REVIEW

WORDS/PICTURES: JIM DAVEY



(Stability Augmentation System)

'The advent of three axis electronic Stability Augmentation Systems (SAS) is the most significant step forward in model helicopters since Dieter Schlüter introduced the Hiller control system to his model – discuss.' That's not quite how Mr. Editor Jon Tanner put it to me but it is a better attention grabber!

Jon has heard me mutter about some of the jargon that is wrapped around these new toys and asked me to have a go at demystifying the subject a bit. This task is not helped by the tendency of the manufacturers to invent 'marketing' names for engineering functions – in some cases different manufacturers use the same name for different functions. I find the use of 'precise' as a style as particularly humorous – who would want 'imprecise'? At times there seems to be a perverse desire to make the whole thing as complicated as possible with no obvious gain for the user. However, here goes.

First a Little Background

A few words on the subject of how I qualified to get sucked into this are in order and will also give some context to a few, more controversial, statements I might slip in here and there.

I have a somewhat eclectic Electrical Engineering qualification (a long story), which included a heavy emphasis on control system theory. I retired (survived?) recently after 42 years in the aerospace industry. Along the way, almost as sidelines to my real job I was responsible for the electrical design of the yaw damper electronics for a commercial airliner, for re-engineering the SAS electronics for a large transport helicopter and for designing a solid state replacement for the mechanical gyros of another large commercial helicopter.

In the model area I have owned and/or flown a bunch of models 'no-bar', ranging from '400' size electric up to 3.2 metres turbine. Most I recall fondly – some I just recall... I have setup and flown V-bar (early), Sk-360, 540 and 720, Micro Beast and Tarot ZYX and have dabbled with a DJI Ace1 – but that's very much another story. I also like to download the user interface for other units to see what controls are available to the user. I'd like to play with a few more so if anyone cares to send them to me care of the magazine...

In my current fleet there are five heli's with SAS onboard.

These days I am mainly into scale, am not into 3D but have been known to do some smooth aerobatics and am addicted to autos.

Next a Short Historical Dissertation

Buried inside these cute little controller boxes is usually a variation on something, which is called a PID controller. PID is the acronym for Proportional, Integral and Differential. The exact objective Jon set me was to explain what the various terms mean that come along with the



SAS. For example what does 'P' or 'I' or even 'D' mean and why do I now have a bunch of different gains to play with?

The math behind this is familiar though oddly enough I can't recall the PID moniker being used when I took those control theory courses, though it was back in the '60s so who knows... Before the word 'math' causes a mass turning to the next article, there will be no math – OK?

So how to explain? Most people, including me, find it relatively easy to grasp physical concepts, so let's start by revisiting the thing we just replaced, the flybar, and start from there. Stick with me on this – you might enjoy the journey.

If you have ever flown 'no-bar' you will recognise that a rotor system responds very rapidly to both cyclic commands and changes in relative wind – especially in model sizes. The contrasting need to make very small and accurate commands in the hover and the need to hold a bunch of forward stick to keep the model moving forward can be... challenging. Well designed and heavy blades help because they slow the response of the disk to both the pilot inputs and changes in airflow, giving the pilot time to stay in front of the beast. Even with good heavy blades, however, you still find yourself with the stick in some funny places, and the concept of 'neutral' has gone out the window.

Back in the day, the full size helicopter pioneers tackled the stability and control challenge from two different but mathematically related ways.

Arthur Young (Bell) invented the Bell bar (or Young's rotor), which is essentially a 'free' gyroscope, which can control the cyclic pitch of the rotor. A mixing linkage combines the pilot's cyclic inputs with those from the Bell bar. The Bell bar tries to maintain a constant plane of rotation in space and (via the mechanical linkage) controls the cyclic pitch of the rotor in such a way as to keep it rotating in the same plane as the Bell bar. The pilot can also control the cyclic pitch of the rotor when he deflects the swashplate – an action, which tries to move the plane of rotation of the rotor away from the state the Bell bar is trying to maintain.

In the resulting difference of opinion between the pilot and the Bell bar, the Bell bar initially won, the result being a stable helicopter that was very reluctant to go anywhere. The solution to this rather fundamental limitation was to introduce friction or hydraulic dampers between the Bell bar and the rotor shaft. These caused the Bell bar to re-align itself with the shaft over a period of time – the Bell bar's 'memory' of where it was would decay or be 'washed out' as the pilot persevered with his cyclic input allowing him to manoeuvre the helicopter. Thousands of Hueys can't be wrong.

Stanley Hiller adopted a physically related approach, but his flybar (as we now call it) looked like a Bell bar with 'paddles' at the ends. Here the 'flybar' alone controls cyclic pitch of the blades, with no mixing of input from the pilot. The pilot controls the cyclic pitch of the paddles (just like a model). Neglecting the pilot for a moment, the flybar works just like the Bell bar except that the mechanical or hydraulic dampers are not required since aerodynamic forces on the paddles will bring the flybar's rotational axis in line with the mast over time. Both systems do a good job of stabilizing the heli against outside influences.

Now let's compare the two approaches from the point of view of the pilot. In the pure Bell approach, the rotor disk responds immediately to pilot input but the helicopter is 'held back' by the action of the Bell bar. In a model this would give 'bounce back' after cyclic commands and you would need lots of forward stick to stay moving – a bit like those tiny coaxials. In the full size the response of the rotor is slower and if the

SAS (STABILITY AUGMENTATION SYSTEM)



A model Bell head from Vario. Mechanical mixers combine swashplate inputs with attitude correcting inputs from the flybar to apply a combined input to the blades



A pure Hiller head from an early Morley. The swashplate controls only the cyclic pitch of the flybar paddles and cyclic pitch of the Blades is directly controlled by the flybar tilt – sorry about the quality of this scan

dampers between the Bell bar and the shaft are properly adjusted all is good.

In the Hiller system, by contrast, the pilot does not generate any blade cyclic change until the flybar has responded to paddle cyclic. In model applications the slow initial response can result in the pilot over controlling, a problem that isn't helped by the fact that once you have deflected the flybar it will take a while to get back to its neutral position relative to the rotor shaft. The control demand does not go away immediately when you neutralise the stick. (By the way, you can adjust at least some SAS to emulate this behaviour - don't!) On almost all flybar equipped models today, the control/stability system is essentially a Hiller inspired system with the introduction of some direct cyclic input from the swashplate to overcome the slow start/stop characteristic of pure Hiller systems. Modellers call this a Bell/Hiller hybrid - and it works very well indeed.

Either way, the flybar is a mechanical computer (Steam Punk is up next!), which accumulates the combined inputs from the pilot and outside disturbances and 'calculates' a cyclic pitch output in response. In engineering terms it performs the function of something known as a 'leaky integrator' (I am not making this up!) and the direct input from the swash to the blades is providing something called 'feed forward' to the overall control function.

Let's dig into that last bit for a moment. You may have come across talk of tuning the response of the helicopter by adjusting the Bell/ Hiller ratio. To achieve an ideal (from the pilot's point of view) control feel, we make the direct input to the blades just enough to achieve the roll (or pitch) rate we want. At the same time we want the cyclic input to the flybar to demand the same roll rate of the flybar plane of rotation. If we get this tuning correct, the fly bar and heli will roll at the same rate and the flybar rotational axis will remain in line with the mast. Done well,



The Bell rotor head as fitted to thousands of Hueys



This is the rotor head from a Hiller XROE-1 Rotorcycle

the effect is: stability as required and immediate and smooth response to controls with apparently instant start and stop.

A moment (or more) of thought will suggest that this level of perfection can only be achieved for one unique set rpm, model weight, blade dynamics and paddle weight/area. Fortunately we can get close enough with a bit of fiddling – or in the case of the F3C flyer a lot of fiddling!

Summarising: for stability, the flybar provides us with a short term attitude reference which 'washes out' over time - heavier paddles 'wash out' slower and make for more 'stability' but also slow response. For control we have the ability to command the flybar to move via paddle cyclic - bigger paddles make for faster response. For 'crispness' we have the ability to feed forward cyclic directly to the blades via the Bell/Hiller mixers. Both Bell and Hiller systems are mechanical variations on a 'PI' controller theme with no D. The use of the Bell/Hiller mix introduces something called 'feed forward', which we will return to later. Note that the behaviour of the paddles is a combination of size, weight and flybar length and they all interact... Right, let's take-off the flybar and get back on track

Homework – Design your Own SAS

Using the mechanical lessons we learned above, let's go design a basic (I was once told by my Engineering VP never to use the word 'simple' in front of the CEO since he would then challenge the engineering cost) SAS. Our bag of parts is going to consist of some rate sensors (gyros) and a small computer.

Since the old flybar did such a great job, perhaps the obvious approach is to create a virtual flybar – which we could call a V-bar if someone else had not thought of that already.

The first snag is that our gyros are not 'free' gyros like the Bell bar; they are classic 'rate'



After all the years, the design still works wonderfully well



The basic Bell system combines the demand from the pilot with the reference provided by the Bell bar and uses the result to control cyclic pitch. Helicopter movement influences the Bell bar via mast mounted dampers



The Hiller system controls cyclic pitch indirectly. Pilot demand moves the flybar aerodynamically via cyclic pitch of the paddles. Helicopter Movement influences the flybar via aerodynamic forces acting on the paddles

gyros. These gyros just tell us how fast we are rotating about each axis (roll, pitch and yaw), which is not enough information. However, if we use our computer to perform a mathematic function called 'integration' of the rate signal from the gyro we can work out how much we have moved away from where we started – in angle terms. If we rotate at 10 degrees per second for one second, we have changed our attitude by 10 degrees.

If we now deflect the swashplate by an amount proportional to the value we have calculated in our integration (and get the direction right) we will start to move the heli back to where it started from. If we continue the integration process, as we return to where we started, the integrated value and the resulting swash angle will return to zero at the same attitude as we started from. If we do this process continuously rather than waiting for the model to have actually moved away from its starting point, we have just built a basic 'l' controller to maintain a steady attitude – in fact we have emulated the Bell bar without dampers. So, we are (possibly) stable, what about manoeuvring?

Let's try another approach. We will make our control stick input scale the rate of roll we want. Say we want full stick to be 360 degrees/sec (one complete roll in 1 second) so half stick is 180 degrees/sec etc. Now let's do another mathematical operation and compare the stick demand with the roll rate that the gyro sees and apply a swashplate deflection proportional (= P) to the difference. In a perfect world the heli would now roll at a rate that would cause the

SAS (STABILITY AUGMENTATION SYSTEM)

difference to be zero. However, in the real world we won't get the rate we want for all kinds of reasons and there will be a (rate) error. However, if we use our handy dandy integrator from a few paragraphs ago to develop a further change in the swashplate position as a result of integrating the rate error over time we can drive it (the rate error) to zero and we will get exactly the roll rate we want. Of course, if we want to stay still (hover) we have exactly the same situation as we discussed above – we are just demanding zero rate by having the stick in neutral. We now have 'P' and 'l' in our system and have built a PI controller.

Imagine the heli in a stable hover, now move the stick a bit and we have created a rate error to input to the swashplate and our integrator. Our integrator starts integrating and will modify the swashplate position to try to match the roll rate we want, exactly. If we got the direction right, the heli will start to roll in such a direction as to generate a rate signal from the gyro that opposes the stick demand that started this. The integration action will cause the roll rate to increase until the stick demand and gyro signal cancel. At this point the integration action will stop with just enough cyclic deflection applied to maintain the required rate of roll. Cool! We are done right? Well no - we just pretty much emulated the Hiller system with all the same warts of that system. Actually, we are a lot better off because the 'P' term gets some swashplate movement as soon as we move the stick, but we are not quite there yet. You can guess were we will go next. Throw in some additional direct movement of the swashplate in response to stick movement and we will emulate the classic Bell/ Hiller mix electronically.

But we can and will do even better than that – plus we need to address that 'wash-out' we talked about earlier.



The basic PID controller moves the swash plate on the basis of a computation based on the difference between the desired and actual rate of motion. Three terms are used in the computation: proportional to, integral of and differential of the error



An improvement on the basic controller is to introduce a feed forward capability. This function improves response and provides a more 'natural' feel

The Real Deal

So far we have created a basic control system that computes a correcting movement of the swashplate based on the difference between the roll rate you want and the roll rate the model is actually executing. The special case is when the stick is in the middle, which means stay still.

At some time you might have asked yourself if you could have used a standard rate gyro to make a flybarless model manageable. Various folks have tried and it does sort of help, but not as much as you would like. There are a few problems. The first is that if you turn up the gyro gain as much as you feel you want to, the whole model will oscillate (the control loop is unstable) unless you have pretty heavy blades, in which case - why bother with the gyro? The second is that we are missing that short term attitude reference that the flybar provided (the 'I' term plus washout). You could try a regular gyro in heading hold, but the relative values of rate and hold terms (P and 'l' since you asked) which work for the tail are way off for the pitch and roll. Also without the washout or decay we talked of earlier, the handling will have 'issues', not least of which will be the ease with which you will roll the thing over on lift off!

Finally, we can introduce a third process which is called differentiation ('D' in the PID) which gives us a measure of the way the error is changing – an acceleration estimate if you will. Throw in all three and we control the swashplate based on the sum of the 'P', 'I' and 'D' values calculated from the behaviour of the error (difference) between where you place the stick and the roll rate of the heli – well pretty much!

If you research PID controllers in the internet you will notice that there are a number of variations on the 'classical' form and that there are whole books written on 'tuning' the things so what hope do we have?

Well, if the designer of our chosen SAS did a good job we will be in pretty good shape. In fact we may not be given access to all the terms in his implementation because it was designed making certain assumption about the models it will be applied to. If our model meets those assumptions we need to do little more than make fine tweaks. Let's discuss the assumptions for a minute.

So far I have avoided much discussion of 'gain'. When discussing our little computer earlier, we were a bit vague about exactly what is output. When the computer drives the swashplate, the objective is to command a particular cyclic pitch angle. Between the computer and the blades are the servos and the linkages. Servos don't vary by much in terms of 'gain', but the linkage is in the hands of the modeller. The closer you get this 'mechanical' gain to the expectations of the designer, the more likely you are to get good results using a default setup.

Servo speed and accuracy do matter. Accuracy is pretty obvious, and speed matters because slow servos will introduce what engineers call a phase lag in the system, which erodes stability margin. Basically if your servos are too slow, the whole thing will oscillate unless you sacrifice gain – just like a regular tail gyro.

You can stray a bit from the recommended gain settings (and I do since I don't do hard 3D) and setup for less mechanical movement. I compensate by increasing the electronic gain in the SAS. Going this way is not usually a problem – at least with the units I prefer – as long as you leave the system enough cyclic range to manage the helicopter. However, go the other way... Too much mechanical gain – the effect of too long servo arms and short blade grip arms can be addressed by reducing electrical gain but sacrifices resolution and can result in poor performance. The more of the system parameters you can adjust, the better chance you have of compensating for an off datum setup. The corollary is that if your particular unit does not offer much in the way of adjustment, you will be best off getting the mechanical setup close to the expectations of the designers.

What Does What?

Unfortunately for this discussion, the facilities presented to the user vary significantly from unit to unit, ranging from a lot of button pushing and a few rotary controls with vague descriptions to sophisticated (ish) graphical user interfaces on a personal computer. Throw in the fact that certain manufacturers seem to delight in inventing their own terms for things and the current tendency to use non-engineering terms like 'agility' and 'brilliance' and it becomes a challenge to provide unambiguous advice on system adjustment.

An additional complication is that the basic PID controller's major weakness is that it has issues with changes, which can lead to non-optimal behaviour - 'bounce back' anyone? There are various strategies to deal with this problem and the acceleration and stop controls seen in many systems are associated with 'tuning' to get the best behaviour in any given application. It is guite possible for the SAS software to have additional control strategies that are effective only during starting and stopping, which modify the basic PID control loop only during starting and stopping for example. The best I can do - given that I have no inside knowledge of the actual software used in any SAS on the model market is to provide some generic guidance, based on the 'training' we have gone through above.

Gain is a generic and generally allows you to vary the amount of swash movement that will occur as a result of the SAS computations. If you rock the model you will see the swash move to correct the movement. The higher the gain the more the swash will move in response to a given heli movement (or in response to stick movement).

On some units you can adjust the effect of the various underlying calculations. For example 'Hiller gain' can reflect the authority of the 'l' term of the PID and is analogous to the relationship between the movement of the flybar and the cyclic pitch. If you hold the model and roll it to a new attitude the swash will deflect to correct the movement and then slowly return to centre. The higher the Hiller gain, the farther the swash will move for a given movement of the model. Hiller gain controls the 'grip' the SAS has over the model and is similar to heading hold gain on a regular tail gyro. The rate at which the swash returns is the 'decay' or 'wash out' rate - slow is the equivalent of heavy paddles. Some units allow you to adjust this - a good thing in my opinion. Taken to extremes, slow decay starts to approach 'attitude hold' characteristics and sounds like a good thing until you go too slow, as the handling can become very 'synthetic'. Too fast and the model will start to feel 'loose' and need a lot of forward stick to keep it moving.

P gain will reflect a direct movement of the swash in response to the gyros – equivalent to rate mode on a tail gyro. It is directly proportional to the rate error. If you can turn the 'l' gain to zero you will see this when you rock the heli. The movement may be quite small (or appear large since it is being masked by the feed forward term!) and the swash will go back to neutral as soon as the heli stops moving. 'P' gain is the same as a rate gyro on the tail. It is actually the basic control term and the right value is important. Too much and the wobble of death (WOD) will strike and too little and the model will wander in gusts and may (emphasise may) get nasty at high speed. High 'P' is rather like having heavy blades and the heavier the blades actually are, the more 'P' (in isolation) you should be able to run so it is all good!

Classically, 'D' gain reflects the rate of change with time of the roll (or pitch) rate error. However, this term can be very prone to noise (= vibration in our case) so a classical PID is likely not being implemented. With the right implementation, however, 'D' gain is beneficial in damping unwanted oscillations, particularly in the fore/ aft direction, and especially after transients and has an influence on how the model will start and stop movements. If you have too much 'D' the dreaded WOD will occur and can be very aggressive!

An important adjustment is the relationship between stick movement and helicopter roll and pitch rates. In my view, getting this right is much better than having it flat out and then cutting things back with the rate switches. Some units allow you to set this directly in degree/sec others may call it something like 'agility' which is irritating – particularly as other vendors use agility to mean something completely different! Unfortunately, not all units allow you this luxury in any type of straightforward manner.

Another very important adjustment is the amount of direct swash movement that results from a stick movement (feed forward). If this is too low we revert to Hiller handling, too high and the model is twitchy and can bounce back after control movements. Some call this the 'Bell gain' and others invent terms again or bury this adjustment in some sort of magical handling characteristic 'name'. You can see it if you give a quick stick movement (with the model not running!) you will see a small full servo speed movement followed by a slower drive to the end point controlled by the 'l' term of the controller. Mathematically, the 'P' term contributes to this effect, but on its own is not generally enough for crisp handling.

Things That Can Go Badly Wrong

Well, there are a few. Perhaps the three that are most talked about are roll over on take-off, weird handling at the end of an auto and pirouette compensation.

The first is an effect of the power of the 'l' term in conjunction with the decay rate. If you hold a

little cyclic on the ground what you should see is the swash move off until the integration rate of the roll error is equal to the decay rate, at which point the swash will stop moving (before the end stop). This may well be rather more angle than you might have bargained for and will lead to excitement if you lift off without paying attention. There is at least one designer out there who thinks it is a good idea to completely inhibit the decay for any degree of stick movement with the result that the slightest stick input will run the swash to the stops. I can guess the thinking behind this – pure rate command control, to make the model fly like a video game - but I feel it is 'ill advised'. (In my old day job the language would have been stronger!) I bought one of these units without realising anyone would put a feature like this in without allowing you to turn it off. I did fly it a couple of times but found it so irritating to deal with that I sold it. As a scale flyer it is not acceptable to be expected to pop the model into the air and catch it like so many 'this is awesome' demo videos we see. Properly setup, a good system will allow you to take-off smoothly - from a slope if need be - with a controlled lift to level before a vertical rise.

Getting back off my rant and moving on to autos. If you are used to a flybar, the degree of control falls off as the rpm drops – this does not happen with electronic control until the rpm is way, way down when other bad things will happen. The PID controller simply applies more cyclic to get the response the stick position is demanding. The result can be an expensive tendency to over control at or close to touch down. Personally, I like autoing a flybarless model, but mine are not setup for hard 3D which may be a clue as to why...

One of the key differences between the behaviour of the flybar and that of the SAS is due to something called 'frame of reference'. The flybar is spinning in space and knows nothing about the way the model under it is pointing. If the flybar was, say, tilted north and you start to pirouette, it will continue to tilt north - at least for a while, which is what we want. The SAS however, is calculating the equivalent of flybar tilt using the integrated signal from the pitch and roll gyros - with respect to the helicopter itself. If the helicopter starts to pirouette the tilt moves with it. This issue can be solved by using the yaw gyro to keep track of the pirouette and adjust where the tilt is pointed. Of course, it has to be corrected in the right direction... On units where you can set this up - do get it right or your first fast pirouette could be entertaining!

SAS (STABILITY AUGMENTATION SYSTEM)

A more general issue comes from the major deficiency of the PID controller. It is very good at maintaining the status quo but can get its underwear in a jam during changes. In our case, changes are changes in rate of rotation. There are various methods of dealing with this, and the various acceleration and deceleration adjustments that we discussed above address this.

If you don't have enough direct feed forward (Bell input in flybar terms) the temptation is to turn up the 'gain' to get faster body rates. The result can be oscillation, bad stopping etc. You can work around this with lighter blades - just like you could with an old Hiller only machine. From this you could convince yourself that the controller works better with light blades - you could, for certain values of the word 'better'. Of course, the downside is that gust response of the head gets worse - a very bad deal in my (less than) humble opinion. Personally I would much rather have heavy blades any day, and a SAS unit that lets me crank in enough feed forward to off load the PID when manoeuvring. But then - I did say I am not interested in 3D...

Parting Thoughts

When I first flew a model with a SAS on it I was impressed how well it stabilized the model and how hard it was to hover in one place with any precision – right back to pure Hiller feel. Times have changed. I now prefer the handling I can get from a SAS (with an exception or two!) and things will only get better.

For the sport flyer it is now possible to buy a full three axis SAS for less money than it cost for a decent gyro five minutes ago – and it will work very well indeed.

The masters at getting a flybar equipped model to perform over a wide envelope are the F3C guys. SAS will come to F3C soon and it will be interesting to see the impact for the rest of us. The innovation pressure so far has been in 3D, which explains why the 'scale' defaults in some systems are – well let's just call them 'unlikely'. The requirement for precision hovering in the F3C schedules is likely to drive both SAS subtlety and the accompanying model setup and blade choice.

The bottom of the column is approaching and I feel that I have only scratched the surface of this subject but hopefully I have been able to inject a little clarity here and there.

I would be happy to answer questions (if I can) by email via the magazine. *MHW*



The Bell Jet Ranger sports a flybarless rotor head

As does the latest 3D models with their very clever SAS