



The MACC Handicap System

Mike Sayers

Overview of the MACC Handicap	1
Racer Handicap Variability	2
Racer Handicap Averages	2
Expected Variations in Handicap.....	2
MACC Handicap Calculation Methods.....	3
Calculation of ZPT and CF	3
Expected Values of HC Based on Past Results.....	3
Pacesetters	4
Choosing Racers for Least Squares Analysis.....	4
Method #1: Calculate ZPT and CF Using Linear Regression.....	5
Method #2: Calculate ZPT with a Specified CF.....	5
Seasonal Adjustments.....	5
Is the System Working?	5
Changes Made Between Seasons	6
Other Ranking Systems.....	6
NASTAR.....	6
FIS Points	7
Revisions.....	7

MACC racers are classified by sex, ability, and (optionally) age. A range of handicaps defines each ability group, where *handicap* is the percentage that the racer is slower than a reference skier on a reference course. With this concept, fast racers have low handicaps and slow racers have high handicaps. This technical memo describes the concept of the MACC handicap and provides technical details for how it is calculated.

Overview of the MACC Handicap

The handicap statistic (HC) is used to place a MACC racer into an ability group, as described in section 7 of the MACC Rules. There are currently seven ability groups ranging from Open Elite (fastest) to slowest (Class D). Each class is defined by a range of HC values as shown in Table 1.

Table 1. Handicap ranges for racer classes.

HC < 18	15 - 28	25 - 38	35 - 48	45 - 64	60 - 80	75 - 96	HC > 90
Open E Women		E1 Women	E2 Women	A Women	B Women	C Women	D Women
Open E Men	E1 M	E2 M	A Men	B Men	C Men	D Men	

HC is the percent that a racer time is longer than the time of a hypothetical reference racer on a reference type of racecourse. For example, an HC value of 34 means the racer had a time that was 34% longer than the time of the theoretical reference racer on a theoretical reference course.

Each MACC racer has two handicaps: one for slalom and one for giant slalom (GS). The class for the racer is determined by the smaller of the two. For example, if a male racer has a GS HC = 34 and a slalom HC = 50, he races in E2 Men or faster, based on his (lower) GS HC value.

HC values for all racers are recalculated after every race weekend. Any racer whose new minimum HC is lower than the limit for the class is automatically upgraded, as described in the MACC Rules. As described later, the HC value reported for each racer is usually an average of the best six recent results.

Notice that the classes have HC ranges that overlap. This is done to avoid having racers “bounce” between classes. For example, if a racer in E2 Men has a handicap of 37, he normally is not allowed to move to the slower A Men class.

A secondary use of the handicap is to provide a tie-breaker at the end of the season in the event that two or more racers in contention for a World Cup trophy have the same number of World Cup points. When this happens, the racer with the lowest sum of the slalom and GS HC values gets the higher prize (see 10.1.1 in the MACC Rules).

Racer Handicap Variability

The relationship between racer times varies for each run. We account for this variability in several ways.

Racer Handicap Averages

The official MACC HC for each racer for a type of course (slalom or GS) is calculated as an average of the HC values for recent runs on that type of course. It is common to have a few runs that are significantly slower due to a mistake; however, it is not so common to have a few runs that are significantly faster. With this in mind, we throw out a few of the highest handicaps before taking the average, while leaving in the lowest handicaps.

The average handicap reported for each racer is calculated using the best six of the most recent eight timed runs for GS (that is, four two-run races if all had valid finishes), and the best six of the most recent nine runs for slalom. DNF and DSQ results are ignored. Race results are used for the current season and previous two seasons. For racers who do not have eight GS times in the database, the best 3/4 of the results are used to calculate the average. For racers who do not have nine slalom times, the best 2/3 of the existing results are used. If the racer has more than a single run, at least one run is dropped. For example, if a new racer has GS handicaps of 34.9 and 39.2 for the two runs in the GS, the 39.2 would be dropped and the official HC for the racer after the weekend would be 34.9.

Expected Variations in Handicap

The reason for using averages is that the average HC reported for each racer fluctuates less than the individual HC values. However, be aware that these averages also vary from weekend to weekend. New HC results are included, and old HC values outside the range (eight runs for GS, nine for slalom) are removed from the average. If several recent HC values are low, the average will drop noticeably. If the HC for a racer drops enough to trigger an upgrade to a faster class, it is possible that the recent HC values that were low for racer will be followed by some higher HC results. This is expected for a racer who has been in the same class for a long time.

MACC Handicap Calculation Methods

The equation used to calculate racer HC is:

$$HC = 100 \times \frac{(T - ZPT)}{ZPT} \times \frac{1}{CF} \quad (1)$$

where T is the racer time, ZPT is zero-par time (the time of a “zero-par reference racer”), and CF is the course factor. For example, if CF = 1, then a HC value of 34 means the racer had a time that was 34% longer than the time of the theoretical reference racer.

The HC range associated with a class does not depend on the type of race. Our intent is that if a racer skis different courses with a similar level of skill, then a similar HC should be obtained. The course factor CF accounts for spread in times on a course between the fastest and slowest racers. In general, the spread is larger on challenging courses. On an icy slalom course, an intermediate skier might be 100% slower than an expert racer. However, on a flat hill with a straight course set, the difference might only be 50%. In order to obtain similar HC values for the slowest and fastest racers on a course, a higher CF value is used for the more challenging slalom course.

Calculation of ZPT and CF

Equation 1 has two independent coefficients: ZPT and CF. For each course, we calculate values for those two coefficients. Following are some methods that could theoretically be used to get ZPT and CF:

1. Use statistical methods to determine the ZPT and CF values that best match the expected HCs for racers and/or pacesetters on the course.
2. Manually enter a CF value and use statistical methods to determine the ZPT that best matches the expected HCs for racers and/or pacesetters on the course.
3. Manually enter a ZPT value and use statistical methods to determine the CF that best matches the expected HCs for racers and/or pacesetters on the course.
4. Manually enter the ZPT and CF coefficients.

The above methods are ranked by compatibility with old data. Method #1 ensures that the new HC values are as compatible as possible with the values expected from past results; Method #1 is the default method used in MACC. Method #2 is used only when results from Method #1 are unacceptable, as will be described below. Methods #3 and #4 are mentioned for completeness, but do not take into account any information about a new course other than someone’s intuition. They have not been used for any MACC races going back to 1998.

Expected Values of HC Based on Past Results

The HC reported for each racer is a little lower than the mathematical “expected value” for future handicaps. By removing the two or three highest handicaps before taking the average, we are introducing a slight negative bias to the average.

After reviewing the entire season results from the 1998 and 1999 seasons, we determined that the negative bias could be canceled with adjustments proportional to the variation in HC values for the six runs used to calculate the average. As a result, the expected HC for both GS and slalom is the average HC plus 0.3σ , where σ is the standard deviation of recent HC values. Prior to the 2019 season, separate

multipliers were used for GS and slalom; in preparation for the 2019 season, the slalom number was changed to 0.3 to match the number used for GS based on looking at results for the past few years in which slalom handicaps had drifted up relative to GS handicaps.

The adjustment is not made in the reported HC averages because the concept of taking an average of the six best recent HC results is easy to understand and easy for racers to calculate if they are so motivated. However, the adjustments are used when calculating new handicaps, as described later.

Pacesetters

We use pacesetters to make sure the ZPT and CF values for the course will be used to predict HC values that are valid for racers outside the current classes. Pacesetters are supposed to have the qualities that:

1. They are relatively consistent in terms of their own HC.
2. They have handicaps that are outside the range of the racers on the course. With the current two-course scheduling, pacesetters for Courses 2 are racers from Course 1. Sometimes pacesetters from Course 2 are scheduled to paceset for Course 1.

The pacesetter results primarily affect the CF, because they define the difference between the average results for the racers on a course and racers who are significantly faster or slower. For example, if a racer on course #2 beats the average time by 3 seconds, should that run be reported as a strike? (A strike is when the racer HC is lower than the limit for the class.) If we have no data except for the times of the course #2 racers, we can make a guess. If we also have pacesetter times from an Elite racer, then we can make a much better guess. If the time is closer to the pacesetter than other racers, then there's a low handicap. However, if the time is still much slower than the pacesetter, then it was just a technically demanding course and the handicap will only be a little lower than normal.

If we don't have pacesetters for the course, then the CF is determined only by the histories of the racers on the course. Although the calculations are theoretically valid, we lose some independent quality control.

Note MACC went from a schedule with three courses to two courses near the end of the 2018 season. Our policy for putting pacesetters on Course 1 is not full settled at the time this revision was posted.

Choosing Racers for Least Squares Analysis

Both pacesetter and racer times are used to calculate ZPT and CF, based on expected values of HC from past results. However, not all racers and pacesetter results are used. Here are some considerations:

1. The software looks at the consistency of HCs for past races. Only racers with consistent HCs (low sigma values) are even considered. We don't care if a racer is fast or slow—just consistent.
2. Because there are many more regular racers than pacesetters, each pacesetter is treated statistically the same as 10 regular racers.
3. A preliminary analysis is done and HCs are calculated for all racers and pacesetters. We compare the calculated HCs to the past results and throw out some of the data:

- Pacesetter results in which the HC is noticeably higher than past averages are thrown out. We try to have three pacesetters for each course, with the expectation that at least two good pacesetter times will be obtained for each run.
 - Racer results in which the HC differs too much from past results are thrown out. This is to eliminate exceptionally bad runs (falls, climbs, etc.).
4. A second analysis is done using only the times accepted after step #3. This typically involves 50% to 80% of the timed runs.

Overall, this approach is designed to maintain compatibility with past results to the maximum extent possible.

Method #1: Calculate ZPT and CF Using Linear Regression

Method #1 is the default used in the MACC software. The statistical method is to calculate a least-squares linear regression of expected HC versus race time T. That is, for each racer or pacesetter, the difference between the HC calculated for the new run and the HC expected from past results is calculated, and ZPT and CF values are calculated that minimize the squared HC differences for all of the racers and pacesetters that are being considered.

The result is an equation that predicts HC values for all racers on the course as a function of their race times (that is, equation 1, with numerical values for ZPT and CF).

A possible outcome of linear regression is that new results do not agree well with expected results. That occurs when the order of racers (fastest to slowest) varies from past results. In the worst case, the mathematical “best fit” is that everyone gets nearly the same handicap. When this happens, the calculated CF will have a high numerical value.

Method #2: Calculate ZPT with a Specified CF

The MACC software supports methods #1 and #2. We try #1, and if the CF is within the range expected for the type of race (0.8 — 1.2 for slalom, 0.5 — 0.9 for GS), then that equation is used. If not, then the software automatically uses method #2 with the CF appropriate for the type of course. Usually the CF is set to the upper limit (1.2 for slalom or 0.9 for GS) and ZPT is calculated.

Seasonal Adjustments

Ultimately, the handicap system exists for just one reason: to determine when racers should change class. The calculation of new handicaps depends completely on the existing handicap statistics of participating racers and pacesetters. The handicap values are updated automatically during the season.

Is the System Working?

Here are some things to look for:

1. Are strikes and upgrades being awarded at all? If we haven't upgraded some racers by mid-season, there's a problem.
2. When a racer is upgraded into a new class, how does he or she do?

- a. We expect that the racer will not be competitive right away in the new class, but will at least have HCs near the cut-off.
 - b. If the racer is way out of competition, with HCs much higher than the range for the class, then the upgrade should not have been made and our system needs improvement.
 - c. If the racer is competitive in the new class, then the racer should not have been in the slower class.
3. Are racers being upgraded for both GS and slalom? If nearly all upgrades are for GS, then we need to adjust our slalom handicaps.

Changes Made Between Seasons

Often, no adjustments are made to the database between seasons, other than to copy the data from the old season database into the new season database, without the data from two seasons ago. However, if problems are identified, the summer between seasons is the time that they can be fixed.

Several times, we applied a transform to all existing HC values for use in the next season. We have done this several times in the past 15 years to lower slalom handicaps such that they better match the GS handicaps.

Other Ranking Systems

The handicap concept has been used in MACC and other racing programs for decades. NASTAR (originally an acronym for NAtional STandard Race) uses the same basic handicap concept as MACC. USSA, FIS, and the World Cup use a point system that indicates how much a racer is slower than a reference racer on a specific type of course.

NASTAR

NASTAR uses handicaps to provide a measure of performance that is intended to be independent of where the course was set up. A NASTAR handicap is calculated with the same formula as MACC, but with the course factor always set to 1.0. The ZPT value is obtained only from the best time of a pacesetter on the course, and is calculated simply as:

$$ZPT = 100 \frac{T_{ps}}{100 + HC_{ps}} \quad (2)$$

where T_{ps} is the pacesetter time and HC_{ps} is the pacesetter handicap. The NASTAR definition of ZPT does not take into account effects of course difficulty, nor does it take into account the times of any racers on the course other than the pacesetter.

NASTAR distinguishes between regular courses and “NASTAR Slalom” courses. The same methods are used to calculate handicap. However, the handicap ranges associated with medals are adjusted slightly so some racers are allowed a higher handicap for a given type of medal.

Handicaps for NASTAR pacesetters are reset each season. Several sessions are held where pacesetters from different regions all run on the same course. The pacesetter that defines the zero par time for the season is the fastest racer participating at a pre-season session (last season, this was held at Copper

Mountain in November). Other sessions are held later for other pacesetters to run on the same courses as pacesetters were at the national pre-season session.

FIS Points

FIS races, including World Cup, rank all internationally registered racers with points that are calculated separately for slalom, GS, super G, downhill, and combined. The points for the World Cup and high-level FIS races are similar in concept to a handicap, but they are scaled differently. First, the points are calculated per race, rather than per run. If the racer does not finish both runs in a two-run race, then no points are calculated for that racer. Second, the racer points are scaled so the best in the world has 0.00 points, and the 31st racer has 6.00 points. Regardless of the event, if a racer has less than 6.00 points, he or she is in the top 30 in the world.

The calculation methods for FIS races are much more complicated than MACC, because they are designed to work with the points of the fastest racers who show up for a specific race. FIS points are recalculated every few weeks during the season. For example, in April 2016 there had been 11 point lists published on the FIS web site for the 2015-2016 season. Each time the racer points are recalculated, the races that are considered cover the past 13 months. Anything older than that is not included.

FIS points are used to qualify racers for high-level races, and to determine seeding within those races. In general, the emphasis is on the fastest results from the fastest racers. Points for an individual racer can fluctuate quite a bit on separate races, especially for the slower racers.

Revisions

The first version of this document was posted November 2016.

The next revision, posted January 2019, fixed some typos and removed references to rules that have changed.