

Monte Carlo Analysis for Accident Reconstruction

Wade Bartlett, 22 December 2007

INTRODUCTION

OK, I confess: I have never liked time and distance analyses. Perhaps the unpredictability of humans scares me more than the variability of physical phenomena like friction and force. The prospect of having to pick one or two values for each of the many variables in a typical time/distance crash makes me want to run and hide. I think I've turned a corner though, and am starting to like time/distance analyses. What changed? The tools I'm using. I have finally tried using *Monte Carlo Analysis* (MCA) for a time/distance case. This is a probability analysis technique first used to great effect by scientists working on the Manhattan Project, who had access to the newly created ENIAC computer. Since then the technique has found wide acceptance with folks involved with risk assessment, financial decisions, meteorology, nuclear engineering, traffic flow, project scheduling, failure prediction, chemical processes, and others. Now its gone to work for me, and it can go to work for you, too.

Through the computing power currently available with even a modest personal computer, it is now possible to evaluate virtually all the possible combinations of all the variables from A to Z for many crashes in such a way as to allow the determination of a result to a specific level of confidence. Having a scientific foundation for the confidence limits for our results is becoming more and more important as judges look to Daubert's "Reliability Standard" in determining who gets to testify in today's courtrooms. There are other means of quantifying certainty, or "confidence limits," but MCA is now my favorite for any analysis that I can conduct as a series of closed-form equations in Excel. Put another way, if I can solve the case by hand-methods with a pencil and a calculator, then I can put it in Excel, and I can use MCA.

After a brief introduction to MCA and a literature review, including a 2003 SAE paper I wrote on how to run MCA in Excel, this article will build on that original paper by working out a momentum-based analysis for a full intersection collision case study. The case examined involved a 90-degree intersection, one speeding motorist, and one motorist who ignored a stop sign. Pre-crash event data was recovered from the airbag control module of one vehicle, providing information on its incoming speed and delta-V, which were incorporated into the analysis. This article will include only a rudimentary discussion of statistics, and the reader is directed to the cited references for further information, any including statistics primer or my 2003 SAE paper which specifically described the steps necessary to implement a Monte Carlo Analysis using a spreadsheet program.

LITERATURE REVIEW

The notion of Monte Carlo Analysis dates back more than 100 years but was quite limited until the advent of current computers. There are a variety of decades-old books either dedicated to MCA [Hammersly & Handscomb 1964] or which discuss the technique [Schlaifer, 1969] and more recent textbooks on the topic [Fishman 2003, Christian & Casella 2004], but none that specifically address its application to accident reconstruction.

The earliest papers specifically describing the use of MCA for crash reconstruction that I have found are a pair of 1994 SAE papers from Wood & O'Riordain [1994] and Kost & Werner [1994]. Several contemporary papers and some earlier ones discussed various techniques for uncertainty analysis. They primarily approached the issue from a mathematical standpoint, and were not directly relevant to MCA, though MCA is clearly a viable means of assessing the uncertainty in an analysis, and may be considered related [Slakov & MacInnis 1991; Neiderer 1991, Brach 1994, Tubergen 1995, Bartlett and Fonda 2003, Fonda 2004].

Wood & O'Riordain [1994] discussed using MCA with an "in house simulation package" to evaluate vehicle avoidance maneuvers prior to an intersection crash. Though their paper does not provide the actual code utilized, it appears that their program was essentially an automated form of the closed-form analysis commonly performed by hand for momentum and time and distance analyses, incorporating some post-impact rotation analysis. Their case study included ranges for a number of variables, but other than one reference for the brake rise time, no specific citations were provided which could assist with selecting appropriate ranges or distribution types for other analyses. The paper discussed using logic tests to discard or accept each run, which they called "Redundancy." This is based on the idea that one can quite often place some ranges on one or another result, either based on physical impossibility, geometry, known vehicle dynamics, or witness statements. These limits can allow us to discard calculated results which do not fit our known true result. For instance, we can immediately discard any sets of data calculated using friction values less than zero. Another example might be that in right-angle intersection crashes with significant central engagement between the two vehicles, we can usually assert with confidence that the departure speeds for the two vehicles should be similar, so one could discard all results where they differ by some selected value, say 4 or 6mph, as being inconsistent with the crash under consideration. If it is known that one car is moving slowly across another's path, but the direction of motion is known to be to the left, for instance, all cases where the analysis indicates it was going to the right can be discarded. Wood & O'Riordan selected the following four limits for their intersection crash: Vehicle 1 had to depart faster than Vehicle 2; the departure speeds had to be within 3.5mph of each other; vehicle closing speeds had to be between 45 and 99 kph (based on crush damage); and the pre-crash cornering acceleration for one vehicle had to be less than 0.55g based on scene evidence. Using these limits, they selected very wide input ranges ($0.7 < f < 1.05$, for one vehicle's deceleration, for instance), but discarded 99.4% of the results as generating results outside their limitation ranges. The results for each variable of interest for the remaining 1591 cases formed what I would call "good looking" or "filled in" bell curves with no significant gaps. Though no statistical analysis was presented, they noted that they re-ran the analysis until the overall average results "stopped changing."

Kost & Werner [1994] described using the Crystal Ball software package, which is an Excel add-on package, to evaluate a vehicle's initial speed based on the energy dissipated during the crash. They describe three common probability distributions (normal, triangular, and rectangular, aka uniform or even), but provide no guidance on selecting appropriate values for the ranges. Crystal Ball is still available, currently costing something over US\$1,000, depending on which package one purchases.
[<http://www.crystalball.com/>]

Since those early expositions, there have been a number of papers which described the technique's foundation, attributes, and applications to reconstruction. Moser et al [2003] discussed a Monte Carlo-style parameter variation application in PC-Crash. Kimbrough [2004] described the use of MCA to analyze a passing-situation which lead to a crash, but not the mechanics of the technique. Moser et al described the use of MCA to evaluate incoming vehicle speeds, and analyzed one test-crash. They used the term "Conditional Sampling" to describe the selective discarding of unreasonable results called "Redundancy" by Wood & O'Riordain. Ball et al [2007] used the Crystal Ball add-on and discussed the effect of distribution selection on results.

In addition to Crystal Ball (mentioned above), there are other Excel-add-on packages available, including RiskAmp (<http://riskamp.com/>) which costs approximately \$250 as of this writing for the Professional version. This version includes Latin Hypercube sampling (LHS) capability for some distribution types which can significantly reduce the number of calculations required for a Monte Carlo Analysis. For simplicity, though that technique will not be utilized for this paper. RiskAmp has a 30-day trial available. Playing with it for a while I've learned several things: (1) With occasional HELP FILE reference, I was up and running in about 15 minutes, so it's pretty simple; (2) It's REALLY cool to watch it dynamically run an analysis (Visit <http://mfes.com/MonteCarlo.html> to see a little AVI screen capture of the package doing a skid-to-stop analysis with two normally-distributed variables); (3) It has loads of distribution choices and it

seems to be only slightly slower than Excel alone for recalculations; (4) It does NOT easily allow the conditional sampling described below, which turns out to be a deal-breaker for me.

In my 2003 SAE paper [Bartlett 2003] I described the specific steps necessary to set up a Monte Carlo analysis in Excel, including some of the “shortcut” keys and functions which made operating in Excel simpler, however that paper did not include a complete collision example.

STATISTICS REFRESHER

Many naturally occurring phenomena (say height of adult male humans, or length of year-old crocodiles) have some middling-“most likely” value at or near which most of the examples occur. The frequency of examples that are different from the mean becomes less frequent the further from the mean we get, either higher or lower. Said another way, we’re more likely to find examples nearer the middle “typical” value than we are to find values very different from it. This is shown graphically in the “bell curve” so often used in schools” Most students’ grades are near “average”, while only a few perform much higher, and only a few perform much lower.

For this paper, I’ll skip most of the normal distribution mathematical discussion, but it is worth noting that the two key values are the mean (or average) value, which is the middle “high spot” on the bell curve, and the “Standard Deviation” (SD) which defines how “spread out” the curve is. In a normally distributed population or data set, the measured values fall within one standard deviation of the mean value (Avg-SD to Avg.+SD) 68.3% of the time, (34.1% of all values on either side of the middle) as shown in Figure 1. About 95.5% of all cases fall within 2 standard deviations of the mean, and 99.7% will be within 3 standard deviations. The two “tails” above and below 3 standard deviations from the mean comprise 0.15% of all cases.

Figure 1: Showing a normal distribution (mean=0, SD=1) and the percentage of cases in each standard deviation segment


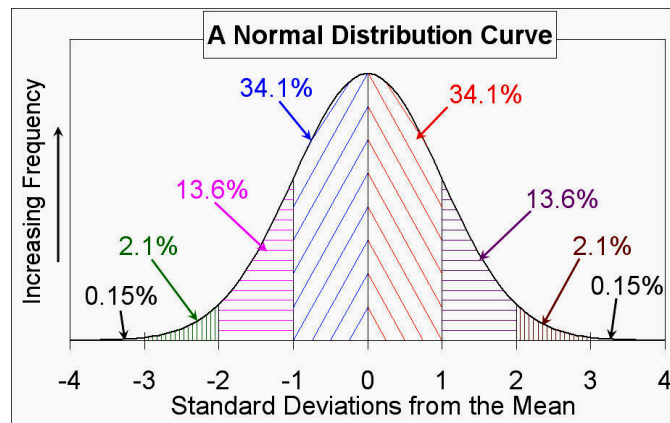
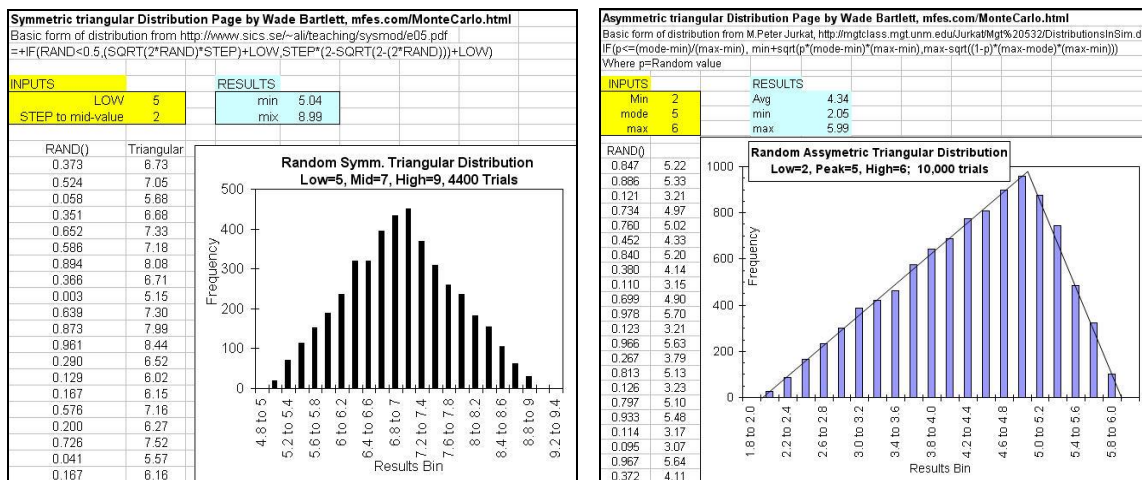


Figure 2: An example of a randomly selected normally-distributed data set having a mean value of 0.76 and a standard deviation (SD) of 0.06. We’ll see these values again later.

Normal distributions have no absolute upper or lower bound, so theoretically there may be values far from those near the middle. Sometimes these highly unlikely values are physically impossible, such as negative friction values. For this reason, I quite often place limits on my normally distributed values as appropriate.

A second type of distribution commonly encountered in accident reconstruction analyses is the uniform or “rectangular” distribution, also sometimes called an “even” distribution. This type can be used in the absence of evidence to suggest that the value is more likely to be near the center of the range. This distribution is more conservative than a normal distribution, as it gives equal probability to all values in the specified range. At the same time, it absolutely precludes values outside the specified range, so the bounding minimum and maximum terms must be selected with great care so as to not exclude possible values. The standard deviation for this type of distribution is equal to $[a / \text{SQRT}(3)]$. The range within one SD of the mean includes 57.74% of all values. [Montgomery & Runger 1999]

The third commonly used type of distribution would be a triangular distribution. This type offers the highest probability to a middle-value, with linearly decreasing probability as you move away from that middle value, and concrete minimum and maximum values. This type of distribution incorporates the analyst’s BEST VALUE as the most likely, but still allows a range of values, with decreasing probability as you get away from the middle. Triangular distributions are available in all commercial Excel Add-on MCA packages, but to my knowledge it is not inherently available in Excel itself. You can code it yourself, though, with relative ease. The equation for a symmetric distribution is shown at this website: www.sics.se/~ali/teaching/sysmod/e05.pdf. An implementation of this in Excel using the random number generator gives the distribution curve shown in Figure 3A below. One assymetric triangular distribution is provided by this online article: <http://mgtclass.mgt.unm.edu/Jurkat/Mgt%20532/DistributionsInSim.doc>, with its implementation shown in Figure 3B.



Figures 3A, 3B: Implementing two triangular probability distributions, one symmetric, one not.

There are numerous other types of distributions, some symmetric (such as three of the above four), others asymmetric (such as the assymetric triangular discussed, or the Chi or Pareto functions). Some have an appearance similar to the normal distribution (such as Cauchy-Lorentz and Student’s t function), while others are very different (such as the Laplace or Double-Wiebull). With the mathematical formula describing the function in hand, most of these could be coded into Excel. Fifty Six different distributions are provided in this online brochure for the REGRESS+ software package: [McLaughlin 2001]. These could each be implemented in Excel, but not without some effort.

There are numerous other types of distributions, some symmetric (as the three discussed above are), others asymmetric, such as the Chi or Pareto functions. Some have an appearance similar to the normal distribution (such as Cauchy-Lorentz and Student’s t function), while others are very different (such as the

Laplace or Double-Wiebull). With the mathematical formula describing the function in hand, most of these could be coded into Excel. Fifty Six different distributions are provided in this online brochure for the REGRESS+ software package: http://www.causascientia.org/math_stat/Dists/Compendium.pdf [McLaughlin 2001]

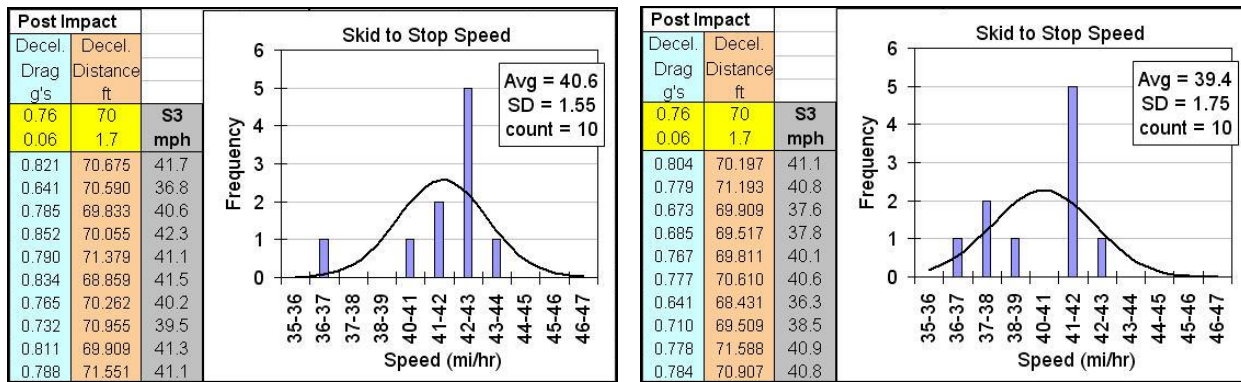
In addition to Crystal Ball (mentioned above), there are other Excel-add-on packages available, including RiskAmp (<http://riskamp.com/>) which costs approximately \$250 as of this writing for the Pro version which includes Latin Hypercube sampling (LHS) capability for some distribution types. Though LHS can significantly reduce the number of calculations required for a Monte Carlo Analysis, it will not be utilized for this paper. RiskAmp has a 30-day trial available. Playing with it for a while I've learned several things: (1) With occasional HELP file reference, I was up and running in about 15 minutes, so it's pretty simple; (2) It's REALLY cool to watch it dynamically run an analysis (visit mfes.com/MonteCarlo.html to view a short AVI video screen-capture of the package doing a skid-to-stop analysis with two normally-distributed variables); (3) It has loads of distribution choices and it seems to be only slightly slower than Excel alone for recalculations; (4) It does NOT easily allow the conditional sampling described below, which turns out to be a deal-breaker for me.

I generally use whichever probability distribution type fits the available data best. Regardless of which types or combinations one selects for their variables though, calculations incorporating multiple independent variables tend to form normally distributed results. This phenomenon is described by the Central Limit Theorem which is described in more detail in all statistics texts and numerous online sources. The analyst should remember, though, that selectively discarding portions of one result may make a normal-result assumption improper, requiring more careful evaluation of the results to determine probability ranges.

INTRODUCTION TO MONTE CARLO ANALYSIS

The basic idea behind MCA is that if we know the range in which our input variables fall, we can randomly select a possible value for each variable from that range and run the analysis to get a result. Doing this a few times will only yield a few potential results, and essentially nothing useful. If we repeat this process often enough, though, say thousands or tens of thousands of times, the collection of results provides insight into the full range of what may possibly have transpired prior to or during the event under analysis. We can use this collection of results to assess not only the minimum, maximum, and average values, but also what range constitutes the “most likely” (the middle 51%), or perhaps the range that constitutes a 95% or 99% confidence level. A simple example follows.

Say we wish to evaluate a vehicle's speed at the start of a nominally 70-foot long skid on dry pavement. Based on a review of hundreds of skid tests summarized at this webpage: mfes.com/friction.html, I might select a normal distribution (a “bell curve”) with an average value of 0.76g and a standard deviation (SD) of 0.06g (written for this article as $0.76 \pm 0.06g$). This means that 95% of all cases will fall between 0.64g and 0.88g, and 99% will be between 0.58g and 0.94g. I was a co-author on a 2002 SAE paper on research into the ranges inherent to many measurements common to reconstruction [Bartlett, et al 2002]. A review of that paper shows that the standard deviation for measurement of a skidmark was found to be approximately 1.7 feet. Using Excel and the steps described in Bartlett [2003], ten random runs generate choppy results, as shown in Figures 4A and 4B for two successive runs of just 10 trials each. In Figure 4A, for instance, we see that the 10 trials produced one occurrence each in three ranges (36 to 37mph, 40 to 41, and 43 to 44mph), two results between 41 and 42mph, and 5 results between 42 and 43mph, while generating zero results in all other ranges. Figure 4B shows similar uneven distribution. But once the spreadsheet is set up, the calculation can just as easily be performed a hundred times, shown in Figure 5, or 10,000 times, shown in Figure 6. Normal curves for each run's specific statistics are also shown in Figures 5 and 6. The more trials one performs, the less variation there will be between runs, and the more “filled in” the normal curve appears. For a simple two-variable calculation such as this, there may not be much additional benefit to be had beyond a few hundred trials. For more complicated analyses, though, tens of thousands of runs may be required to get results to satisfactorily converge.



Figures 4A, 4B: Two individual runs of 10 speed calculation trials. Both figures were generated by using normal distributions for drag (mean=0.76, SD=0.06g) and distance (mean=70, SD=1.7ft), showing normal distributions based on each run's resulting mean and standard deviation.





Figure 5: Extending the Monte Carlo results from Figures 4A & 4B to one run of 100 trials, with normal curve based on the results overlaid. Note that histogram is beginning to fill in the bell curve, but is still uneven in places.

Figure 6: Extending the Monte Carlo analysis of the skid-to stop problem to 10,000 trials.



The complete result set can then be cut-and-pasted (values only, not equations) into a separate worksheet for further analysis. By sorting the results from low to high, they can be plotted as shown in Figure 7, which may be more informative on some levels than the histograms above. This chart is essentially a form of the Cumulative Distribution Function. The confidence ranges of interest can then be defined (50.1% and 95% are shown), and either read off the chart, or picked out of the list. From the chart (or table) the “more likely than not” speed from the skid described above would be 38.8 to 41.0mph. The 95% confidence range would be 36.6 to 43.0mph. Said another way, we can be 95% certain that the true speed was between 36.6 and 43.0mph. The Excel spreadsheet used to generate Figures 4 through 7 can be found online at

<http://mfes.com/montecarlo.html>




Figure 7: Ordered results of skid-to-stop calculations from 10,000 trials shown in Figure 6, showing confidence ranges.

INTERSECTION CRASH ANALYSIS

The crash analyzed here involved two passenger cars. Vehicle 1 was a red late-model Chevrolet sedan with a lone occupant traveling north on a secondary road, and crossing a through-road at the edge of town. Vehicle 2 was a blue Toyota sedan with four teenage occupants which was traveling on the through-road in a 35mph zone. Several witness statements indicated that the red Chevrolet (Vehicle 1) did not appear to react to the stopsign or intersection at all. The Chevrolet entered the intersection and struck the blue Toyota (Vehicle 2) on the left side near the front left wheel. Post impact, both vehicles rotated somewhat, coming to rest approximately 70 feet from the area of impact, near the edge of the pavement. The Chevrolet's wheels were all still free to rotate, but one Toyota wheel was crush locked, giving it a higher Rotation Factor. Figure 8 shows the general crash scene diagram, and Tables 1 and 2 shows the value ranges used for the principle analysis variables.




Figure 8: showing the scene diagram for the intersection crash in the primary example.

Table 1: Momentum Analysis Variables for Vehicle 1:

<u>Variable</u>	<u>Low - High</u>	<u>Nominal / St.Dev</u>	<u>Distribution Type</u>
Weight (pounds)	3357 - 3957		uniform
Veh1 Post Impact Travel Distance (feet)	65 – 75		
Road Friction (g's)		0.76 / 0.06g	Normal
Rotation Factor	0.6 – 0.7		Uniform
Approach Angle (degrees)	0		Uniform
Departure Angle (degrees)	45 – 55		Uniform

Table 2: Momentum Analysis Variables for Vehicle 2:

<u>Variable</u>	<u>High/Low</u>	<u>Nominal / St.Dev</u>	<u>Distribution Type</u>
Weight	3461/3861		Uniform
Road Friction (g's)		0.76 / 0.06g	Normal
Rotation Factor	0.81 – 0.91		Uniform
Approach Angle (degrees)	74 – 78		Uniform
Departure Angle (degrees)	35 – 45		Uniform

Using the nominal values listed above, the nominal momentum solution can be performed. The calculated Chevrolet's incoming speed is 36.4mph, while the Toyota's incoming speed is 49.2mph. Since their total weights are essentially the same, they both experience approximately the same delta-V, calculated to be 29mph. The vector diagram for this solution using the nominal values is shown in Figure 9. (see www.mfes.com/colm_2005.pdf for more info on Conservation of Linear Momentum and vector diagrams).




Figure 9: The vector diagram of the nominal solution for the momentum analysis. (From the ARPro software package.)

So far, there's nothing new. But how do we assess the "most likely" range? We could use the traditional HIGH/LOW combined analysis technique, but the chances of getting all variables to line up like that are pretty slim. Now I'll turn to Monte Carlo Analysis: plugging in all those ranges in an excel spreadsheet as described above and in SAE 2003-01-0487. The Chevrolet's calculated incoming speed comes out to have a mean of 36.4mph with a standard deviation of 2.97mph, while the Toyota's incoming speed is calculated to have a mean of 49.1mph and a standard deviation of 2.86mph, with the bell-curve of possible speeds as shown in Figure 10.




Figure 10: Histogram showing all possible calculated speed results.

Subsequently, the Chevrolet's airbag control module was downloaded. It showed the Chevrolet gently slowing (44-44-43-43-42mph) during the five seconds prior to the collision. Due to uncertainties in speed and timing, I decided to only accept calculations which produced an incoming speed of 38 to 42mph.¹ This incoming speed for the Chevrolet is within the range of our initial calculations, but it is near the higher end, as we would hope and expect due to the intrinsically conservative nature of COLM analyses by virtue of neglecting external forces during the impact. Additionally, the Chevrolet's module recorded lateral and longitudinal delta-v, showing a total velocity change of at least 27.7mph. Limiting our examination to only those results which met the two criteria recorded by the airbag control module, and which generated departure speeds within 4mph of one another, we get approximately 1300 accepted scenarios per 10,000 trials. Collecting the results of 6 runs generated just under 8,000 accepted trials, as shown in the histogram of Figure 11. The conditionally sampled result set shows the range covering 95% of all cases for the Toyota's incoming speed to be 42.9mph to 52.9mph, which is not only slightly lower but is somewhat "tighter" than the whole "unfiltered" data set.




Figure 11: Histogram showing results after "conditional sampling": These include only cases where the calculated incoming Chevrolet speed was between 38 and 42mph, its change in velocity was more than 27.7mph, and the departure speeds for the two vehicles were within 4mph of each other.

¹ Despite the apparent simplicity of that dataset, airbag control module data is never that simple, and leaves room for uncertainty. There have been numerous papers published through SAE, ITAI, NHTSA, and others on this topic, which is beyond the scope of this paper.

With the speeds of the vehicles settled on, we can proceed with the part I have always liked the least: The time and distance analysis. We will run a time and distance evaluation for each scenario or trial run that we accepted from the previous analysis. We'll start by determining the changes which would have resulted from the Toyota's having been traveling at the speed limit at the time the Chevrolet crossed the stop-bar. Another group of variables must be defined for this stage of the analysis, shown in Table 3:

Table 3: Showing data used for the time-and-distance portion of the analysis.

<u>Variable</u>	<u>High/Low</u>	<u>Nom./St.Dev</u>	<u>Distribution Type</u>
V1 (Chevrolet) StopBar to Impact travel distance (feet)	37 – 45		Uniform
V1 Distance to clear impact (feet)	11 – 19		Uniform
Assumed V2 speed (mph)		35	None

For each trial already accepted, the distance for the Chevrolet to travel from the stop sign to the impact was selected from the range of possible values, and then using the travel speed calculated earlier and assuming that speed was constant, the time to impact was calculated. Using the Toyota's calculated speed for that run, its distance from the crash when the Chevrolet crossed the stop bar was calculated. Next up, the time it would have taken the Toyota to cover that distance had they been traveling the speed limit of 35mph. Then the extra time that would have afforded the Chevrolet to cross was calculated, and this was used with the Chevrolet's speed and distance-to-clear to determine if the crash still occurs or if the Chevrolet clears the intersection in time for the Toyota to travel behind it.

Using the nominal values, the wreck does NOT happen. However, the Toyota's slower speed and increased arrival time only allows about an extra 0.03 seconds of extra time. In other words, the Chevrolet only clears the Toyota by about 2 feet. The time for the Chevrolet to cover the distance from the stopbar to the impact is so short that there would have been very little chance for the Toyota driver to have reacted in any meaningful way. Running the Monte Carlo Analysis 100,000 times yielded 12,631 accepted trials, but since the Chevrolet's speed was limited to values higher than the average of all runs, the result changes: Evaluating only the successful trials showed that 52.5% of the time, the crash DOES still happen even if the Toyota had been traveling at the speed limit. In other words, the crash still happens more often than not. Of course, the "new" impact area is always further back on the Toyota's side, and the overall severity (of the initial impact, at least) is reduced.

If the situation had been such that the Toyota driver might have had time to react prior to the impact, her reaction to the event could have been incorporated, including reaction time, brake rise time, average pre-crash deceleration, and even steering maneuvers if the analyst was ambitious and so inclined.

CONCLUSION

This paper outlined the basic premise of Monte Carlo Analysis as it applies to crash reconstruction, reviewed some of the literature on the topic, and provided a case study. The case study showed how MCA can be applied to a complicated intersection crash utilizing ranges for all the major values in a momentum and time-and-distance analysis in order to assess actual likelihood that a speeding motorist's speed affected the occurrence of a crash. No other currently available tool that I know of combines MCA's easy availability, analysis technique flexibility, and conditional sampling capability. I expect I'll be using it more and more often, and maybe some day I'll even come to enjoy Time and Distance problems.

ACKNOWLEDGEMENTS

I would like to thank Bill Wright, Fred Hochgraf, Bill Messerschmidt, and Jeremy Daily who offered helpful critiques of the paper during its creation.

REFERENCES

- Ball, Jeffrey, David Danaher, Richard Ziernicki, Considerations for Applying and Interpreting Monte Carlo Simulation Analyses in Accident Reconstruction, SAE paper 2007-01-0741, 2007
- Bartlett, Wade, William Wright, Oren Masory, Raymond Brach, Al Baxter, Bruno Schmidt, Frank Navin, Terry Stanard, Evaluating the uncertainty in various measurement tasks common to accident reconstruction, SAE paper 2002-01-0546, 2002
- Bartlett, Wade, Conducting Monte Carlo Analysis with Spreadsheets, SAE paper 2003-01-0487, 2003
- Bartlett, Wade and Albert Fonda, Evaluating uncertainty in accident reconstruction with finite differences, SAE paper 2003-01-0489, 2003
- Christian P. Robert and George Casella, Monte Carlo Statistical Methods, 2nd Ed., Springer Verlag Publishers, ISBN-10: 0387212396, ISBN-13: 9780387212395, 2004
- Brach, Raymond M., Uncertainty in accident reconstruction calculations, SAE paper 940722, 1994
- Fishman, George, Monte Carlo Concepts, Algorithms, and Applications, 1st ed., Springer Verlag, 1996. ISBN: 978-0-387-94527-9, 2003
- Fonda, Albert G., The effects of measurement uncertainty on the reconstruction of various vehicular collisions, SAE paper 2004-01-1220, 2004
- Hammersly, J.M., and D.C. Handscomb, Methuen's Statistical Monographs: Monte Carlo Methods, John Wiley & Sons, 1964
- Kimbrough, Scott, Determining the relative likelihoods of competing scenarios of events leading to an accident, 2004-01-1222, 2004
- Kost, Garrison, and Stephen M. Werner, Use of Monte Carlo simulation techniques in accident reconstruction, SAE paper 940719, 1994
- McLaughlin, Michael, A Compendium of Common Probability Distributions, Regress+ Ver 2.3, Appendix A, accessed 12DEC2007 at http://www.causascientia.org/math_stat/Dists/Compendium.pdf, 2001
- Montgomery, D. and G. Runger, Applied Statistics and Probability for Engineers, 2nd Ed, John Wiley, 1999
- Moser, A., H. Steffan, A. Spek, W. Makkinga, Application of the Monte Carlo methods for stability analysis within the accident reconstruction software PC-CRASH, SAE paper 2003-01-0488, 2003
- Neiderer, P.F., The Accuracy and Reliability of accident reconstruction, Automotive Engineering and Litigation, editors Peters and Peters, Vol 4, John Wiley and Sons Inc, 1991, p257-304
- Schlaifer, Robert, Analysis of Decisions Under Uncertainty (Chapter 13), McGraw-Hill, Library of Congress Catalog Card Number 69-19203, 1969
- Slakov, G.A. and D.D. MacInnis, The Uncertainty of Pre-Impact Speeds calculated using Conservation of Linear Momentum, Proc. Canadian Multidisciplinary Road Safety Conference VII, 1991, pg242-255
- Taylor, Barry N. and Chris E. Kuyatt, "Guidelines for Evaluating and Expressing the Uncertainty of NIST Measurement Results," National Institute of Standards and Technology, Technical Note 1297, U.S. Government Printing Office, Washington D.C., 1994
- Tubergen, Renard G., The technique of uncertainty analysis as applied to the momentum equation for accident reconstruction, SAE paper 950135, 1995
- Wood, Denis P., and Sean O'Riordain, Monte Carlo simulation methods applied to accident reconstruction and avoidance analysis, SAE paper 940720, 1994
- Wach, Wojciech, and Jan Unarski, Determination of Vehicle Velocities and Collision Location by Means of Monte Carlo Simulation Method, SAE paper 2006-01-0907, 2006

APPENDIX: FURTHER EXCEL DISCUSSION

Excel utilizes a pseudo-random number generator, rather than a more elegant quasi-random number generator. This limitation is only pertinent if CPU time required for the calculations is of interest, and is only apparent for low trial numbers. The limitations of the pseudo-generator can be overcome simply by performing more trials.

Results of Excel's InvNorm() function have been shown to begin to diverge from a true normal distribution at the extreme ends of the distribution tail [http://www.puc-rio.br/marco.ind/quasi_mc.html, referenced 12/8/07]. This feature will be most important for analyses where the highly unlikely "tails" of the probability curve are of most interest, and will not meaningfully affect accident reconstruction analyses which are primarily concerned with the majority of the results (not the outliers).

Microsoft has stated that the random number generation algorithm in pre-2003 versions of Excel did not perform well in "randomness" tests when run out to more than a million values. Though this is not a significant issue with regard to this type of analysis, the newer editions ('03 and '07) have a more robust random number generation algorithm which passes those tests. More details can be found online at Microsoft's website: [<http://support.microsoft.com/kb/828795>].