continuous stream of 'east' guide commands that were sent for the next two minutes – nothing happened! The guide star was never returned to its starting position and in fact continued to drift further west as time went on. The diagnosis? The guide cable was bad and the east guide commands were never being received by the mount. Replacing the guide cable solved the problem. This wasn't the only possible explanation, but it was the obvious thing to try because the ST-4 guide cable is a common point of failure.

Here's another easy one, albeit one you might never see:
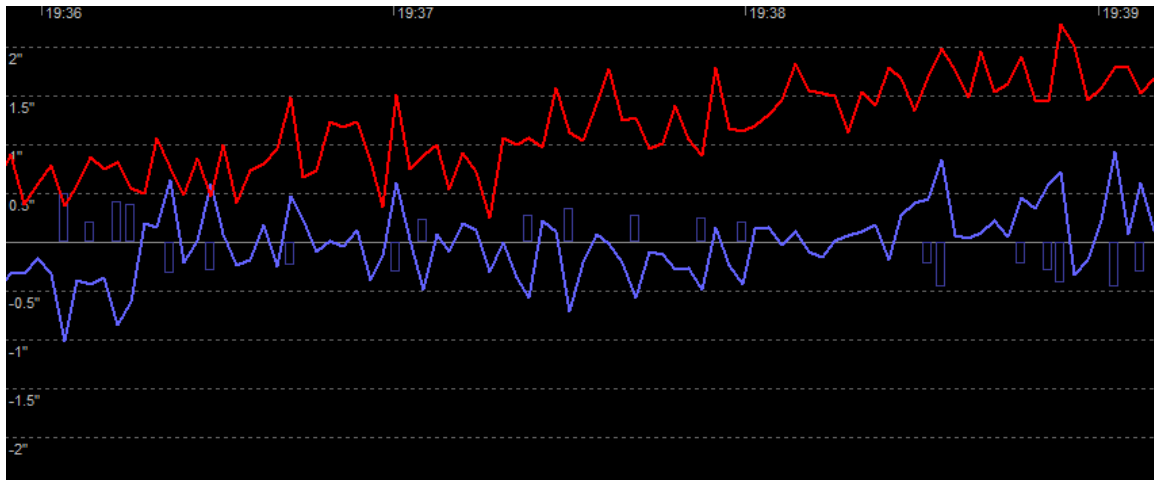
Again, the session started out normally and the guiding results were looking pretty good for RA (blue). But the guide star was drifting to the north (Dec, red) and wasn't being restored – so is this another bad cable? No. Notice there are no red rectangles shown, which means PHD2 was never generating any guide commands to restore the star position. Why did this happen? Because the user had chosen a Dec guide mode of "north-only", which told PHD2 not to generate any south guide commands. Changing the guide mode to 'auto' or 'south' solved the problem. In fairness, this wasn't a dumb mistake because the user might have had a good reason to guide in only one Dec direction and didn't know yet which direction that needed to be. Even so, it's an instructive example, and we've seen numerous situations where users have blundered around in the brain dialog and inadvertently set the Dec Guide Mode to a weird value.

We can finish the warm-up session with another example, one you probably will see at some point:

|	RMS	Peak
RA	1.61" (1.02 px)	-6.23" (-3.97 px)
Dec	0.96" (0.61 px)	-3.87" (-2.46 px)
Total	1.87" (1.19 px)	
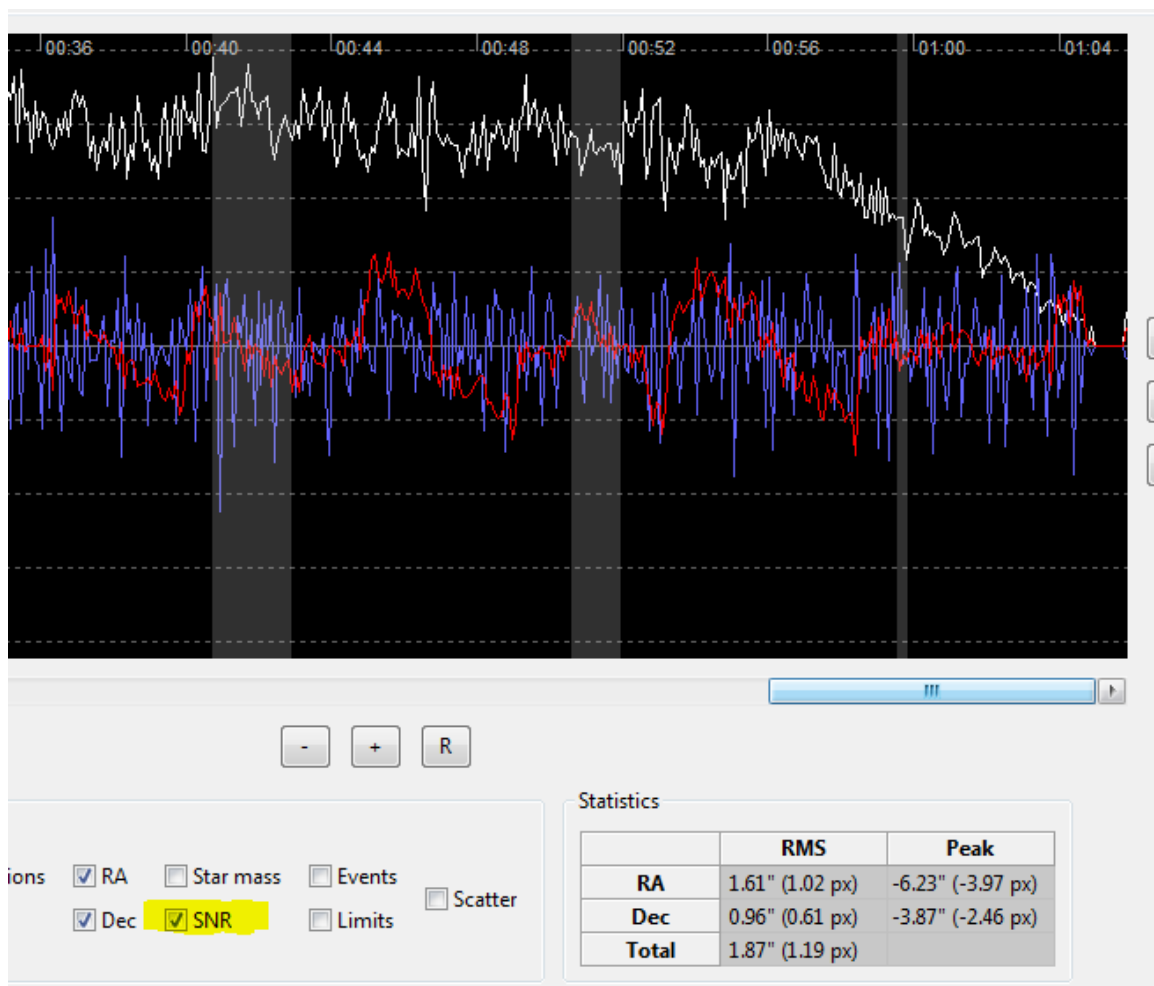
In this case, we've checked the SNR box near the bottom of the window to see how the star's signal-to-noise ratio was behaving. That's shown by the white line at the top of the graph. In this case, the SNR started dropping at around 0:56 and got steadily worse – probably because of clouds. This will often lead to deteriorating guiding results as well, so it's always good to check this if you're starting to investigate a problem.

The point of these easy examples is just to help you get used to looking at the graphs and asking the basic questions – how was the guide star moving on the sensor, what guide commands were generated as a result, and how did the mount respond to those commands? Of course, you may not see such easy problems, but the questions and basic methodology can be applied to more complicated situations.

Sudden Large Excursions – "Gremlins"

At some time or another, most imagers will be confronted with sudden large movements of the guide star – the earlier graph shows one such example. The first thing that beginners often ask is a chicken-and-egg question – is this something that just happened or was it caused by some sort

of rogue guide command from PHD2? In every case I've seen, the answer is the former, and you'll need to track down the source of this "gremlin." Answering the chicken-and-egg question is easy though, you don't need to take my assurances about it. Just enlarge the PHD2 guiding graph and look in the immediate region of the big move – did it start with a ridiculously large guide command or are there simply a bunch of guide commands that follow it as PHD2 tries to wrestle the mount back into position?   Let's look again at the earlier example, but this time we'll zoom in and look closely at what happened when the RA guiding "went crazy":
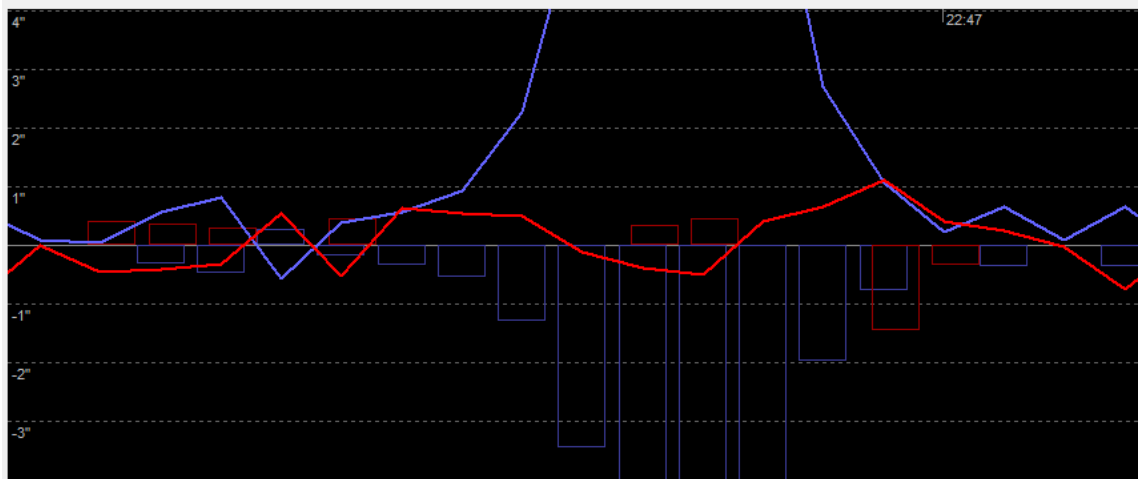
We see that guiding was going normally on the left-hand side of the graph before the big upward spike. In particular, there was no huge guide pulse that caused the star to move so far on the sensor. Instead, you see PHD2 reacting *after* the guide star moved off-target by sending a stream of 9-10 RA guide pulses in the opposite direction. So "something just happened" to cause this problem – some sort of mechanical event that started the whole thing, not something caused by PHD2. Unfortunately, this is a very common problem, particularly for set-ups that aren't permanently installed in an observatory. Worse still, the guide log doesn't give you much help in figuring out what caused the original deflection, and the possibilities are seemingly endless. One thing that can help is to get a mental calibration of how little motion is required to trigger these events, particularly if you're using a longer focal length set-up. Position your monitor so you can see it from near the telescope, then start looping on a star. Now gently push on various parts of the guide scope assembly and tug gently on the various cables. You'll usually see a little goes a long way, and it's very easy to create large guide star excursions.

Here are some of the more common causes of these problems:
1. Any sort of looseness in the mount, tripod, pier, or scope assemblies
2. Dragging cables
3. Wind gusts
4. Anything that jostles the scope, camera, pier, or tripod such as moving around near the scope

Dragging cables are a particularly common problem, which is why experienced imagers do a careful job of routing and securing them. Especially in cold weather, these cables become stiff and unyielding, so if they touch or rub against a stationary surface you will probably see guiding problems.

Obviously, the list of possibilities is endless. Some of the reported causes can be pretty funny assuming, of course, it's someone else's problem. Here are just a few real-world examples:
1. The family cat poking around in the observatory at 3:00 in the morning
2. Owls landing on the end of the telescope tube
3. Leaving a rolling observatory chair near the end of the declination axis (ok, that was me)

Harder Problems

Going beyond these fairly simple examples, you may encounter problems that are a bit more sophisticated. Often, you can make an initial judgment about whether you are under-correcting (falling behind) or over-correcting (creating oscillations or saw-tooth patterns). Those types of problems are covered next.

Problems with Under-correction

The most common form of under-correction is something you're likely to see – declination backlash. Here's a fairly typical example:
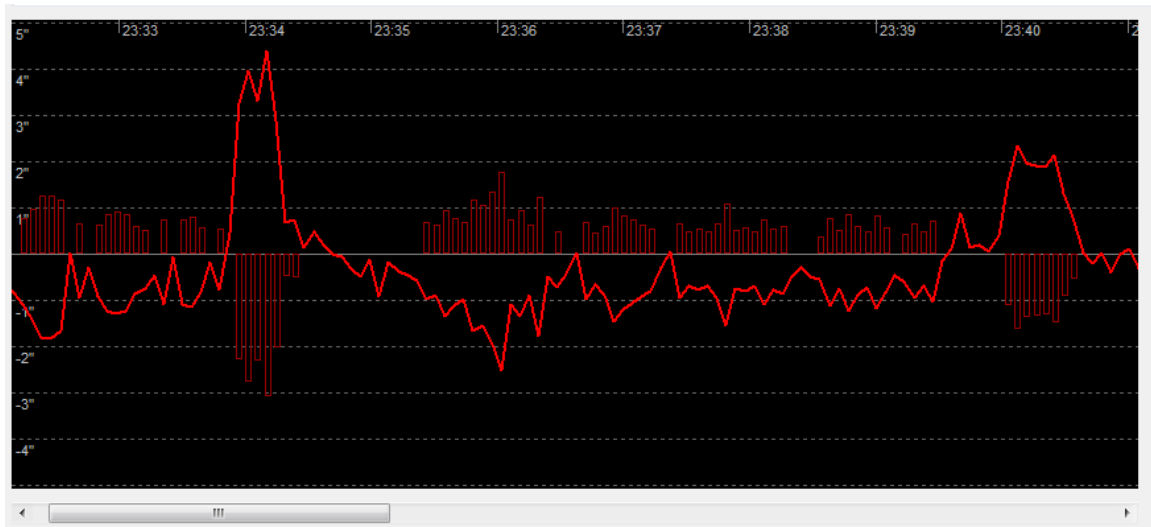

To emphasize the point, I've chosen to display only the Dec behavior. You can see that the mount becomes quite unresponsive whenever there's a change in direction for the guide commands. When that happens, it takes a substantial amount of correction in the reverse direction to get the guide star moving back to the target point. This is a pretty standard symptom of backlash. Most geared mounts have this problem to some extent, and it's caused by looseness in the gear mesh. Some amount of looseness is required in order to avoid binding of the gears, so it becomes a matter of how much is too much. The guide algorithms for declination do a pretty good job of minimizing direction reversals, and the PHD2 declination backlash compensation feature can help to control it if the inherent backlash is not too large. Improving the mount behavior is always a good first step assuming you have the means to adjust the gear mesh. If not, you may need to guide in only one direction for declination, as described in the PHD2 help document. Note that this problem almost never occurs in RA if you're using a guide speed of 1X sidereal or lower. That's because the RA drive system won't have to actually reverse direction. Instead it will just slow or pause for the length of the guide pulse and then continue rotating in the same direction.

Although you're unlikely to see RA backlash, there's another problem that looks somewhat similar:
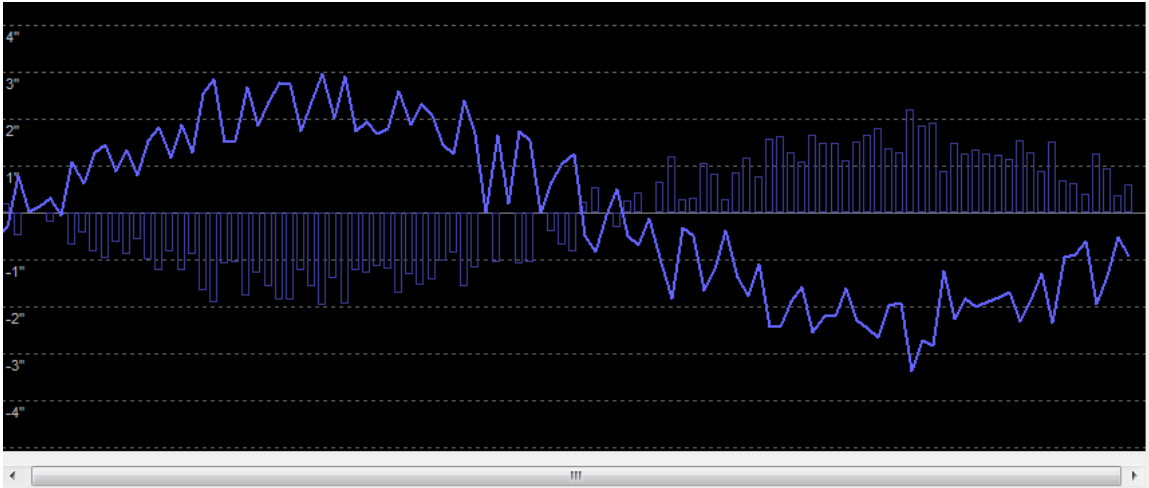
This time, we're only looking at the RA data and we can see that PHD2 was not able to keep up with the guide star movements. But we know it's not backlash for the reason just described. The clue is the sine-like appearance of the curve, which is very suggestive of RA periodic error. In this case, the periodic error in the mount was huge, and the combination of a long exposure time and conservative guiding parameters meant PHD2 was always "behind the curve." This is why applying periodic error correction in the mount is always recommended, assuming the option is available. If that can't be done, you'll probably have to use shorter guide exposure times to avoid falling behind. A good way to confirm the diagnosis is to run the PHD2 Guiding Assistant and measure the behavior of the mount with guiding disabled. You'll need to let it run for at least as long as the worm period in your mount – usually something in the 6-8 minute range or thereabouts.

A similar kind of problem can sometimes be seen in Declination, like this:
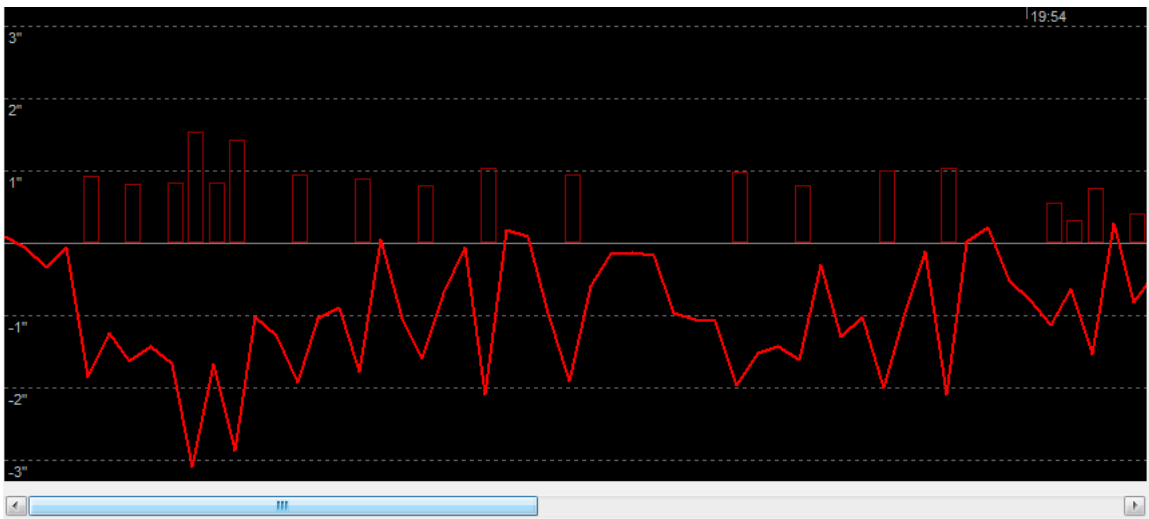
Here again, PHD2 was unable to keep the guide star near its target position. All of the guide pulses were in the same direction so we know it isn't a backlash problem. Instead, this is a symptom of very large polar alignment error. The corrections needed to keep the guide star on target were large and needed to be applied frequently – and that wasn't happening here. For Declination, there's nothing wrong with guide corrections all in the same direction, that's even a good thing. But if the corrections aren't getting the job done, something else is wrong. To check the diagnosis of a large polar alignment error, use either the Guiding Assistant or the Drift

Alignment tool in PHD2 to measure the error. Bad polar alignment, at least, should be a problem that's easy to fix.

Other kinds of things can lead to guiding under-correction, but these are usually more complicated and require a good understanding of the guiding algorithms. As a general rule, you can consider decreasing a min-move setting or increasing an aggressiveness setting if you're completely convinced the problem is not caused by an issue with the mount.   You should also be sure you're using a guiding algorithm that is appropriate for the axis. The Resist-Switch, LowPass, and LowPass2 algorithms apply a lot of damping and are best used for Declination. However, they can result in under-correction when applied to the RA axis. Conversely, the Hysteresis algorithm is usually well-suited for RA but can lead to too many direction reversals and oscillation when applied to Declination. These are generalizations, of course, but it's best to stick with the default algorithm choices until you have a clear idea of what you're doing.

Problems with Over-correction

The most common source of apparent over-correction is "chasing the seeing." Here's an example:
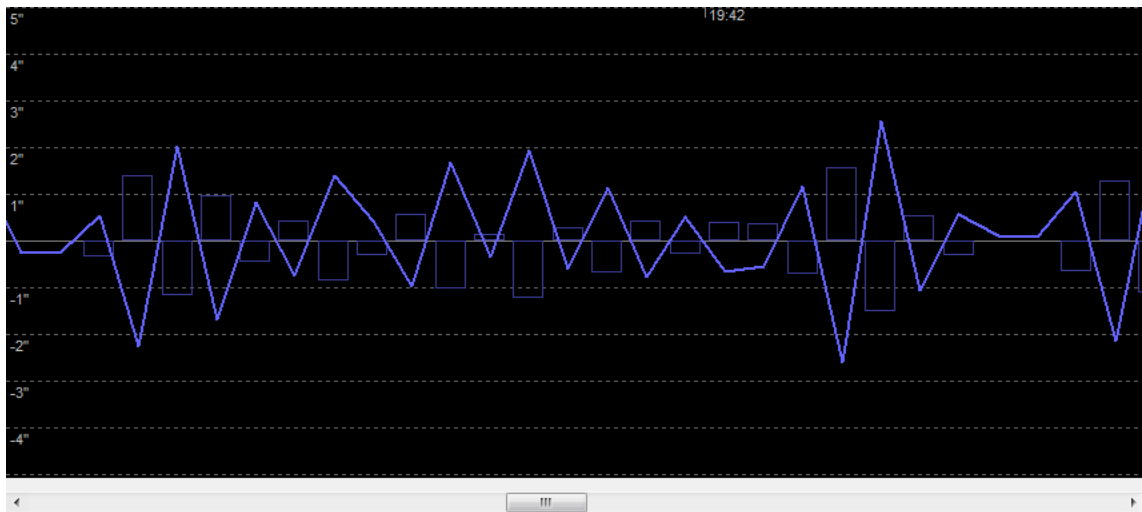
In this case, we can see that RA corrections were being applied after every exposure, and they frequently reversed direction. The result was a saw-tooth pattern and no real stability in the RA guiding. A common cause for this in RA is using a min-move setting that is too small for that night's seeing conditions or an exposure time that is too short (e.g. 1 sec). By trying to react to every guide star deflection, PHD2 was falling victim to the under-sampling problem described in the appendix. Increasing the min-move setting or perhaps the exposure time can help to improve this situation. You can also run the Guiding Assistant for a couple of minutes to get a measurement of that night's seeing behavior and a suggestion for an appropriate min-move setting. If you think the min-move setting is correct, you can slowly reduce the RA aggressiveness setting or increase the RA hysteresis value. But you need to be careful about making these changes, being sure you're taking enough time to evaluate what each change has done. This adjustment process will be described in a later section.

Over-correction can also occur in Declination, again driven by a min-move value that is too low:
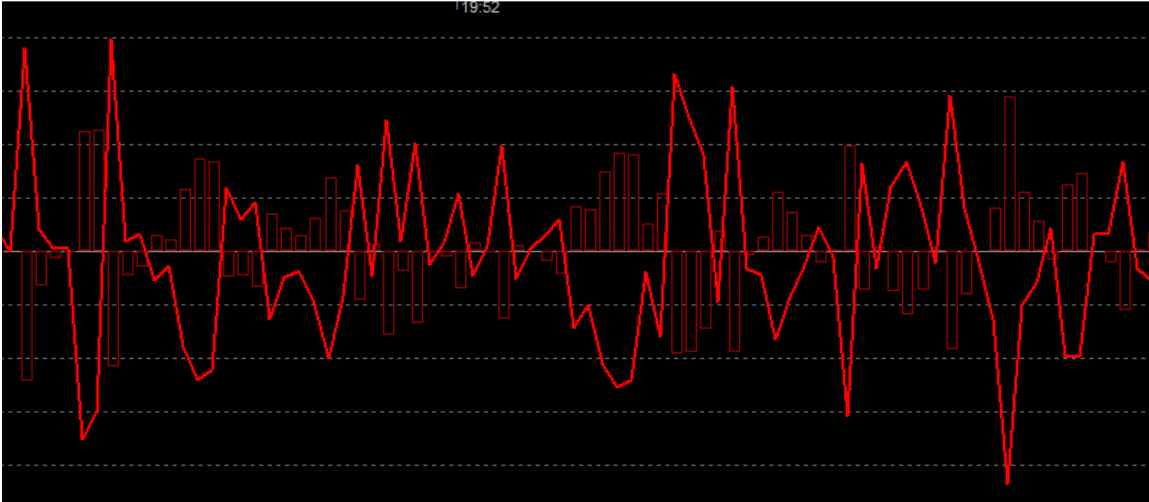
There are two likely sources for all this oscillation in Declination: seeing effects and a futile effort to correct for them. In this case, the user had specified a min-move of zero! Remember, the Dec motor will be turned off for much of the time, and you really only want to issue guide commands for slow-and-steady changes. Infrequent commands may be needed to correct for larger seeing excursions or small mechanical issues, but these are comparatively safe and not likely to create the sort of instability seen in this example. As mentioned earlier, you may also trigger over-correction in Declination if you're using a too-aggressive guiding algorithm such as Hysteresis.

Another form of Dec oscillation can result from setting a declination backlash compensation in the mount. Here's a result of that:

Two things were happening here. First, whenever PHD2 issued a Dec guide pulse, there was a large over-shoot. Second, this caused the Dec guiding algorithm to pause for a bit, not wanting to drive the system into instability with increasingly large oscillations. Once the algorithm decided the star wasn't going to bounce back on its own, it issued another guide pulse in the opposite direction and the cycle continued. This was all caused by having a backlash compensation parameter set in the mount. Once that was set to zero, the oscillation disappeared. The backlash compensation feature in PHD2 is different because it will automatically adjust to avoid these oscillations, and it seems to work pretty well for mounts with modest Declination backlash. Backlash settings in the mount almost never work when you're also guiding.

Over-correction in declination can sometimes be triggered by one problem and then exacerbated by something different. Unfortunately, there's no law to say you can have only one problem at a time. Consider this example:
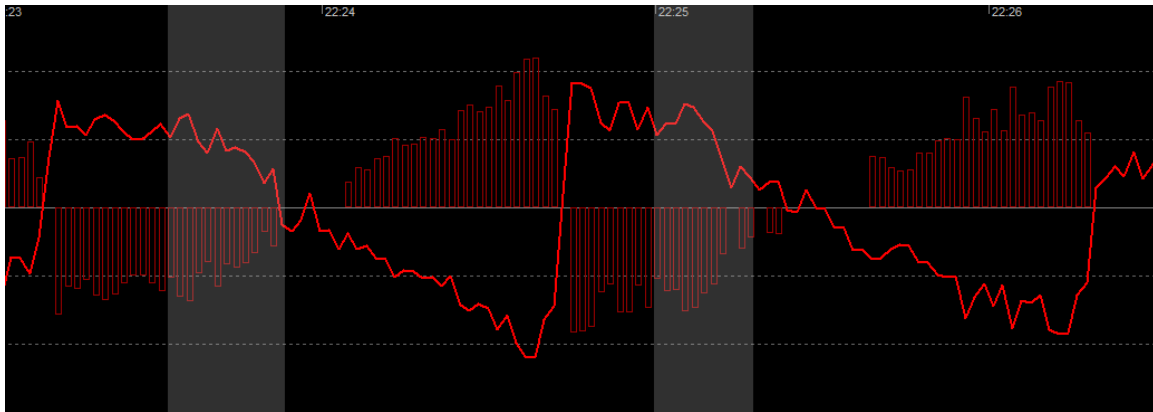
This starts out looking very much like declination backlash – whenever PHD2 reversed the direction of guiding, there was a delay before the mount responded and actually moved in the correct direction. And this is backlash, no question about it. But notice that when the mount did start moving in the right direction, it subsequently over-shot the target. This part of the behavior isn't backlash but is likely a second problem. In the initial phases of the direction reversal, the Dec gears were changing direction and weren't fully engaged for a period of time. That's the backlash part of the problem, where the Dec drive train had effectively entered a dead zone. But once the gears did engage, the rotation of the Dec motor was still not able to overcome a resistive force in the drive train. This resistive force is often called "stiction", a shorthand form of "static friction." At this point in the guiding process, the Dec motor continued to turn in the right direction, but the energy wasn't being translated into the desired rotation of the entire Dec axis. Instead, it was causing other sorts of deflection or flexure in the gear train that resulted in a temporary storage of this energy. Once the static friction was overcome, the stored energy in the other parts of the drive train was released, and the axis turned abruptly in a spring-back action. That was the likely source of the over-shoot in this example. Reducing the backlash as much as possible would be the first step to resolving this problem. Beyond that, you might need to look at things like Dec balance, lubrication, fasteners in the gear train, or anything else that can add to the static friction or mechanical elasticity on the Dec axis. Usually, this stuff is considered in the design and manufacture of the mount, so there may not be much you can do about it. If you can't find a mechanical solution, you might need to use lower aggressiveness settings or even guide in only one Dec direction.

Considerations for Adaptive Optics Devices

Generally speaking, amateur-grade adaptive optics (AO) devices can only deal with some of the problems that create imperfect guiding. AO's can usually mask or at least improve the behavior of an under-performing mount because most of the guiding adjustments are accomplished by moving a small tip-tilt mirror – not 70 pounds of telescope gear. Basically, the mount is rarely asked to do anything beyond basic tracking. Problems with backlash, stiction, imperfect sidereal tracking, and mount controller software are largely eliminated with AO use. In addition, because an AO is inherently an off-axis-guider, it also eliminates differential flexure problems. These are significant benefits and explain why many serious imagers use AO's.

What an AO can't eliminate is guide star movement due to seeing, at least not under normal conditions. Unless you can find an unsaturated guide star that produces a good SNR with exposures of 1/20[th] or 1/50[th] second or faster, you will still under-sample the seeing. For most AO users, these wonderful conditions basically never occur. Remember, you don't get partial

credit for under-sampling the seeing – the guide star position measurements and the resultant corrections will be inherently inaccurate, just as they are with normal guiding (see appendix). For that reason, AO users are still advised to use exposure times of longer than 1 second to avoid chasing the seeing. Somewhat ironically, many users of AO devices report getting the best results when they have very good seeing and/or an under-performing mount. In both of those cases, the performance of the underlying mount is the limiting factor in guiding, and an AO can substantially improve it.

Analyzing logs from AO sessions is not very different from what's been described before. Of course, you can tolerate a higher rate of guide commands and a bit more oscillation or over-shoot because there's no real hardware penalty involved – no backlash or stiction, for example. But you still need to watch for over-correction due to seeing, and you should probably try using one of the history-based guiding algorithms such as hysteresis on both RA and Dec axes. By adjusting the aggressiveness and hysteresis settings, you should be able to get reasonable results.

Making Guiding Parameter Changes – "Fooled by Randomness"

Trying to fine-tune guiding by making changes in the guiding parameters requires patience, something most of us have only in very limited amounts. The problem is what we've described before – much (or most) of the guide star movement is caused by seeing, a not-quite-random physical process we can do nothing about. Here's a typical parameter-tuning process we've probably all engaged in:
1. We notice the real-time guiding graph is looking pretty choppy. Not good.
2. We make an educated guess that we need to lower the aggressiveness and maybe boost the hysteresis parameter.
3. Guiding improves almost immediately – voila! Make a note of these settings, they're golden.
4. 5 minutes later we look at the graph again and now it looks worse – it looks like the guide commands are falling behind a bit.
5. We put the guiding parameters back where they were originally.
6. Guiding improves almost immediately – what the heck!!!

This is mostly just an exercise in chasing near-random behavior because the seeing conditions are changing. The only way to make any sense of these things is to look at much longer time intervals and to adjust parameters only in small increments. Even then, you are likely to see major differences in seeing on an hourly, nightly, and seasonal basis. In my experience, most people's guiding problems have little to do with the PHD2 guiding parameters and a lot to do with all the other mechanical and physical things we've talked about. There's nothing wrong with trying different guide settings, but it should be done with a clear understanding of what the parameters do and what specific behavior you're trying to change. You really want to develop a hypothesis about what the problem is, then make an adjustment that will confirm or disprove your thinking. Wholesale changes, something we see all too often, are never a good idea. You also need to accept that some nights just have bad seeing and no amount of fiddling around is going to get you good guiding results. At that point, you're probably better off going back in the house to watch a rerun of Gilligan's Island.

Differential Flexure – "The Dog That Didn't Bark"

If you're like most people – having a life, for example – you're probably not looking for reasons to analyze guide logs. Instead, you'll probably get into them only because you see problems in your images, typically elongated stars. So a logical process for you to follow would be to note the time the problem image was taken and then look at the guiding performance in that specific time window. And you may find things like what we've discussed in this document or some variants of them – in other words, clear evidence that guiding problems caused the bad image. If you're guiding through the main imaging scope, this will usually be the case. But if you're using a

separate guide scope assembly, your careful log analysis may reveal …… nothing. The guiding might be very good or at least completely consistent over a long period of time, yet you are seeing elongated stars in your deep-sky image. This is usually a sign of differential flexure, and it is not a welcome discovery.

We can cover the mechanical basics in short order. Any sort of telescope is going to flex and sag a bit as it moves around to different parts of the sky – that's just gravity and physics. And the whole scope doesn't behave as a single unit, that would be too easy. All the individual components - the focusers, the extension tubes, the cameras, the mirror - are going to flex and sag by slightly different amounts depending on their mass and where they are located. If you're using a large telescope, especially an SCT, the amount of flexure can be substantial, keeping in mind that your camera is typically measuring things with an accuracy of 3-9 microns (a human hair is upwards of 15 microns in width).   Now think about what happens with a separately attached guide scope and camera. Obviously, the same physics apply, but the amount and location of flexure will be different from that of the main scope. That is what is meant by the term '*differential* flexure.' So if you're guiding through one scope and imaging through the other, the guiding system is not going to see or correct for exactly the same motion that shows up on your main image.

If you're shooting at long focal lengths, say 2000mm or above, you'll be lucky if you can take long exposures using a separate guide scope. It can happen, some people do it, but most cannot. And the successful minority are often just lucky or benefit from mysteriously offsetting errors. This is why most imagers use off-axis guiders for long focal length setups. That way, whatever flexure is present in the main scope will be visible to the guiding system, and corrections will be made. This is often a bitter pill to swallow because off-axis guiding brings its own complications and expense, but there's often no good alternative. Many imagers go into denial about this and say "I've got everything clamped down really tight, nothing is moving, it can't be differential flexure." Yes, it can be. Every single mechanical interface on the system - clamped, threaded, geared, set-screwed, whatever - has some possibility for movement. And all you need is a movement of a few microns to create a problem. The flexure may not even be in the guiding system itself, it could be part of your main system. There's no easy way to know which optical system is causing the problem, only that the two assemblies are flexing by different amounts at different rates.

If you suspect you're seeing the results of differential flexure, you can do a simple experiment to confirm it. Just take a sequence of short exposures with guiding on, as you normally would. Choose an exposure time such that the stars in each raw frame are at least acceptable. Now stack those images without first aligning them. The stacked result is likely to show very elongated stars, and the size of the elongation tells you how much differential flexure you have. Sometimes, the star centers will shift by roughly equal amounts from one frame to the next, something you can see by quickly "blinking" through the raw frames. It's worth doing this test because elongated stars can be caused by other things. Optical problems like poor collimation, tube currents, or other sorts of thermals near the scope can also result in elongated stars. These will more often show up on short exposures, so it's good to confirm the differential flexure diagnosis before tearing into all the mechanical connections.

To conclude this topic, if you use a separate guide scope assembly, there will be some amount of differential flexure, period. If you're fortunate, it might not show up with the exposure times you want to use. This is even quite likely if you're imaging at relatively short focal lengths (e.g. below 1500mm). Or you may be able to tighten and adjust things so the differential flexure is essentially invisible in your images. But it is still there, lurking in the background – it's just gravity and physics conspiring against you.

Appendix

Image Scale and Measurements in Arc-Secs

The real-world things that affect guiding – bad seeing, mechanical displacements, flexure, etc. – create guide star movements that are best measured in units of arc-secs. That's why PHD2 wants to know your camera pixel size and focal length – it can then convert the pixel-sized movements it sees on the guider images into units of arc-secs. That's what we call "image scale" – how linear measurements on the camera sensor translate into angular measurements. Just to be clear, 1 arc-second is only 1/3600 of a degree - the angular size of a 0.004mm object held at arms-length. Discussions of guiding performance using units of pixels are completely meaningless, yet you will often see forum discussions that do that. Why is it useless? Let's take a simple example, one that periodically shows up on the PHD forums. Let's suppose you've been working with a separate guide scope having a 300mm focal length and a guide camera with 5u pixels. You've been working with this configuration for a while and observe that most of the star motion in the real-time guiding graph falls within a range of 0.3 px. That seems like a pretty good number, so maybe you even brag a little on the Yahoo forums. But now you switch to using an off-axis guider set-up, and the focal length of the guide scope is 2000mm, not 300mm. Suddenly, the real-time graph is showing star movement all over the place, huge swings of 2 pixels rather than 0.3. It's the same mount, what the heck? Must be the stupid guiding software. Well, not really – you aren't taking into account the image scale. With the guide scope, the image scale was 3.4 arc-sec/px – so the 0.3 px movements were really 1.02 arc-sec movements (3.4 x 0.3). But what about the new off-axis guider set-up? The image scale for that is 0.52 arc-sec/px, and the "horrible" 2 px movements are really just displacements of 1.04 arc-sec – essentially the same as before. You can use the calibration step calculator in the brain dialog (guide tab) to do the math for you and compute your image scale. But bottom line – always think of performance in terms of arc-secs. Of course, that's why the graphs we use with PHD2 always default to displaying things in arc-secs.

Astronomical Seeing and Guiding

You can't get very far looking at guiding performance without coming to grips with astronomical seeing. This is a complex subject, not something to try to deal with here. But the Cliff's Notes version goes something like this. "Seeing" is the term given to the positional jitter and sudden brightness changes of stars we see (or image) through a telescope. It is atmospheric turbulence, caused by the movement of thermal cells in the Earth's atmosphere, and there's basically nothing to be done about it. Light is refracted as it passes through each atmospheric cell, so when you look at a star, you're really looking through a column of air that is behaving like a column of little lenses. That might be ok except that the refraction of the light by each cell depends on the temperature of that cell, and the cells generally have different temperatures. And of course, the atmosphere is very dynamic, so these elements are all moving around at various speeds, coming into and then leaving the column of air you're looking through. Thinking of it this way, it's a wonder we can image anything. Particularly with longer focal lengths, this atmospheric seeing is the single biggest source of the guide star movement we see, and we're stuck with it. Can't we guide it out? The short answer is 'no.' The longer answer is also 'no.' The movement of the atmospheric cells means the guide star position is changing at rates of 10's to 100's of times per second. You aren't going to be able to measure it and react to it nearly fast enough, even using amateur-grade adaptive optics devices. Professional observatories are able to do it to a large extent by employing very expensive measurement devices, artificial stars, and mechanisms that can both deform the mirror and shift the image at very high frequencies. That's not us.

From a guiding perspective, we are "under-sampling" the seeing behavior. By the time we've taken an exposure, downloaded the image, computed the location of the guide star, and then

transmitted a guide command, the star position on the sensor has moved – probably 10's or 100's of times. Basically, we're always dealing with outdated information about the guide star position, so the guide commands are inherently inaccurate – and that doesn't even take into account whatever shortcomings the mount has in precisely executing the guide commands it receives. Essentially, the star movements we *can* correct for - drift, periodic error, atmospheric refraction, etc - are hiding in a sea of noise created by the seeing conditions. This is a profound limitation of conventional guiding and a major reason why elaborate guiding algorithms or process control models have trouble producing any significant improvements over simpler algorithms. It's also the reason why using a combination of very short exposures and very aggressive guiding parameters invariably results in poorer guiding results.

Seeing and Exposure Times

The high-frequency, seeing-induced star motion seen by the guide camera is strongly affected by the length of your guide exposure. Look at the following plots of guide star motion as the exposure time is increased:

0.5 sec guide exposures, Dec

1 sec guide exposures, Dec

2 sec guide exposures, Dec

4 sec guide exposures, Dec
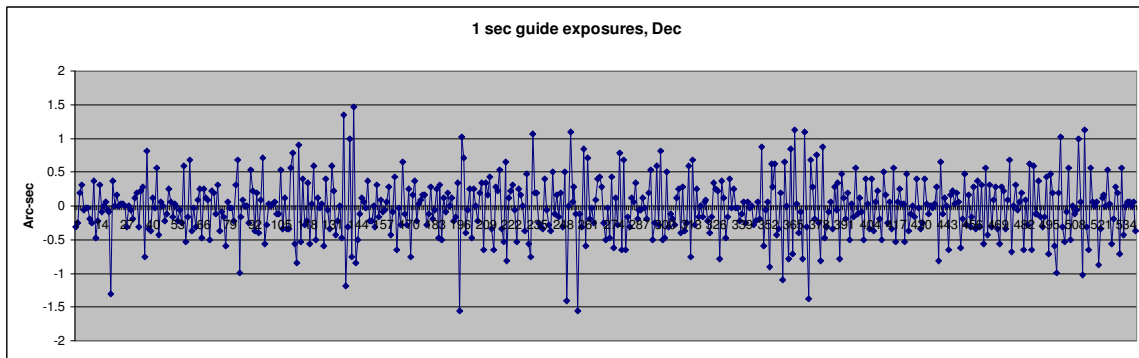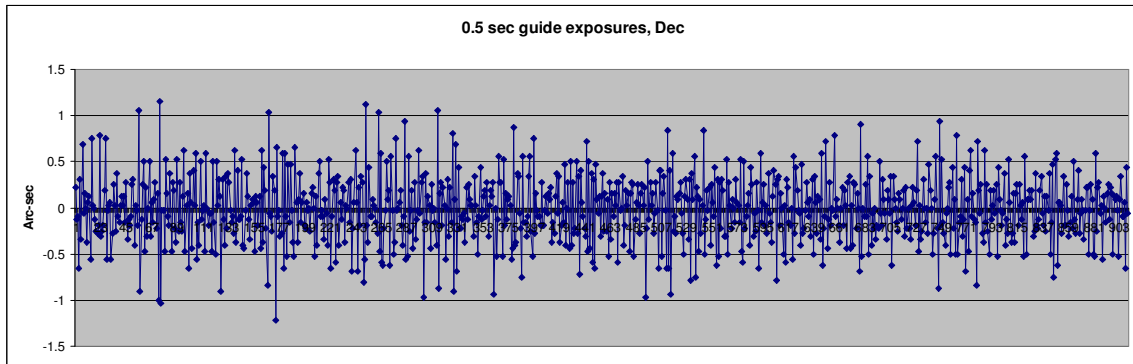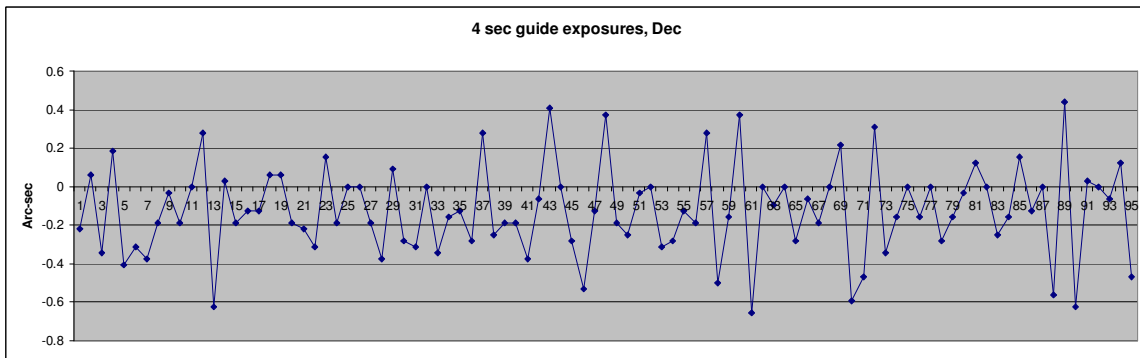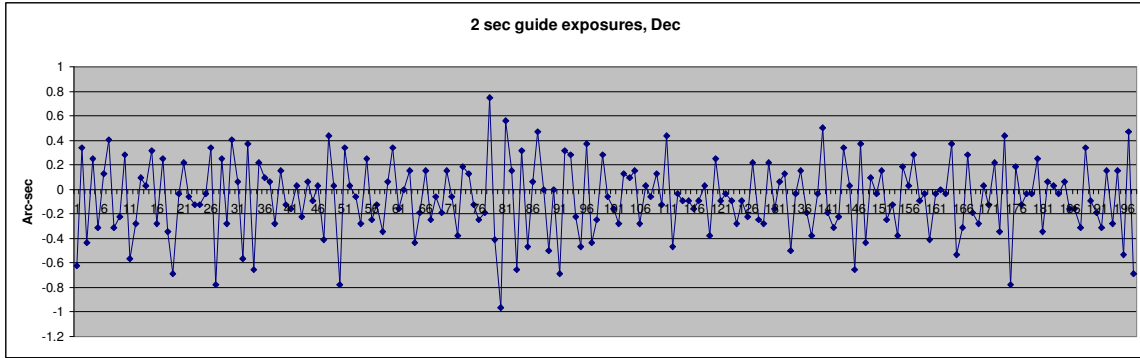
As the guide exposure is increased, the range of star motion decreases – the "envelope" of star motion goes down by a factor of 2X when switching from 1-sec exposures to 4-sec exposures. Essentially, the camera sensor is averaging the changing light pattern of the star and smoothing the result. These measurements are still inaccurate because of under-sampling, but the longer exposures make it easier for PHD2 to isolate and identify the lower frequency errors that really can be improved through guiding. Obviously, there is a practical upper limit to the exposure time. Typically, it will be limited by the length of time your mount can run on its own without needing a correction. Small errors from periodic error, drift, flexure and other sources need to be corrected before they become large enough to ruin an image. Finding the right balance will always depend on both the seeing conditions and the quality of the equipment. As a starting point in PHD2, we typically recommend using exposure times of 2-4 seconds.

Conclusion

Hopefully, this tutorial will help you get a better understanding of the guiding results you're getting and perhaps some clues for how to improve them. Even if you have trouble figuring out a problem, you can ask more informed questions on the support forum and be better able to evaluate the responses you get.

If you have comments or suggestions for improving the tutorial, please post a message on the Open-PHD-Guiding forum or send me a message at bw_msg01@earthlink.net

-- Bruce Waddington
V 1.0   December 2015
V 1.1   January 2016