

Just a little faster, please

John Cochrane

There comes a time that you want to fly a little bit faster. Maybe you've been to a contest or two and you've seen what amazing speeds the top pilots achieve – and often in surprisingly bad conditions. Maybe you want to go for a badge, or just cover a little more territory in your cross-country flying. We are glued to this sport by obsessive self-improvement, and a little more speed soon becomes the focus of that obsession.

A better pilot or a better glider?

Many pilots think that the key to going faster is to spend a lot of money on new gliders. They don't go to contests because "I won't be competitive in this old thing." In fact, small differences in pilot technique outweigh huge differences in expensive fiberglass. You usually see new gliders at the top of the scoresheet only because great pilots tend to put the money and effort into flying the latest gliders. The top pilots would still win if they had to fly 20 year old gliders.

To see what a little thinking and practicing can do, let's set a goal of cutting down 3 circles per hour. This doesn't seem like much, maybe one circle every other thermal. How many of us do not, three times per hour, take a circle that gains nothing; maybe searching for a thermal that isn't there, indecisive about leaving, or centering poorly? That seems like an achievable goal for a season's practice.

Each circle takes about 25 seconds; 3 times 25 seconds divided by an hour is 2 percent or 20 contest points. In Figure 1, I used the US handicaps to plot performance against cost. *Cutting 3 circles per hour is worth about \$20,000!* It is like moving up one generation in gliders, for free. And given the choice, wouldn't it be a lot more fun to be a better pilot in a worse glider than to be a poor pilot in an expensive glider?




Figure 1. The cost of performance

Speed and Modern MacCready theory

Now, how to go faster? I have spent a lot of time watching fast pilots, listening to them, reading articles by and about them, and trying to understand what they do and what they say they do, which are not always the same thing. Techniques have changed since the classic writings by Moffat, Reichmann, and Byars and Holbrook, and distilling the essence of the Brigiadori is not easy. I'll point out some of the innovations that

I see. I have also updated the classic MacCready theory to take account of the fact that thermals are random and height is limited¹. This mathematical theory seems to accord well with what fast pilots do.

The MacCready value is still the key to in-flight decisions. It's easiest to think about it as, "what is the weakest thermal I would stop for right now?" It is a good discipline to think about this out loud, set a conscious policy. Don't stop for less, and leave as soon as your averager is less than this value.

The MacCready value also determines the cruising speed. If you have decided not to stop for less than 2 m/s thermals, then you should cruise at the 2 m/s MacCready setting.² Yes, this is still true. There is lots of talk about "flying slower than MacCready", which we will think about in detail, but it is a mathematical fact that if you will not stop for less than 2 m/s thermals, you should cruise at a 2 m/s MacCready setting.

More generally, the MacCready value answers the question "how much higher would I have to be in order to finish one minute sooner?" Our game is trading altitude for time, and the MacCready value is the price of altitude in terms of time. My rules above derive from this concept. If it takes 2 meters of height to finish one second sooner (or 120 meters to finish a minute sooner), then you take any thermal greater than 2 m/s, and you spend altitude at the same rate.

That's all fine, *given* the MacCready value, but what is the right MacCready value? What is the relative price of altitude and time at any moment in the flight? How aggressive *should* you be? Now we leave the land of mathematical certainty. This is what that long experience in watching weather and learning what thermals lie ahead tells the experts. But we can work out the answers in some simple and stylized situations, and these parables are useful ways to organize our thinking about the right MacCready value for a real flight.

MacCready.

If you know the strength of the next thermal, and that you can get to it, then this is the MacCready value for the glide to that thermal. If you know that the next thermal will be 2 m/s, then you set the speed ring for 2, and fly the appropriate speed to fly.

Reichmann.

Reichmann refined this calculation. Thermals are often weaker at the top and bottom than in the middle. Reichmann showed that you should use the weaker "initial" thermal strength as the MacCready value for the preceding glide. If you fly a bit faster, you will have to make up your altitude in that weaker lift, not in the booming lift near the top of the thermal.

You should always take any thermal greater than the current MacCready value, and Reichmann applied this idea to the last thermal: stay in the last thermal until it weakens so much that it equals the initial climb of the next thermal. Thus, Reichmann's rule: Initial climb in the next thermal = MacCready setting = final climb in the last thermal.

Random lift and finite altitude.

¹ See "MacCready Theory with Uncertain Lift and Limited Altitude," *Technical Soaring* 23(3) (July 1999) 88-96, also available on my webpage at <http://faculty.chicagobooth.edu/john.cochrane/soaring/index.htm>. Also see Robert Almgren and Agnes Tourin, "[Optimal soaring with Hamilton-Jacobi-Bellman equations](http://www.courant.nyu.edu/~almgren/)", available at <http://www.courant.nyu.edu/~almgren/>. Anthony Edwards points out that what we call "MacCready theory" existed long before MacCready, who invented the ring and won a world championship with it, and the extensions to uncertain lift and limited altitude also have a long history. See the "Armchair Pilot" in the June/July 1980 *Sailplane and Gliding*.

² The 2 m/s MacCready speed is not the speed for 2 m/s sink, it is the speed where an *additional* second of time gained will cost an *additional* two meters of altitude.

These calculations are obviously simplified. Most importantly, we really don't know where the next thermal will be and how strong it will be. We want to know the right MacCready value, given the *chances* of finding thermals of various strengths, and given the altitude in hand to search for them.

Figure 2 presents an answer to this question, flying a discus on a good day in Northern Europe and the Eastern US. I specify that thermals rise to 2000 m. I specify the *probability* of finding thermals shown in Table 1. For example, if you travel 2 km, you have a 20% chance of finding a 0.5 m/s thermal, a 10% chance of finding a 1 m/s thermal, and so on. In 10 km, there is a 90% chance of finding a 0.5 m/s or better thermal, as 61% chance of finding a 1 m/s or better thermal, and so on. There are enough weak thermals that you are pretty sure of staying up. There are a few really good thermals, but you'd better not go barreling around counting on them. Still, you want to adjust your strategy so that if you find one, you can take advantage of it. I solve the dynamic program for maximizing average speed, valuing landouts as they are in competitions.

Thermal Strength m/s	2 km Point	2 km Cumulative probability	10 km Cumulative probability	20 km Cumulative probability
0.5	20	37	90	99
1	10	17	61	84
2	5	7	30	52
3	2	2	10	18

Table 1. Thermal assumptions. "Point" gives the probability of finding each thermal in the first 2 km. "Cumulative" gives the chance of finding a thermal this strong or stronger in the indicated number of Km.




Figure 2. Optimal MacCready value vs. altitude, using the thermals of Table 1, and no waterballast. Dashed line: Discus with no 0.5 m/s thermals.

You can read several rules from Figure 2.

1. *Steadily reduce the MacCready setting as you get lower -- fly more slowly and take weaker thermals.*

The optimal MacCready value rises from less than 0.5 m/s at 200 meters to 2 m/s at 2000 meters. The reason is simple: *range*. If you stop for nothing less than 2 m/s at 300 meters, you are soon going to meet a nice farmer. We knew this of course. Even the earliest explanations of MacCready theory added advice such as “take anything to stay up when you’re low.” If you should “take anything when you’re low” there must be a middle range where you should “take mediocre thermals”, and that is Figure 2’s advice.

2. *Leave weak thermals to go find better lift as you get higher.*

Many books warn that after a low save, it’s important to recharge your psychology and not work your 1 m/s thermal all the way to cloudbase. Figure 2 refines this idea. You might get low, and find a 1.5 m/s thermal. You take it. Around 800 meters, though, you should start getting impatient, leave and try to find something better. You might not, and have to take another 1 or 1.5 m/s thermal, but at 800 meters, you are more likely to find something better than you are likely to have to accept something worse. When you do find that something better, you’ll have the altitude to use it. Cloudbase is the *worst* place you can be. If you run in to a 5 m/s thermal at cloudbase, you can’t take it anywhere!

Conversely, Figure 2 means that if you find a lucky, strong thermal, you should stay in it even as it weakens. A pilot who finds a 3 m/s thermal should be grateful for the gift. If it cools off to 2 m/s, he should still stay in the thermal up to 1800 m. Don’t leave until the thermal you are in is weaker than the next thermal you are likely to find, and the Mc setting you are willing to cruise at to get there.

Many pilots and books describe flying in a “height band” for the day. Together, the last two points dynamically define such a band. If you get less choosy as you get lower, you are more likely to climb up. If you get more choosy as you get higher, you are likely not to do so. The band is not, as is often described, fixed. You don’t, say, climb at 0.1 m/s if you’re a bit lower than you’d like, nor do you bail out of 6 m/s if you reach some “upper band.”

3. *MacCready settings are substantially lower than best climbs.*

In my calculation, the best thermals of the day are 3 m/s. Yet the optimal MacCready setting never goes over 2, and will be more like 1.5 through the typical range of the flight.

The basic principle behind the calculations in Figure 2 is this:

4. *The MacCready value now should be the same as you expect it to be farther ahead.*

If you know you are going to be desperate up ahead, you should start conserving altitude now. Suppose that you are at 900 meters. Looking ahead 10 km, you think there is half a chance you will find a 2 m/s thermal. However, there is half a chance that you will not find any thermal, wind up low, and be quite happy to take 1 m/s. *Your MacCready value now should be 1.5 m/s.* This is a good principle to use in thinking about what MacCready value to set. I used this principle to ask the computer to work back from the finish to find the right MacCready values for any combination of altitude and distance to go.

5. *Weather, pilots and gliders*

The curve in Figure 2 moves around according to the weather, the glider, and the pilot. Obviously, the curve shifts to the left in weak weather and to the right in strong weather. Less obviously,

- The shape of the curve depends on how good thermals are at lower altitudes. If thermals are weaker lower down, you become more conservative sooner, and accept weaker lift to stay in the good band. Thermals tend to be weaker down low in wind, in mountains, at the end of the day, when there is a wind shift with altitude, and when a strong circulation layer develops as with cloudstreets or with strong capping inversions.

- The shape of the curve also depends on how *frequent* thermals are, especially at low altitudes. If the “Discus” curve seems aggressive to you, it is because I have programmed in a quite high chance of staying up, by assuming a 20% chance of finding a 0.5 m/s thermal every 2 km. This opportunity to save the flight and keep going encourages what would otherwise be pretty aggressive low-altitude behavior. If we keep the 1,2, and 3 m/s thermals of Table 1 but eliminate these 0.5 m/s savior thermals, the curve shifts to the less aggressive dashed curve shown in Figure 2. This may explain why the English advocate pressing on at what seems like very low altitudes, while pilots in the western US get panicked at 3000 m. In the western US, you either climb at 5 m/s or you don’t climb at all. Ridges on which to “save” a flight at low altitude can also allow aggressive flying.
- A lower performance glider must fly more conservatively. The calculation for a Ka6 gives MacCready settings about 0.5 m/s lower than the Discus. The Ka6 pilot must stop to take weaker thermals to bridge the gaps between longer thermals, and he must cruise at lower MacCready. Classic MacCready calculations, which assume that everyone will be able to get to the same thermals, understate the advantages of higher performing gliders.
- A less skilled pilot should fly more conservatively, shifting the curve to the left. If you are less skilled than the top pilot is, *you* will increase *your* points by following a more conservative strategy than he follows. Top pilots will find a thermal that you and I will miss. We need to give ourselves a little more room. Following top pilots leads to starting too late, watching them disappear over the horizon, and then struggling home.
- The curve depends on how you feel about landing out. If you want to minimize the probability of landing out, you set the MacCready to zero always. This is really slow. To fly any faster, you must accept some larger probability of landing out. In Figure 2, I valued landouts according to the distance points in contest rules. If you’re flying in a contest that gives more distance credit, fly more aggressively. If your philosophy is “win or landout”, not “maximize expected number of points”, fly more aggressively. If your personal dislike of landing out goes beyond contest points, fly more cautiously, especially as you get low.

Refinements

Centering time

On most flights, it takes at least a couple of turns to center the thermal. A pretty good pilot can start climbing at the thermal’s maximum rate in 4 turns – 2 minutes. The rest of us flog around longer than that. Table 2 shows what 2 minutes of centering time does to the achieved climb rate.

Height Gain (m)	Thermal strength (m/s)				
	0.5	1	2	3	5
250	0.40	0.67	1.0	1.2	1.5
500	0.44	0.81	1.4	1.7	2.3
1000	0.47	0.89	1.6	2.2	3.1
2000	0.48	0.94	1.8	2.5	3.8

Table 2. Achieved climb rate if it takes 2 minutes to center a thermal.

As you can see, a few minutes of centering time has a dramatic effect on achieved climb rates! The effect is larger for *stronger* thermals, and for *smaller* height gains. Managing this centering time is the next crucial piece of flying strategy. For many thermals, the decision to stop doesn’t depend so much on *how strong* you think the thermal is, as *how easy it will be to center*. If you feel the right kind of surges and can roll right in to a 3 m/s thermal for 900 meters, that is better than having to center a 5 m/s thermal for the same height gain.

Many modern flight computers include an average climb for the whole thermal – from the minute you switch in to climb mode or start circling. These “reality meters” are wonderful checks on your enthusiasm. When I bought a flight computer with this feature, I was amazed that what I thought of as a “3 m/s day” was often really a 1.5 m/s day. I felt a lot better about my seemingly wimpy intrathermal speeds.

Centering time affects classic calculations such as Reichmann’s, that presume you know what the next thermal will be like and where it will be. The *lower* of average climb and initial climb (after centering) determines the MacCready setting. The “initial climb” rule considers how much lower you will arrive at the next thermal if you fly a little faster. The “average climb” rule considers how many more thermals you will have to center if you fly a little faster. The actual price of altitude you face – the correct MacCready setting -- is the lower of the two climb rates.

When it takes time to center thermals, it is worth *staying* in a thermal somewhat weaker than you would *stop* for. The curve of Figure 2 breaks into two curves, a weaker “stay” curve and a stronger “stop” curve. The difference is strongest higher up. The speed decision is based on the lower value.

Many pilots follow rules such as “don’t stop unless you can gain at least 300 meters.” Like any rule, this one is meant to be broken, but it contains a grain of truth. It’s worth stopping at any altitude if the thermal is strong enough, and especially if it feels smooth so that you will not have to center it. But to stop in any thermal you must amortize the centering investment in a decently long climb.

Misconceptions

Of course, “fly the MacCready speed” does not mean we chase the vario needle around. Lags in the instrument and the pilot means that most pilots fly relatively constant speeds, unless long stretches of lift or sink ahead are clear. You choose that relatively constant speed based on the MacCready value.

Pilots often criticize MacCready theory, noticing that the exact speed you fly isn’t that crucial. 10 km/h one way or the other will not make a great deal of difference. However, 20 km/h will make a big difference. More importantly, while *gliding* a few km/h too fast or slow won’t make much difference, *choosing thermals* one m/s too low or insisting on thermals one m/s too strong will make a huge difference to your speed. *Deciding when to stop and when to leave thermals, and achieving the best average climb rate, are the most important determinants of cross-country speed.* This decision is as much a part of “MacCready theory” as is the decision of what speed to fly – the MacCready values in Figure 2 apply equally to each decision.

It is a common misconception that you should use MacCready settings that are systematically lower than the worst thermal you would take, in order to get more range. It is a mathematical fact that if you are cruising at a MacCready 1, you will always do better by stopping for a (smooth, easily centered, bottom-to-top average) 2 m/s thermal, at least for a short climb until you can cruise faster.

However, the misconception contains a grain of truth. When you add up the effects of low initial climb rates, centering times, and the fact that the *average* thermal you will climb in is stronger than the *weakest* thermal you would take, the correct MacCready value is a lot less than the peak averager reading in the best thermal of the day that you brag about in the bar after the flight. So, yes, pilots now use MacCready settings much lower than they used to. It does not mean “MacCready theory is dead”, it means “use the right climb rate.”

Course deviations

The MacCready value governs other decisions as well, including course deviations. It is surprising how far off course you should go. For example, by going 30 degrees off course, you have to fly 13% further. If you average 100 km/h, 5 km 30 degrees off course costs you 23 seconds. At a MacCready setting of 1m/s, this is worth it if you gain more than 23 meters. Just about any cloud or haze dome will net you 23 meters. (You don’t have to *gain* 23 meters, you just have to gain 23 meters over the pilot who flies straight.) If it nets you

50 meters, constantly zig zagging 30 degrees off course from cloud to cloud will give you a much better speed than going straight. As an extreme, going 1 km perpendicular to the course line will cost 36 seconds. It's worth it at MacCready 1 if it nets you 36 meters.

If the MacCready value is low, it's worth trading a lot of time for a little altitude by larger course deviations. If the MacCready value is high, time is precious so you should drive straight ahead. Of course in stronger lift you will gain more by flying through thermals, so the two effects can cancel. Pilots at Uvalde, where lift is strong, close together, and well-marked often take as much as 45 degree course deviations to hop from cloud to cloud with little circling. Conversely, since higher performance gliders use higher MacCready settings in the same weather, now we understand why ballasted and open class gliders fly straighter courses.

Final glides

The standard final glide calculation assumes equal lift and sink. How should you approach a final glide given that thermals are random and you might land out? There are two schools of thought on this.

First there is the "start the glide early and low" school. Doug Jacobs has offered this advice, and seems to start his glides one thermal before everyone else. Bill Bartell advises you to start thinking about the final glide when you hit MacCready 0. You can often do better than the still air glide by course deviations and porpoising in thermals. Starting a final glide low also keeps alive the option of stopping in a superb thermal if one comes along. How many of us have struggled to make "final glide" in a 1.5 m/s thermal, only to blunder in to a now useless 3 m/s while bashing home!

Second, there is the "make sure you don't blow the contest by landing out" school. Dick Johnson has offered this advice. If there is lift there is also sink. How many of us have not also set up a nice 30:1 plus 150 meter final glide, only to have it all evaporate and either end up struggling low, or landing on the way home? Being a little more conservative than the standard calculation – say climbing in a 1.5 m/s thermal to a MacCready 2 glide – might cost a minute or so, but it buys good insurance against this kind of disaster.

Who is right? To get a handle on this question, I went back to the computer, and Figure 3 gives the computer's answer. The lines give the optimal MacCready settings for each altitude and with 15, 30, 100, and 275 km left to go on the task. The 275 km out line is the same line shown in Figure 2. The 30 km out and 10 km out lines show how the calculation advises you to fly the final glide.

Start by comparing these lines to the dashed lines, which give the standard still-air calculation. Over 800 meters, the 30 km out line is about 300 m below the corresponding dashed line – it advises you to fly about 300 m below final glide. There is lift that you can use to porpoise in. If you do not find lift, you can still glide in at a lower setting, and, by the assumptions of Table 1, you are nearly certain to find a weak thermal save the flight. This line verifies the advice of the low-and-fast school.

However, the 15 km out line is much more conservative, and only about 50 meters below the still-air calculation. When I eliminate the 0.5 m/s thermals (not shown), that line is right on the still-air calculation, and adding any chance of sink makes it lie above the still-air calculation. At 15 km out, the program trades the slight advantage of a few km/h more speed for a short time against the small probability of a disastrous landout, and advises a cautious final glide. In sum, this calculation balances the two schools of thought: *start final glides aggressively, but finish them conservatively.*




Figure 3. MacCready values on final glide, Discus with thermals as in Table 1. Dashed lines give the standard still-air calculation.

Now, compare the 15 and 30 km lines with the 100 and 275 km lines, and let's make sense of the interesting bow shape. In the bowed region, the 15 and 30 km out lines say to fly more conservatively on final glide than you would be on course. In these situations, the out-on-course MacCready setting would not get you home, but the slightly lower MacCready settings will practically guarantee a glide home if you should not happen to find a thermal along the way. The program trades off the small loss in points from flying slowly for a few km against the scoresheet disaster of landing out if you don't find a thermal ahead.

The calculations are far from the last word, but the curious way they come out make clear the tradeoffs you have to think about. On final glide, you balance large chances of a small speed increase against small chances of a costly landout. Managing this tension correctly wins contests. Final glide strategies are a particularly fertile area for quantitative analysis. As with safety issues, which trade similarly small probabilities, it is hard to learn this balance from personal experience since the disasters are infrequent.

Weather is especially important on final glides. Even the most aggressive pilots take high final glides when they have to go through rain on the way home! The chance of sink is just as important as the chance of lift. You fly more conservatively if the weather is more *uncertain*. (I learned this sharp lesson from Liz Schwenkler when she beat me home on a MacCready 0.5 final glide. "No lift means no sink," she said, and she was right.) Porpoising may be harder down low than when up high, and the presence of weak thermals with which to save the flight are crucial for the low and early strategy.

Finally, this calculation like all final glide calculations assumes there are plenty of fields in which to make a safe, last minute landing should a final glide go wrong. One must be much more conservative if this is not the case, as is often true flying in the U.S. An off-field landing from a final glide is made from very low altitude, without a deviation to look at the fields or a chance to plan the pattern well. Contests are full of

serious crashes from landouts a few km from the airport, or from arriving too slowly at the home field. “I just made it over the fence” is not funny!

Upwind and downwind

We all know we should take upwind turnpoints low and downwind turnpoints high. *How* low, and how high? Again, MacCready values determine the answer. As you approach an upwind turnpoint, the whole curve of Figure 2 shifts to the right, and as you approach a downwind turnpoint it all shifts to the left. As a result, your height band naturally goes down near an upwind turnpoint – unless you happen to find a 5 m/s thermal! See “Upwind and downwind” on my webpage³ for details with a handy chart.

What’s next?

When you learned to follow the towplane, you and your instructor analyzed the task. Then you flew to learn to do in the air things you understood on the ground. By the time you got your license, following the towplane became automatic, and you probably would have trouble explaining how to do it to a beginner.

Cross-country flying works the same way. You start with the basics, thermaling and navigation. This article is about the intermediate stage, getting up to speed on course. You have to think about and analyze these decisions on the ground, and then use your flying time to learn to make them in the air, and then to make them subconsciously. We fly to learn to make in the air decisions that we understand on the ground. This is not easy, and requires dedicated practice. I know this from experience: I write articles on theory, yet from lack of practice I still end each flight with a list of silly decisions to mull over.

Great pilots have made this all automatic. They often have trouble describing what they do as you might have trouble describing how to follow the towplane. They fly thinking about weather, psychology, and contest tactics. Our job is to get to that stage!

For technical types, I have only scratched the surface of what the mathematical technique – dynamic programming – can do to advance the theory of soaring flight. Explicit treatment of centering times, thermals whose strength and character changes with altitude, better thermal models, upwind and downwind turnpoints, objectives other than expected value of contest score, comparing the program solutions to flight recorder data of top pilots’ decisions, and many more questions only await enough wintertime programming to be solved.

³ http://faculty.chicagogsb.edu/john.cochrane/research/Papers/upwind_and_downwind.doc