

Use of Genetic Algorithms with Multiple Metrics Aimed at the Optimization of Automotive Suspension Systems

Scott A. Mitchell, Stephen Smith, Alberto Damiano, Joel Durgavich, Rosalyn MacCracken
School of Computational Sciences, George Mason University, Fairfax VA

Copyright © 2004 Society of Automotive Engineers, Inc.

ABSTRACT

Suspension models are highly multivariate and require a nonlinear system to model the movements and interaction of the parameters within the suspension system. Multiple metrics must be considered to determine an optimal result.

This paper describes a system for the use of a Genetic Algorithm for the optimization of automotive suspension geometries, a description of the suspension model, and the scoring mechanism. The results of this model evaluate the impact of multiple independent metrics. A combined objective function score is determined with the assistance of a user selectable weighting of metrics. The optimization algorithm is also compared to a discrete grid search.

INTRODUCTION

A suspension system must be modeled to determine the suitability of the design. In doing this, the designer/engineer can diagnose where there are deficiencies in the overall system. A simple 2-dimensional model, called Suspen, was developed by Scott Mitchell in 1991 to aid in the design of the Virginia Tech Formula SAE cars.

After the initial development of the Suspen model, certain limitations in the system became apparent. It was decided to improve on the original model and correct these deficiencies. A new version of the Suspen model was developed from the ground up. The model was created under the specifications of the Foundations of Computational Sciences course offered at George Mason University.

This paper focuses on the current Suspen model and the creation of a new system for the optimization of suspension geometries with a genetic optimization algorithm. A new scoring technique has also been created to evaluate the effects of the inputs into the model.

SUSPEN MODEL

In choosing a suspension design, it was necessary to select a system that needs a high degree of optimization and is widely used. The expected users of the software would be racers, in particular Formula SAE teams, where SLA (Short Long Arm) suspensions are used almost exclusively.

This mathematical model projects the 3D suspension system into two dimensions, the Front View. This is computationally analogous to Carroll Smith's "Paper Dolls" [3]. This projection retains a number of very useful metrics, such as camber, track width, roll center (RC) location, as well as other parameters. Lengths used here are in inches, however the model itself can use any length unit as long as they are consistent.

All members are assumed to be rigid. This includes the tire, which is modeled as a solid member on its centerline. Off the centerline, it is only drawn for visualization purposes. This simplifies the analysis of tire location. Physical tires are not rigid members, as load is applied or relieved the effective radius of the tire changes. However, individual wheel bumps are modeled so that tire deflection can be modeled in a simplified manner.

The Suspen model uses four inputs to drive the simulation. The chassis position is defined by bump/droop and roll measurements. Each wheel can be individually bumped as well. The suspension is generally modeled through a sequence of these positions. At each position an *instance* is generated. This instance defines suspension locations for that position. Camber, Track and other metrics can be derived from the instance.

MODEL PARAMETERS – The Suspen symmetric SLA suspension model is defined by the following parameters. The no unit measurement for length is defined in the code, but must be used consistently. For these tests inches were used.

Upper Chassis Width - the distance between the point where the right and left upper A-arm connects to the chassis See Figure 1 for this and other parameter definitions.

Lower Chassis Width - the distance between the point where the right and left lower A-arm connects to the chassis

Chassis Height – the vertical distance between the points where the upper and lower A-arms connect to the chassis.

Static Height –the vertical distance from the ground plane to the lower A-arm chassis connection point at static ride height.




Figure 1: Chassis Model Dimensions

Upper A-arm Length –the 2D length of the upper A-arm. See Figure 2 for Upper A-arm length definition.

Lower A-arm Length –the 2D length of the lower A-arm. See Figure 2 for Lower A-arm length definition.




Figure 2: A-arm Dimensions

Upright Length – the 2D distance from where the lower A-arm attaches to the upright to the point where the upper A-arm attaches. See Figure 3 for Upright Length and following parameter definitions.

Upper Upright Offset – the distance from the point where the upper A-arm attaches to the upright and the tire centerline.

Lower Upright Offset – the distance from the point where the lower A-arm attaches to the upright and the tire centerline.

Lower Offset Height – the vertical distance from the point defined in the Lower Upright Offset and the bottom of the tire.

Tire Rolling Radius – the radius of the tire as loaded.

Tire Width – the width of the tire. This parameter is currently used only for visualization purposes.




Figure 3: Tire & Upright Model Dimensions

SCORING MECHANISM

There are many things that must be considered when determining the fitness of a given suspension design. These are not all the same magnitude or even the same units. So coming up with a single fitness value is difficult.

The basic approach here is to design with a unit-less measure of fitness for each test and then to combine these with a weighting function. The unit-less measure is created by a function. One can visualize this by thinking of the normal or bell-shaped curve. The ideal value is at the mean. Further away from the mean on either side causes a reduction in score. For the typical test the user supplies an ideal value (C) and two acceptable bounds ($Bound_L$ & $Bound_R$). These values are used to define a scoring function.

Several functions were analyzed and compared in evaluation speed and accuracy using the Genetic optimization algorithm scoring. Most of these scoring functions used different orders of the normal distribution. We chose the Order -1 Normal Distribution, Equation (1), because it converged the fastest.

$$score_i = \begin{cases} 1 - e^{-\frac{C - Bound_L}{3(C - x_i)}}, & x_i \leq C \\ 1 - e^{-\frac{Bound_R - C}{3(x_i - C)}}, & x_i > C \end{cases} * 100\% \quad (1)$$

The function is smooth and continuous. The ideal score is 100%. The score at a bound is 28.3%. And it approaches 0% as the value goes to infinity.

MODEL SCORING – A suspension is scored by generating a sequence of instances based upon a sequence of body positions. These instances define the position of the chassis, linkages, and tires. Each scoring metric picks the instance(s) that apply and scores the result. Each metric score is combined by way of a weighting function (2). W_i is the weight for the i^{th} scoring metric and $score_i$ is the score from that metric. The resulting score is normalized.

$$TotalScore = \frac{\sum (W_i * score_i)}{\sum W_i} \quad (2)$$

An example of scoring is shown in Figure 4. This shows the space of a two-degree optimization using two metrics. The parameters being optimized were the Lower and Upper upright offsets. The two test metrics were static camber and track width. The diagonal represents the camber angle where the two offsets must be close to equal. The track width test generates a much broader peak running roughly vertically in the image. The ideal point is where the ridges intersect. These metrics are further explained below.

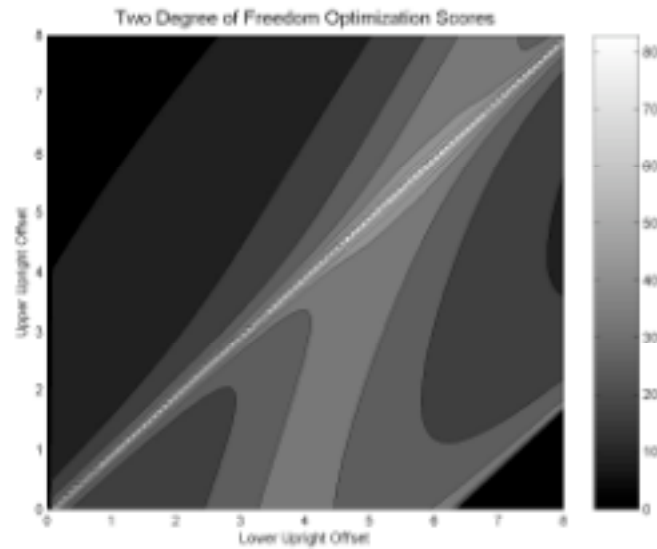


Figure 4: Scoring Test using Order -1 Normal Distribution

SCORING METRICS – Multiple parameters are scored. Example scoring metrics are defined below:

Static Camber- The Static Camber test measures the tire camber when the vehicle is at the static location. For all tests conducted here the ideal static camber was -1.0 ($C = -1.0$) degrees and the bounds were at -2.0 and 0.0 degrees ($Bound_L = -2.0$ and $Bound_R = 0.0$).

Track - The Track test measures a vehicle track width at the static location. The ideal track was set to 48.0 inches and the bounds were at 45.0 and 49.0 inches.

Scrub - The Scrub test measures the change in a vehicle track width as it moves through the defined path. The ideal scrub is zero. The upper bound is defined at 1.0 inch.

Lateral RC Movement – This test measures the lateral component of the roll center as it moves. The difference between the extremes is measured and scored. The tolerance is at 24.0 inches ($C = 0.0$ and $Bound_R = 24.0$, $Bound_L$ is unused because the value is always non-negative).

Vertical RC Movement – Similar to the Lateral Roll Center Movement Test, this measures the vertical component of Roll Center movement. This is measured relative to the chassis. The tolerance for vertical movement is 0.5 inches ($C = 0.0$ and $Bound_R = 0.5$).

Jacking – This metric estimates the effects of jacking force. It uses the roll center height, track, and the roll for each position the chassis moves through. Roll is used as a representation of lateral force. The largest jacking estimate is scored. 1.0 is used as the bound ($C = 0.0$ and $Bound_R = 1.0$). This bound corresponds 8.0 inches of roll center height at 3.0 degrees of roll for our ideal track width.

Total Camber Change – This measures the camber as the vehicle moves through its travel. The difference between the extremes is scored. In our camber optimization tests the bound for camber change is 0.5 degrees. A bound of 6.0 degrees was used for the final optimization.

Camber (Bump) – This test measures the camber similar to the measurement in the Static Camber test however this is measured at the extreme bump travel. For these tests the ideal camber was defined at 0.0 degrees with bounds at ± 1.0 degree.

Camber (Droop) – This is the opposite of the Camber (Bump) test; it measures the tire camber angle at the extreme droop travel. Because the droop condition is associated with the low tire load, it uses a comparatively lower weighting.

Laden Camber – This measurement extends the above camber tests into the extremes of roll travel. The test analyzes camber angle of the tire that gains weight

transfer due to cornering forces. The ideal camber was defined as -1.0 degrees and bounds at $+0.5$ and -2 degrees.

UnLaden Camber – This test is used to include the contribution of the tire that loses loading due to weight transfer. As a result of the unequal contribution of the laden or unladen tire these tests use different weights when the scores are combined.

GENETIC OPTIMIZER

In redesigning the Suspen model, it was determined that an optimization of the model would perform at a more efficient level if a genetic algorithm was incorporated into the model. Genetic algorithms (GAs) may also be referred to as evolutionary programs. Classical GAs were first developed in the early 1970's by John Holland and were later expanded by De Jong in the 1980's [10]. Simply put, GAs attempt to encode the process of "survival of the fittest" and use that process to determine optimal solutions for a given problem and solution space. GAs use the information in the current solution population in an attempt to produce an intelligent random search. Through a small set of operations, GAs can efficiently search a large solution space for optimal solutions. However, one huge drawback of GAs is that they are not guaranteed to find the optimal solution and may instead converge to local optimal solutions.

The steps that describe a genetic algorithm are [10]:

- 1) Produce an initial population with solutions generated at random
 - a) Calculate the fitness of each solution
- 2) Repeat the next steps until a solution meets the exit criteria
 - a) Determine the probability of selecting any given solution for mating, solutions with a higher fitness have a greater probability of being selected
 - b) Produce a new generation using two operations
 - i) Duplication/Reproduction, the direct copy of a solution
 - ii) Crossover/Mating two solutions with the possibility of mutation
 - c) Calculate the fitness of each new solution

Chromosomes can be expressed as a structure of properties stored as real values. This representation of a chromosome is more natural for applications for which conversion to a binary structure is not feasible due to the size of the solution space or constraints on the problem. When using a structure chromosome, the crossover operation is changed slightly. Instead of a positional crossover, an arithmetic crossover of some or all of the individual genes can be performed. The simple arithmetic crossover [10] with a crossover proportion of a ,

where a is a value from 0 to 1, for gene g in chromosome one and two is shown in Equation (3).

$$\begin{aligned} \text{new } g1 &= a * (g1) + (1 - a) * (g2) \\ \text{new } g2 &= (1 - a) * (g1) + a * (g2) \end{aligned} \tag{3}$$

The added benefit of a simple arithmetic crossover is that the resulting chromosomes will be guaranteed to be within the solution bounds. Other possible crossover operations exist; however, some crossover operations may lead to solutions lying outside the solution space.

The selection of an individual chromosome for crossover is proportional to the chromosome's fitness. Generally, the fitness of each solution in a GA population must be greater than zero and fitter solutions have larger numbers. Because of this restriction, the determination of a solution's fitness within the population is many times a two-step process. First, the raw fitness for every solution is determined. Second, if any of the raw fitness values is less than or equal to zero, all the raw fitness values must be shifted so that all the final fitness values are all greater than zero. Once all the final fitness values are determined, the probability that a single chromosome will be chosen for mating is determined by the ratio of a chromosome's fitness divided by the sum of all the population fitness values.

An additional technique that is often implemented in GAs are Mutations. During the reproduction step, a gene has a random chance of mutating. That mutation can be any number of things, a new random value, a bound, or something else. Mutations are useful in breaking outside of the bounds of the existing solutions and local minima.

Although many GA applications exist, our team implemented its own GA application so that we would get a better understanding of how the algorithm works and is implemented. GAs have been used in successfully examining NP-Hard problems like the Traveling Salesman problem and in adaptive learning algorithms [10]. It is interesting to note that many of the operations within the GA process can be performed in parallel; so, GAs can take advantage of parallel and multi-processor machines.

A simple Grid Optimization algorithm was used for comparison. The Grid algorithm breaks each variable (degree of freedom) into equally spaced points. Each combination of points is evaluated and the best is chosen. The advantage of the Grid algorithm is that it covers the search range. However in order to cover that range it requires n^D evaluations, where n is the number of grid elements per degree of freedom and D is the number of degrees of freedom. The number of combinations gets very large as D increases.

As a baseline for the optimizer evaluation, a simple discrete grid optimizer was implemented. Four test cases

were examined. In Test A, a two-degree of freedom known model that allowed the upright offsets to be changed was scored on track and static camber. In Test B, the same model allowed the two A-arm lengths to be changed and was scored on a more complex set of tests. An optimal result from Test B, unlike Test A, does not have a 100% score.

These tests were scored and timed. The results are plotted in Figure 5. The Genetic Algorithm produces ideal results faster than using the Grid Algorithm.

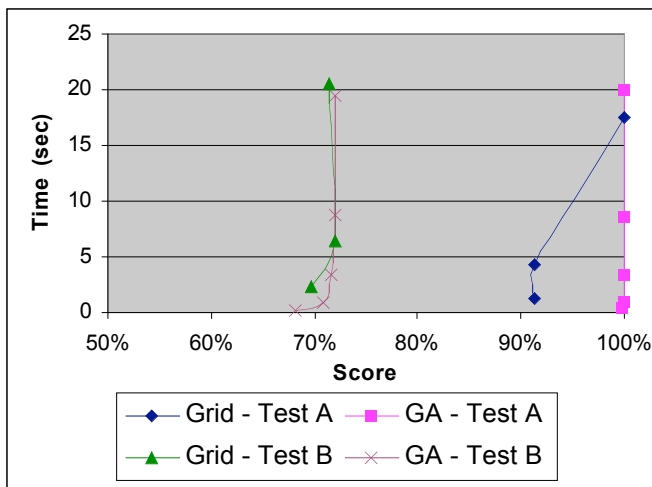


Figure 5: Two-Degree of Freedom Optimization Tests

Both preceding tests were again executed using a three-degree of freedom optimization that allowed both A-arm lengths and the upright length to be changed. As can be seen in Figure 6, the Genetic Algorithm is significantly faster. This is due to increased dimensionality of the problem.




Figure 6: Three-Degree of Freedom Optimization Tests

As the dimensionality increases the problems become intractable: so the Grid method is impractical. For

example, a 25-element grid on a full 11-degree of freedom optimization will take 34,400 years to complete.

ANALYSIS OF THE MODEL OPTIMIZATION OUTPUT – The two and three-degree of freedom tests described above showed that the optimizer could find the correct solution for these limited dimensionality cases. To analyze optimization of the full 11-degree of freedom system, two special cases were tested.

Special Case One - Case one was an optimization for Camber Change in Bump and the resulting suspension design should tend toward an equal length parallel design. A test was set up that moved the chassis through 1.5” of bump to 1.5” droop. The individual models were analyzed by the amount of total camber change occurred over that range of motion. Two other scores for track width and camber were included to keep the model ‘reasonable’ but were given low weights. The bounding models allowed a great flexibility in design. Given a population of 1000 chromosomes and 50 generations, the design shown in Figure 7 was the result. The parallelism is shown in Table 1




Figure 7: Case one - Suspension Visualization

Table 1: Case one results

Upper A-arm Length	Lower A-arm Length	Difference
18.4946”	18.5449”	0.27%
Chassis Upper Width	Chassis Lower Width	
7.4537”	7.0852”	5.2%
Chassis Height	Upright Height	
6.9194”	6.8491”	1.0%

Over the range of motion defined, the camber changes only 1/10th of a degree. It is apparent that this model is converging toward equal length and parallel.

Special Case Two - Case two was an optimization for Camber change in roll. The resulting model should tend toward miniscule chassis size with a swing axle configuration. Using a test similar to case one, the algorithm minimized the changes of camber through the range of -3.0 degrees to +3.0 degrees of roll. The resulting model is shown in Figure 8.




Figure 8: Case Two - Suspension Visualization

The camber change is 0.08 degrees over that range. It could improve by shrinking the chassis even more. However, the GA algorithm can have difficulty finding solutions at the boundaries. This is due to our selection of crossover function. It ensures that new values are linear combinations of existing values. So a value not contained in the initial population, would never be generated.

To overcome this, mutations were implemented into the algorithm. The mutation is randomly chosen from a new random value, a boundary value, or performing the crossover operation with a boundary. A mutation rate was 1% was used. Too large a value reduces the Genetic Algorithm to a random optimization.

Running the above case again with mutations yields the model as seen in Figure 9. Camber change for this model is only 0.02 degrees over the same range, so this model is an improvement.




Figure 9: Case Two result with Mutations

However, we note that this is not the ideal design described above. It appears that the GA algorithm does not converge to this design because all three chassis dimensions must all be very small to score well. Chassis dimensions as small as 0.05 inches still cause more camber change than in the above solution.

Real-World Design - The true test is that the optimizer should be able to produce an optimum suspension for a real-world vehicle. A test was developed that moved the chassis through 3 degrees of roll and 1.5" of bump and droop. The individual models were scored on a sequence of tests that should be representative of what a designer would be looking for as a rear suspension on a Formula SAE car.

The bounds and resulting solution are described in Table 2. The tests that the optimization was based upon are described in Table 3. A visualization of the suspension is shown in Figure 10. A dialog showing the score of the resulting model is shown in Figure 11.

Table 2: Optimization Bounds and Solution

	Min. Bound	Max. Bound	Optimized Solution
Chassis			
Upper Chassis width	.01	20.0	8.236508
Lower Chassis width	.01	20.0	4.649678
Chassis height	.01	8.0	5.318161
Static height	4.0	9.0	5.872972
A-Arms			
Upper A-Arm length	8.0	30.0	14.470783
Lower A-Arm length	8.0	30.0	18.416723
Upright/Wheel			
Upright length	2.0	10.0	6.594033
Upper Upright Offset	0.0	9.0	5.042054
Lower Upright Offset	0.0	9.0	3.059438
Lower Offset height	2.0	7.5	5.121896
Tire Rolling Radius	9.0	13.0	10.464683




Figure 10: Optimized Model

Table 3: Optimization Test Metrics

Test	Ideal	Min.	Max.	Wt.
Static Tests				
Camber	-1.0°	0.0°	-2.0°	0.30
Track Width	48.0"	45.0"	49.0"	1.50
Scrub	0.0"		1.0"	0.50
Roll Center				
Lateral RC Movement	0.0"		24.0"	0.50
Vertical RC Movement	0.0"		0.5"	0.30
Jacking	0.0		1.0	1.50
Camber				
Total Camber Change	0.0°		6.0°	1.00
Camber (Bump)	0.0°	-1.0°	1.0°	0.80
Camber (Droop)	0.0°	-1.0°	1.0°	0.25
Laden Camber	-1.0°	-2.0°	0.5°	5.00
Unladen Camber	1.0°	-0.5°	2.0°	0.60

The optimized model has a number of characteristics in common with professionally designed systems. The lower chassis suspension pickups are closer together

than the upper ones. The A-arms are as long as practical. And the upper A-arm is shorter than the lower.

Test	Score	Weight	Weight Score
Static Camber	54.52	0.30	16.36
Track Width	99.66	1.50	149.49
Scrub	48.62	0.50	24.31
Lateral RC movement	99.99	0.50	50.00
Vertical RC movement	90.19	0.30	27.06
Minimum Jacking	52.29	1.50	78.44
Total Camber Change	26.84	1.00	26.84
Average Static Camber (Bump)	44.48	1.00	44.48
Average Static Camber (Droop)	44.48	0.25	11.12
Average Laden Camber	49.03	5.00	245.16
Average Unladen Camber	13.31	0.60	7.99

Total Test Score: 54.72%

Figure 11: Optimized Model Scoring

The performance parameters of the optimized model are shown below. The Roll Center position (Figure 12) is where this model excels. The vertical component changes by only 0.07 inches across the tested range. At the same time, the lateral component moves 3.07 inches. And the vertical position is low enough to keep jacking effects low.

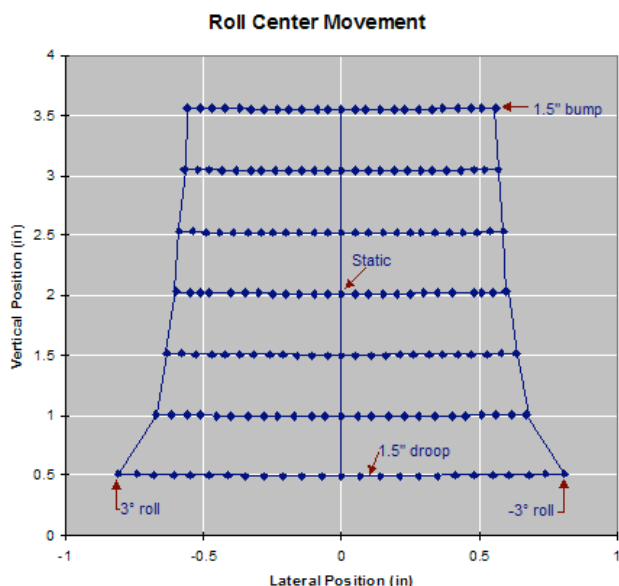


Figure 12: Roll Center Position

The camber change is linear through roll (Figure 13). The total change over the test range is 6.04 degrees. The camber curve is a balance of ideals for braking, cornering, and accelerating; and none of them match the ideal that was tested. However, several of the individual camber tests score near 50% as seen in Figure 11.

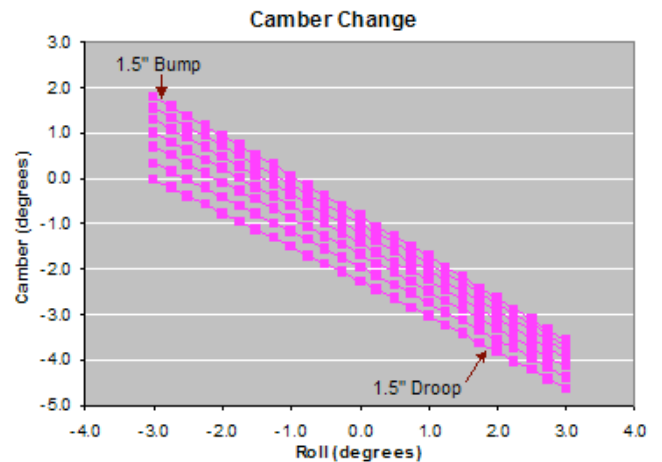


Figure 13: Camber Change

The track width is within 0.18 inches of the specified ideal from the scoring test. The track change is limited to 0.5 inches over the tested range.

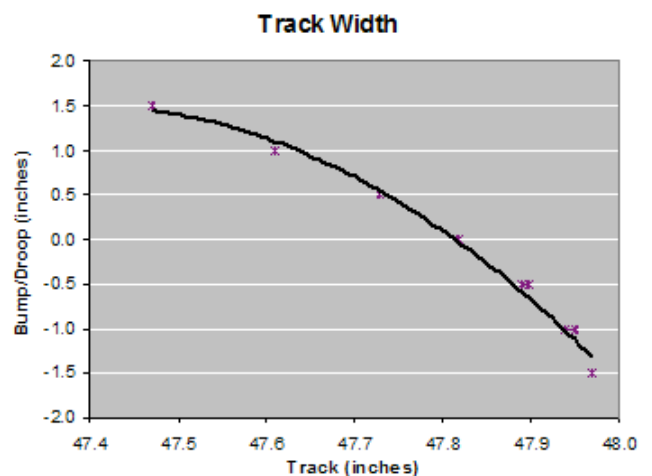


Figure 14: Track Width

CONCLUSION

The Genetic Algorithm and Scoring Mechanism were shown to work effectively and significantly faster than the Grid Optimization technique. The improvements allow intractable problems to be examined. The Scoring Mechanism does an excellent job in allowing sufficient variability for optimal results and is an elegant solution to the problem of merging disparate metrics.

For future work, the underlying models can be expanded to handle both front and rear suspensions simultaneously. Asymmetric SLA and McPherson strut suspensions could be developed as well. The Genetic Algorithm settings can be tuned. The duplication rate, mutations, and convergence testing can be modified to produce more efficient results.

REFERENCES

1. Milliken, William F. and Douglas L. Milliken, Race Car Vehicle Dynamics, SAE International, 1995.
2. Bastow, Donald and Geoffrey Howard, Car Suspension and Handling, Third Edition, pp 1-14, Pentach Press Limited, 1993.
3. Smith, Carroll, Tune to Win, Aero Publishers, 1978.
4. Gillespie, Thomas D., Fundamentals of Vehicle Dynamics, SAE International, 1992.
5. Mabie, Hamilton H. and Charles F. Reinholtz, Mechanisms and Dynamics of Machinery, Fourth Edition ed., pp 20-26: Wiley and Sons, 1987.
6. Kreuzer, Martin and Lorenzo Robbiano, Computational Commutative Algebra I, Springer 2001
7. Introduction to Genetic Algorithms:
http://www.doc.ic.ac.uk/~nd/surprise_96/journal/vol1/hmw/article1.html
8. What is Genetic Programming?:
<http://www.genetic-programming.com/gpanimatedtutorial.html>
2003.
9. DeJong, Kenneth A and William Spears, Learning Concept Classification Rules Using Genetic Algorithms, Machine Learning 13, pp. 618-625, 1993.
10. Michalewicz, Zbigniew, Genetic Algorithms + Data Structures = Evolution Programs, Springer-Verlag, 1992.

11. Dawid, Herbert, Adaptive Learning by Genetic Algorithms, Springer-Verlag, 1999.

DEFINITIONS, ACRONYMS, ABBREVIATIONS

Some of the basic terminology in GAs used is:

Population: the group of possible solutions at a given time

Chromosome: a single solution

Gene: an individual trait within a chromosome

Duplication: the direct copy of a given solution; also known as reproduction

Crossover: the exchange of the information between two solutions; also known as mating

Fitness: the numerical score for how good a solution is; the fitness may need to be calculated relative to the other solutions

Pair Selection: the selection of two solutions for crossover; the probability a solution is chosen for crossover is related to its fitness in the population