

“Now that it’s built, where do I put the antennas?”

Sherman Knight, duworm@aol.com

Antenna placement is a problem we experience daily. The rabbit ears on the TV are able to pull in a signal, but if I move them a little, the picture would get much better. The FM rock station starts to fizz and pop when the car stops, but clears up when the car moves just a couple of feet. The building, overpass, the nearby hilltops all have an impact on the quality of your Cell phone or the number of bars. Your GPS system is lost inside a building. All of these are antenna placement issues. If you want clear radio reception in your car, don’t drive into a parking garage.

Unfortunately, like the car driving into a parking garage, aircraft are moving which means the quality of the signal at the receiving antenna is constantly changing. When carbon/Kevlar weave fabrics were first used in fuselages, no one

was ready for the problems it caused 72 MHz antennas. A bunch of solutions appeared, none of them pretty. The most common was a 36 inch long antenna dangling from the tail and swinging in the wind. In the end, it became the accepted method of antenna placement. No matter what you called it, it was an ugly compromise.

Years later, 2.4 GHz came along and changed the antenna rules again. A 2.4 GHz antenna is very short so dangling them out the end of the fuselage was no longer an option. Fortunately, as long as the aircraft is built with materials transparent to a 2.4 GHz signal you can place the antenna anywhere inside the aircraft. Materials like fiberglass, wood, Styrofoam and plastic are “transparent” to the 2.4 GHz signal.

Unfortunately, a 2.4 GHz signal is extremely poor at “bending” around materials that are not transparent. Things that can conduct an electrical current, like metal and carbon fiber reflect or block the signal. Your body and other items with a high water content, actually absorb the signal. Because of 2.4’s inability to “bend” around them, situations occur where the signal cannot be seen by the receiver. All of the 2.4 GHz systems have to deal with this issue. Fancy signal processing, antenna with cool names or more powerful signal amplifiers cannot “cheat” this law of nature.

Kit manufacturers were caught off guard again. There was no way to install a 2.4 GHz antenna in some of these carbon fuselages. Some tried whiskers with varying degrees of success. It didn’t take



1. Contemporary 2.4 GHz receivers, all with multiple antennae. The Spektrum 9300 on the left has a satellite receiver as well.

long for model makers to start building “2.4 GHz ready” fuselages that replaced the carbon in the nose of fuselages with fiberglass/epoxy. Nonetheless, because of the cramped space inside many sailplanes, antenna placement is still difficult.

Diversity and Redundancy

Antenna placement became more flexible when some of the 2.4 GHz manufacturers started providing receivers with longer, multiple antennas. Internet chatter (by individuals, not brand representatives) concerning the reasoning behind two antennas was interesting to follow. One camp insists that their system is so good that it does not need a second

antenna. The other camp claiming that two are always better than one. Today, nearly all the manufacturers include a second antenna on their more expensive receivers. (See #1 above.)


But, no 2.4 GHz receiver actually “needs” more than one antenna. There is a lot of poor information (a combination of misinformation and speculation) on line concerning redundant receivers and antenna diversity. There is nothing complicated about the concepts. Frankly, both are just a little common sense.

(See #2, next page.) The first diagram (A) is an example of the “donut” reception pattern of a single antenna running through the middle of the donut. As you

can see, when the antenna is pointed at the transmitter, there may be tremendous signal loss.

The second diagram (B) demonstrates a second antenna placed at a 90 degree angle to the first. The resulting reception pattern is known as antenna “diversity.” The diagram should tell you everything you need to know. There is no magic here. 2.4 GHz does not “need” a second antenna, but it sure can’t hurt to have one.

Some receivers, like the Spektrum 9300 in #1 above, have a second satellite receiver. This is known as receiver redundancy. We already have examples of redundancy in some RC applications.



2. Antenna reception patterns.

Larger RC aircraft with a backup battery have redundant electrical systems. Aircraft with two engines have redundant power systems.

Another example of redundancy is easily demonstrated with bridge design. Image #3 shows a bridge without redundancy. Remove a single leg and the bridge fails. When non-redundant systems fail, they fail catastrophically. The second bridge, #4, is an example of redundancy in the bridge supports. You can remove a leg or several sections of the bridge and the bridge continues to perform its task. Like the bridge in #3, 2.4 GHz does not “need” a second receiver. Like the bridge in #4, it can’t hurt. Similar to antenna diversity, there is no magic.

Combine antenna diversity with receiver redundancy and you have the best of both worlds. If your system allows for receiver redundancy you might as well use it.


Antenna Construction

With 2.4 GHz, the demands placed on antenna construction have also changed. A 72 MHz antenna could be made from just about any electrical conducting material as long as it was over



3. Above, no redundancy. 4. Below, redundant system.





5. Co-axial cable construction.

two feet long. 2.4 GHz antennas are a different animal and are only 31 millimeters long. Today's longer 2.4 GHz antennas use a co-axial cable that allows you to move the antenna (the final 31 mm) away from the receiver.

Co-axial cable has a core wire surrounded by an insulator and a wire mesh jacket. (See #5 above) The wire mesh jacket shields the internal wire from radio signals. To work as a 2.4 GHz antenna, the final 31 mm of the wire mesh jacket have been stripped off the co-axial cable. This is the portion of the co-axial cable that actually "sees" the signal. You can shorten or lengthen the cable as long as the unshielded portion of the antenna remains 31 mm long. **Do not** clip off any of the unshielded portion of the antenna. If you do, it will significantly reduce the antenna sensitivity, reducing its range.

A radio signal will weaken as it travels from the unshielded 31 mm, down the core wire to the receiver. Obviously, you want to keep this loss to a minimum. Generally, the smaller the wire and the longer the length, the greater the signal loss. Signal loss also occurs at connectors, both crimped and soldered. A kink in a co-axial cable can significantly contribute to signal loss.

One of the factors that determine the quality of a co-axial cable is the distance between the center wire and the outer shielding. Imagine the center core wire traveling down the center tube from a paper towel roll. As long as the tube is kept straight, the tube maintains a common diameter and the imaginary wire running up the middle maintains a common distance from the shielding. Now bend the tube until it kinks and the tube flattens at that location. The distance between the inner core wire and the outer shielding is significantly reduced. A kink causing this sudden loss of thickness will increase signal loss of the antenna at that location.

The Corrupted Packet

Now that we have some basics down, let's shift back to the antennas in my aircraft. All of this data was collected from "2.4 GHz ready" fuselages and does not discuss antennas that exit the aircraft such as whiskers. I have permanently installed data loggers in each of my sailplanes and have kept notes of what works best. The data logger works because of the nature of digital communications.

Today's digital data is not sent in a continuous coherent stream. It is broken into "packets." The transmitter is only sending packets about 10 percent of the time. Yes, there are gaps between the packets. In addition, some of the packets may not be received or might be corrupted. In a direct sequencing system, corrupted packets are typically caused by something blanketing the antenna. The data logger records these corrupted packets as "fades." Recording corrupted packets for an individual antenna is useful in determining best antenna placement.

Unfortunately, this technique does not work with true frequency hopping systems. In addition to corrupted packets from a blanketed antenna, when a frequency hopping system hops to a channel that is already occupied, the packet is assumed to be corrupted and will be rejected. Because frequency hopping systems reject these packets as part of their interference avoidance architecture, it is impossible to tell if the packet loss is from a blanketed antenna or a signal collision during a hop. If the 2.4 GHz spectrum is crowded, the packets rejected from signal collisions may be very high compared to packets lost from antenna blanketing. The short version is that the data here would not help in antenna location.

The wireless transmission systems used by all the RC manufacturers expect to have corrupted data packets. So many packets are transmitted in a very short period of time that random missing packets have no impact on the quality of the data stream necessary for flying RC aircraft. Because of the nature of the data stream required for RC, human reflexes and a bunch of other stuff, random corrupted or missing packets will not be noticed by the pilot in the field. If you are flying on 2.4 GHz now, you are suffering from corrupted packets and don't even know it.

In an effort to combat corrupted packets, every manufacturer sends redundant

multiple copies of the same packet to provide the most coherent data stream possible (another example of redundancy). If the system rejects a packet, the duplicates will usually make up for it. Keep in mind that if the packets quit coming (transmitter died or you flew your aircraft behind a metal building) or a large number of simultaneous packets are lost or corrupted (antenna shielded by carbon or a big metal engine block) you might have issues.

In addition to reducing the impact of random corrupted packets by sending redundant packets, the significance of lost data packets is further reduced with antenna diversity and receiver redundancy. With proper antenna placement, a corrupted data packet received by antenna A may not be corrupted on antenna B. The impact is reduced again when receiver A may receive a corrupted packet, but receiver B does not.

The benefit of diversity and redundancy is demonstrated by the data logger. The data logger records corrupted packets as “fades” for each antenna. If all of the antennas suffer a fade at the same time, the data logger records it as a “frame loss.” It is not unusual to have hundreds of fades in a flight and no frame losses. Diversity and redundancy work.

Before some of you start jumping off a cliff, a corrupted data packet will not cause your servos to imitate a Mexican

jumping bean. Corrupted packets are simply rejected at the front end of the receiver and never make it to the servo amplifier. The receiver just holds the control surface in the last position until a new uncorrupted packet is received. Therefore, they have no impact on servo movement. ***Corrupted packets cannot cause an un-commanded servo movement in a 2.4 GHz system.***

The short version is that every 2.4 GHz system will have corrupted data packets. A perfect data stream is impossible; but at the same time, a perfect data stream is not required.

Antenna Location

By comparing the number of corrupted packets before and after antenna relocation, the best location and attitude of an antenna can be found. Record the information from different antenna locations, repeat the results in several different aircraft and a picture begins to develop. Share the information with others that can confirm the results in their aircraft and the picture becomes more clear.

Surprisingly, little thought is given to antenna placement when today's pilots assemble their airplane. Typically, where to install the antennas is left until the very end. And then the thought process is often limited to, “Ah, I'll just cram it in there.” Unfortunately, the data indicates that haphazard antenna placement will

significantly increase the number of fades.

To minimize corrupted packets, design your antenna placement with the following five rules in mind.

1. Place the antenna (the last 31 mm) in an area where it will not be blocked by carbon, batteries, electrical wires or metal or carbon pushrods.
2. Electrically isolate the antenna from all other wires, metal or carbon.
3. Orientate the antenna to achieve antenna diversity.
4. If antennas are too close, something blocking one antenna may also block the other. Install diverse antennas as far apart as possible.
5. Do not kink any antenna cables.

If you are a sailplane guy, you are probably thinking, “How are you gonna achieve all that in a lawn dart fuselage?” It’s a good question. By using the plastic tubes that come with the Spektrum 9300 receiver (or some other RF transparent tube), you can satisfy all five design rules.

The effect of the tubes on fade count was discovered accidentally. Because the receiver is usually the last item squashed into the fuselage, antennas are kinked, pulled, stretched and wrapped around servo wires and batteries. It isn’t pretty. I realized that by installing the tubes before the servo tray and battery are installed, the antenna would wind up where I want

it, not where I forced to stuff it. Later, simply slide the antenna into the tube just before the receiver is rammed into place. A little time prepping and installing the tubes makes the final installation easy.

There is something more about the tubes. I wish I knew why it happens. Simply placing the antenna inside the tubes results in a reduction of fades no matter where you place the tube inside the fuselage. (This is true in messy sailplane installations where antenna are rubbing the inside of the fuse, servos and the wiring harness, and may not have the same impact where the install has much more room.)

Now take the advantage of the tube and improve it again by installing the tubes in a way that each of the design rules are satisfied. After installing four tubes in a way that satisfied all five of the rules, the fade count on a ten minute flight dropped from the mid to high three figures to the low double or single digits. Although this reduction in fades seems high, there was no change in the frame losses. It was at zero before and stayed that way. The pilot did not notice any change.

Even though the system functioned the same, personally, I want as few fades as possible in my aircraft.

The Installation

Image #6 is the antennas install in a new Tragi 801 also known as the “Cluster.” The red and blue indicate the antenna

tube. The red portion is the last 31 mm of the antenna, the part that “sees” the signal. All five of the design rules are satisfied.




1. The antennas are not blocked.
2. The tube electrically isolates the antenna from all wires.
3. The antenna orientation is diverse.
4. The antennas are as far apart as possible.
5. None of the antenna cables are kinked.

The planning for this install started with looking at the aircraft in flight. If you want to put all the antennas in the nose, start by figuring out what is the last portion of the nose of the fuselage you see as the aircraft turns away just before it is blanked by the wing. Then determine what is the first part of the fuse you see as the nose appears from behind the wing. (This same analysis works for any type or aircraft.) The antenna in the bottom of the fuse just ahead of the wing is the last seen. The antenna in the top of the fuse just ahead of the canopy is the first. These two antennas get the most attention in my installation.

This aircraft uses a 2100mAh LiFe battery so the space ahead of the battery was open and became the location of the third antenna. The last antenna was at 90 degrees to the one in the nose and was mounted between the servos. (See #7 and #8, next page.)

TRAGI 801 INSTALL

- Blue and red lines indicate tube
- Red is antenna
- All 31 mm of the antenna is enclosed in a tube




6. Above, Tragi 800 install. 7., 8, and 9. From top on right, showing tube installations.

As you can see, some of the tube shapes are fairly complicated. Fortunately, it is simple to pre-shape the tubes by sliding them over a bent piece of piano wire, heating the tube up, cool it in water and then slide the tube off the wire. (See #10, next page.) Cooling in the water is VERY important. The water provides a lubricant so the tube slides off the wire more easily. You wind up with a compound curved tube that fits and only need to be tack glued in place. Make sure you rough up the outside of the tube or wrap masking tape around the tube where you expect to use glue.

The tube in #9, right, and #11, next page, runs under the edge of the servo tray and then down and over the bottom of the fuselage. A similar tube was formed for the






10. Tube shaping.

antenna that curves around the upper fuselage just ahead of the canopy. Notice how there are no kinks in the antenna wire because of the tube.

As you can see, planning where to place the tubes has to be done long before cramming the receiver into the fuselage. The result is a much better installation and in the case of 2.4 GHz, fewer corrupted packets.

In Closing

I have seen installations of all the 2.4 GHz brands. In one, the antennas were wire tied to the wiring harness (I was dumbfounded), in another they were wrapped around the receiver and held in place with a rubber band (I was speechless.), and in many, the antenna is just stuffed into all the other wires (You have got to be kidding me!). Amazingly, they



11. Curved antenna tube under servo tray.

worked OK. I think it just shows how well the 2.4 GHz systems work even when the pilot tries to kill his aircraft with a poor installation. Although sloppy installations seem to work, do you really want to decrease your signal strength, reduce your sensitivity and limit your range?

This exercise confirms that lost packets at a particular antenna (fades) are covered by the second antenna in nearly every case. A typical flight results in many fades on an individual antenna but in few or no frame losses at the receiver. In addition, the number of fades can be further reduced with good antenna placement. The fewer fades, the less likely you will have a frame loss.

Try this in your own aircraft and let me know what you think. Send me an email at duworm@aol.com.



How High RT Altimeter

Winged Shadow Systems has redesigned its popular How High Altimeter. The new How High RT features a wider altitude range (35 to 9999 feet above ground level), faster sampling, lower minimum voltage, and an improved capture-on-command function. A board-mounted connector allows plug-in cable changes for a wide range of applications.

In its most common application, the How High RT simply plugs into any R/C receiver (like a servo) and reports peak altitude after each flight using a series of light flashes. No computer or additional equipment is needed. Like the original, it offers one-foot resolution, automatically adjusts for field elevation, and can report in feet or meters.

A significant new feature is the real-time output capability. The How High RT is plug-in compatible with the Hitec Aurora 9 and Spektrum DX8 telemetry systems. Live in-flight altitude is presented on the transmitter display when the altimeter is used with these radios. For do-it-yourself logging and telemetry systems, a fully documented serial output mode is also provided.

An available clip-on battery board, the Smart Bat, transforms the How High RT into a self-powered unit weighing

just slightly more than a nickel coin (5.2 grams) -- including the battery. This self-contained combination is ideal for free-flight models, rockets, falconry, or plane-to-plane portability.


The new How High RT is fully compatible with the optional See How display for convenient post-flight altitude viewing. When used with the display, the altimeter can capture up to nine altitudes on command from any R/C transmitter.

All products are designed and manufactured in the USA. Complete details (including free instruction sheet downloads) are available at www.WingedShadow.com

How High RT Altimeter, \$39.90
Smart Bat Battery Board, \$9.90
See How Display, \$34.90

Winged Shadow Systems
P.O. Box 432
Streamwood IL 60107
USA

(630)837-6553
E-mail: support@wingedshadow.com
Web site: www.WingedShadow.com



Thermal Structure

Ingo Renner from *West Wind*


Reprinted from *Free Flight/Vol Libre* 2/98 Apr/May, pp. 12 & 17

This article comes from a series of lectures given at the Gawler club in Australia called "Flying Further and Faster." The basis for the lecture series began in 1971 when Helmut Reichmann, then World Standard Class champion, participated in the 1970-71 Australian Nationals at Benalla. The lectures include work and input from a variety of sources, including Ingo Renner (a four times World champion).

Thermals are like fingerprints. They are all different but at the same time have sufficient common features to be all in the same class of events. Thus if we draw the structure of one thermal then it will be unique. However there will be sufficient common features to say that most other thermals will be similar.

The structure of the thermal illustrated was established by simultaneously flying a number of aircraft equipped with recording instruments through the thermal. As well as up and down current strengths, temperature and humidity measurements were taken.

The thermal shown had one core. Other thermals observed in the same way were found to have many. Some had as many as twelve! The form was slightly asymmetric. This was attributed to the effects of wind. Stronger winds tended to move the weaker lift surrounding the core more than the core; that is, the core will shift to the upwind side of the thermal as a whole.



The superadiabatic layer near the ground is usually 100 to 300 feet high. On very hot days this may extend up to 600 or 800 feet. This area is very chaotic. In this layer the thermal is made of gusts and is not organized into a steady stream which it becomes higher up.

At low heights (below 1000 feet), there is a strong inflow of air which will drift the sailplane into the thermal. Very little centering action is needed by the pilot and the glider will be drawn into the strong core. Once the sailplane gets into the lift for more than a quarter of the turn, simply keep it turning and let the inflow help you into the thermal core.

The core will have the same strength all the way up and generally will be of constant diameter. Average diameters are 500 to 600 metres. A bank angle of 40° is necessary to keep the sailplane in the core. 45° may be needed for 15m sailplanes and 50° for Open class. With adequate bank and correct position, the circle can be completely inside the core.

A sailplane circling at a 40° bank angle and 46 knots will make a turn of 136 metres diameter. 45° is needed to achieve the same diameter at 51 knots. If the speed is increased due to a higher wing loading then the bank angle must be increased to achieve the same size circle. The thermal tends to be taken at all levels at the same time. That is, if there are many sailplanes in the thermal, they will all leave about the same time regardless of height.

The sink area around the thermal at the levels where it is organized is quite strong.

Sinking air spreading out from the top spreads over a large area and is relatively weak.

There is a wind shear and turbulence in the top section near the inversion layer.

Temperature measurements indicate that by half the height of the thermal the temperature has equalized to that of the outside air. That is, theoretically the thermal should stop!

It does not do so because the mass of moving air has considerable inertia. A thermal column 200 metres across going to 6500 feet will contain over 80,000 tons of moving air! Such a mass cannot stop or change direction quickly. We can conclude from this that the strength of the thermal is more closely related to the height that it goes to rather than other possible factors. The table of thermal strength compiled by Mike Hancy in 1973 based on likely height and temperature has shown a good correlation with results.

The cross-section of the thermal indicates that the sailplane will pass through two distinct areas of turbulence before encountering the core. The first between the more or less neutral air and the strong sink surrounding the thermal should alert the pilot that a thermal is near. The speed director will indicate to fly faster in this sink. The feel of the sailplane is very important to the pilot. Thermal structure from page 12 in this situation. If it feels appropriate the pilot should ignore the speed director.

As the sink area is comparatively narrow by the time the sailplane has accelerated, it will have passed through the sink and the second area of turbulence and into the weak lift surrounding the core.

Horizontal gusts in this area may also complicate the indications showing lift or sink that isn't there! A gust filter in your vario system will help keep it honest. Once in this area it is advantageous to be at good speed — about 5 to 10 knots faster than the usual circling speed (for most sailplanes 55 to 65 knots depending on their wing loading). This allows good aileron response to roll into a turn the moment a decision is made. If the sailplane is not slowed to this speed, then many good thermals will be missed altogether, as the sailplane will have passed through entirely before any good indication shows on the vario.

It needs practice to develop the skill and anticipation necessary. A good, well set up speed director on medium response should

have indicated to slow the sailplane to these speeds. It is necessary to respond to speed director 'up' indications much faster than "down" indications. The feel of the sailplane on coming into the weak lift area should be the best guide to there being a likely core sufficiently close to catch.

While the thermal core is substantially vertical, many factors will cause it to snake about with height. This is similar to the wobbling of a tethered balloon. Wind shears may even break the thermal into two. Generally a strong core will punch straight through most wind shears. Because of this snaking, it is necessary to continually work to keep the sailplane in the best part of the core. Pilots who do that well consistently hold strong lift right up into the top neck of the thermal. They may even get lifted into a thermal dome well above the general inversion layer. From this position an excellent performance can be obtained until the sailplane sinks into the thermal layer again.

If, when you are near the top of the core, the lift becomes irregular, but still with very strong gusts, it is better to leave it than persist. By staying on you will find that the average becomes only half of what it previously was and thus you have been working a thermal that you would not have stopped for at some lower height! Much time (and speed) can be lost in this way.

It is best to try to stay above half the convection height. Thermals are well established by that level and easiest to work. Also, at height there is no anxiety about landing out, so full concentration can be applied to making the best decisions, flying efficiently and working thermals effectively. Set the speed director to that which you are happy to take when low, that is, 2000 feet. But take a good thermal at any height — it is a mistake to ignore thermals until you are at the lower part of your height band!



Constant-Delta Normalized Landing Method

Josh Glaab, louis.j.glaab@nasa.gov

There are two main methods for scoring landings. One tacks-on the landing scores after the flight times are normalized and the other adds the flight times to the landing points before normalizing scores. The benefit of the non-normalized landing method is that landings are always worth the same round score no matter how long the flight was. But there are two drawbacks to the non-normalized landing method. One drawback is that the landings do not account for large changes in landing conditions, due to ground-thermals, or wind-changes, etc. While it is true that all have to land a landing each round, that round could span 60 minutes and conditions can change significantly during that time. We normalize flight times to accommodate for changes in thermal activity (due to changes in the atmosphere), it would seem reasonable to do the same for landings.

Another drawback of the non-normalized landings is that you have to hit a landing to maximize your round score. No matter how many minutes a flyer may beat the rest of his flight group, if the winning pilot does not score well in the landing circle, his score will not be near the max round score. Another scenario is when nobody in a group can get back to the landing due to scratching out time in poor conditions. Even though the pilot that flew the longest, and landed on the field, was the best pilot for that round, his round score would still be penalized by not having any landing points. As a result, he will lose points to other pilots that do hit the landing in other flight groups that have good thermal conditions and easy landing approaches. Non-normalized landings also preclude a scenario where a pilot could decide to continue to work very light lift, and sacrifice a landing attempt, to have a maximum round score. In this sense, a pilot could out-soar the competition and overcome other pilot's landing points. This would be more in-line with a "soaring" event and make it less of a "landing" contest.

One drawback of the normalized landing task is that the landings have a variable effect on your round score. Note that your round

score is what results from whatever normalization/addition process is being used for each round. For example, let's consider two pilots. Pilot A gets a perfect 10:00 flight and a 99 point measured landing (out of 100 points). Pilot B also gets a 10:00 flight, but completely misses his landing. Note here that the term "measured landing" is what the pilot reads on the spot landing tape. The result would be that Pilot A gets a 1000 point (max score) round, and Pilot B would get an 858 round. The difference between the two round scores would be 142 points! That is 42 points more than a 100 point customary non-normalized landing score. In the non-normalized scoring method, Pilot A would get a 1099 and Pilot B would get a 1000. It gets even more severe for short flights. For this example, consider a 5-minute flight time for both Pilots. Pilot A would still get a 1,000 point score, but Pilot B would get a 752. The landing that Pilot A made would be worth a 248 point delta compared to Pilot B!

What is required is a landing task that would provide an order of magnitude improvement in the time-frame used for comparing landings. As stated previously, the current time-frame used to consider all landing conditions constant is the length of time that the round takes to complete (app 60 minutes). As a result, this requirement specifies that the prevailing landing conditions be reflected in the landing scores at least every 6 minutes, which would account for frontal-type wind-shifts. Two orders of magnitude would be desirable (every .6 minutes) to attempt to account for ground thermals. In addition, another requirement is to enable a pilot to maximize his round score while not getting ANY landing points and out-soar the competition.

My proposal would change the normalized landing tasks to effectively replicate a non-normalized landing result for pilots who have identical flight times. It accomplishes this by scaling the landing score by the flight time by a specified amount. Pilots would still read the landing tape (measured landing), as they do now and report that score, but the resulting landing points would be calculated based on the flight time and the amount of round score that a CD wants to allocate for landings. The resulting landing points would

then be added to the flight time for the normalization process to calculate round scores.

To explain this more, consider a case where a CD wants to have the maximum round scores due to landings be worth 100 round points. In this case, the difference in round scores for two pilots, one with a perfect landing and the other without any landing score, would be 100 round points. To accomplish this, the maximum landing score would be calculated by multiplying flight time by 6.666/minute (or .11111/sec). For a 1 minute flight, the landing would only be worth 6.6666 points. For a 10 minute flight, the landing would be worth 66.666 points. For another example, a 5 minute flight would subsequently have a max landing score of 33.333 points. The maximum landing score would then be multiplied by the ratio of the pilots' measured landing to the maximum measured landing (i.e. an 80 pt landing would be 80% of a 100 pt max landing tape).

Consider the example pilots from above again. For a 10-minute flight, Pilot A would get 600 flight points, then he would get $66.666 \times (99/100) = 65.999$ landing points. His combined flight and landing score would be 665.999. Pilot B would get a combined flight and landing score of 600. The normalized round scores would be 1000 for Pilot A, and 900.90 for Pilot B. If Pilot A got a 100-pt measured landing, then Pilot B would get a 900 point round score. This preserves the current 100 point delta provided by the typical 100-pt non-normalized landing method. Proceeding further to consider the 5-minute flight example, Pilot A would get 300 flight points and only $5 \times 6.666 \times 99/100 = 32.9997$ landing points for a combined score of 332.9997. Pilot B would get a combined flight and landing score of 300 (recall Pilot B did not hit the landing tape at all). The resulting round scores would be 1000 for Pilot A and 900.90 for Pilot B. Again, the maximum round score delta due to the landing would be 100 points.

Now consider a situation where Pilot A gets 8 minutes with a 90 point measured landing. In this situation, Pilot A would get $8 \times 6.666 \times 90/100 = 47.9995$ points. Pilot B could maximize his round score by flying for 8 minutes PLUS another 47.9995 (48) seconds. This would allow pilots to win the round AND maximize

their round scores without getting any land points at all! Consider another situation, Pilot A begins his landing approach at the target time and is confronted by a severe ground thermal. Pilot A manages a 4 point (out of 100) measured landing. Everyone else in Pilot A's flight group misses the landing completely due to the turbulent landing conditions. Pilot A would still max his round score. If all of the Pilots in Pilot A's flight group got perfect 10 minute flights, then Pilot A would get a 1000. All the other pilots would 995.57 point round scores.

For another scoring example, consider a case where both pilots get the same landing score, but significantly different flight times. Pilot A gets a 8 minute flight with a 50 pt measured landing, Pilot B gets a 10 minute flight, also with a 50 pt measured landing. Pilot A would get 480 flight points plus $6.6666 \cdot 8 \cdot (50/100) = 26.666$ landing points for a total of 506.666 points. Pilot B would get 600 flight points plus $6.6666 \cdot 10 \cdot (50/100) = 33.333$ landing points for a total of 633.333 points. Pilot A's round score would be 799.999. Pilot B's round score would be 1000. In the non-normalized landing method, Pilot A's round score would be 850, Pilot B's round score would be 1050. It is interesting that the round score delta between the two flights is 200 points for both scoring methods.

Going further, one could consider the landing tape itself. Contest Directors (CDs) could make a landing tape with 50 increments and vary the size of the increments. If you only have 50 increments, the landing score would be computed slightly differently than above. The measured landing score (let's say 48 out of 50 points for a 10 minute flight) would be $10 \cdot 6.6666 \cdot 48/50 = 63.999$ points for a 10-minute flight. CD's could also elect to have landings only worth 50 round score points. In this case, instead of multiplying the flight times by 6.6666 to get the max landing score, 3.3333 would be used instead. Varying the length of the increments could be used to adjust the complexion of the contest. Larger increments would favor precise touch-down times. Shorter increments, would favor landing precision.

One minor drawback to this method is that the resulting round score would not quite be based on 1 point per second flight time

due to the effect flight time has on the landing score. However, considering a 10 minute task, a pilot would have to be more than 6 seconds off of his flight time to see his landing decreased by 1%.

For those of you who have your heads spinning with all this math, please relax and take deep breath. We haven't even broke out the calculus yet (just kidding)! The pilots' objective is still the same as with a non-normalized landing: max the flight, get the landing. But, keep an eye on the competition, if all are down (or almost down) you could maximize your round score without any landing points by flying longer. Since landings are scaled by flight time, the time needed to out-soar your competitor's landing is less for shorter flights. It is also true that you don't have to really stress over a landing for a short flight since again, the landings are scaled with flight time. If a CD calls for a small delta in round scores due to landings (say 50 points) the emphasis will be on precise timing and less on precision landings. If a sudden wind shift/ground-thermal hits the group on approach and all have bad landings, recall that the landings are normalized and that your round score will be OK and you are on an equal footing with those who have great landing conditions.

Before the advent of computers being used for scoring, this method would simply not be possible. Now, however, even the most minimum computer would do the calculations in microseconds. To implement effectively, however, the resulting round score sheet needs to be updated to show flight time, measured landing, calculated landing score, combined flight score, and resulting normalized round score to help all get familiar with the system.

The proposal defined herein satisfies the requirements previously stated to account for significant changes in landing conditions and moves the competition away from being a "landing contest" to be more of a "soaring event." I think this method is worth strong consideration as well as test drive, and I know just the contest to do that (HRSF/BRASS on 5/14 and 5/15/2011)!

Please send me your thoughts and comments!



Reprinted from Eastern Soaring League Newsletter, Nov. 2010