The Seven S-Curves of Adelaide - Part 2

Posting Date: 26 September 2012

This is a supplementary article in a series discussing testing at Gorski Consulting to document roadway effects on potential loss-of-control collisions at curves on Adelaide Street , north of the City of London, Ontario, Canada. A previous article introduced the roadway and the characteristics of each curve through a series of photographs. This article, the second in the series, reviews the instrumentation and testing methodology used to obtain the data from drive-through tests at the site.

Gorski Consulting is a forensic, motor vehicle accident reconstruction firm that provides a variety of analyses both for clients involved in litigation as well as safety research examining the broad field of evidence surrounding real-life collisions. One of our interests involves the evaluation of roadways in relation to loss-of-control collisions.

Recently Gorski Consulting has been undertaking several research projects examining curves on rural highways and how these relate to the incidence of loss-of-control collisions. One of the studies that will be discussed in this series of articles involves testing conducted along seven S-curves that exist along a 12-kilometre distance of Adelaide Street, north of the City of London, Ontario, Canada. Testing was conducted in September of 2012 wherein a vehicle was driven along the study area and data was collected with multiple video cameras and an accelerometer. This article discusses the details of the instrumentation and methodology used to gather the data.

Instrumentation

Our process of data acquisition involved using a test vehicle which was instrumented with six video cameras at the following locations:

Camera 1: Canon SX200 was hand-held and used as the "master" camera that enabled all the videos to be synchronized. Matching the audio track from each video was found to be the easiest method of synchronization. In some instances where a camera did not record a recognizable audio track the cameras also had to be synchronized using the video track. In such instances the master camera was brought along side each camera and both cameras videotaped the same action, such as a rapid hand movement. During the video editing the video from each camera was synchronized by lining up the rapid hand motion in the time line so both videos showed the same motion at that same time-code. This process was repeated for every camera until all were synchronized with each other.

Camera 2: A Sony MiniDV camcorder was mounted to the instrument panel and pointed forward through the windshield to capture the length of the road as the vehicle approached it.

Camera 3: A GoPro22 camera was mounted directly in front of the speedometer/tachometer and other instruments so that the status of these gauges could be documented.

Camera 4: A GoPro24 camera was mounted in the centre console area and pointed down onto the face of an iPhone which was displaying the output from the iPhone's accelerometers.

Camera 5: A GoPro28 camera was attached to the right portion of the roof near the vehicle's B-Pillar and pointed to the right to capture the location of specific references such as curve warning signs. When a curve warning sign was located in the middle of the view of this camera this was the point when documentation of the accelerometer's data was begun.

Camera 6: A GoPro23 camera was mounted to the left end of a lateral bar that itself was attached to a bike rack at the back of the test vehicle. The camera was pointed forward and downward toward the left rear wheel and left rear tail-light of the test vehicle. While the test vehicle was in motion this camera documented the lateral distance between the left wheels of the vehicle and the centre-line of the highway. This view allowed the documentation of the change in lateral position of the test vehicle on the road as this was an important factor when evaluating the lateral acceleration of a vehicle through a curve. This camera also captured the status of the illumination of the left brake light of the vehicle.

The photo at the top of Page 3 shows Camera 2 as it was anchored to show a forward view through the windshield. The bottom photo on Page 3 shows the GoPro22 camera mounted to face the speedometer and other gauges of the instrument panel.

The two photos on Page 4 show GoPro24 as it was anchored to face downward onto the face of the iPhone. The iPhone was anchored to the centre console with velcro.

Anchorage of GoPro24 was difficult as adhesive tape would not stick to any surfaces of the test vehicle. Eventually a suction cup anchor was taped to a block of wood and this block was force fitted into the contours of the cup-holder and this provided satisfactory stability of the camera.

As can be seen in the photos of the face of the iPhone there was considerable reflection of sunlight that made it difficult for the camera to record the numerical values that were being displayed. The problem was rectified by placing a small cardboard box over top of GoPro24 and the iPhone so that also all of the reflections were eliminated. The photos on Page 5 show the box as it was positioned for testing.


4 5 6 7
RPM
X1000
0 20 40 60 80 100 120 140 160 180 200

The two photos on Page 6 show how GoPro23 was anchored to the lateral tube that was anchored to a bike rack at the rear hitch of the test vehicle. The photo on page 7 shows how GoPro23 enabled a view of the left tires in relation to the centre-line of a roadway.


BUICK

Testing Procedure

In an initial run the test vehicle was driven northward along Adelaide Street using a normal manner where the driver did not select a particular speed and was not looking a the speed reading on the speedometer. The cruise control was not activated. As the driver approached each curve he simply performed a typical action that he would do under normal driving conditions. Once the vehicle passed through Curve 7 at the north end of the road the vehicle was turned around and the same driving procedure was performed travelling southward through the curves.

Upon review of the test run it was noted that the vehicle typically strayed laterally within the travel lane while passing through the curves. It was noted that the test driver travelled at a speed that was substantially higher than the speed limit and that the speed changed depending on the conditions on the road.

In a second round of tests the driver set the vehicle's cruise control to 90 km/h, or 10 km/h above the posted maximum speed limit. The driver then maintained that speed, without braking, while travelling northward through the full set of curves. The test driver also purposely attempted to keep the vehicle centred within the lane even while passing through the curves. The driver also purposely attempted to reduce the speed and magnitude of an any steering motions so that the changes in steering action was minimized to what was essential to keep the vehicle in the centre of the lane with a minimum of corrective steering. The test drive was then completed travelling southbound. Then the procedure was repeated, travelling northbound and then southbound, for a second time.

The videos from the six video cameras were then imported into the Adobe Premiere video-editing program and assembled so that all the views were synchronized and displayed on the computer screen. By examining all the views all at once on a computer screen it is possible to detect a variety of useful facts. However this is also demanding on the computer and will sometimes cause the program to freeze as it attempts to catch up with the demands. However, once all the videos have been synchronized alternatives can be developed where only a selection of views are used in any single video project. For example, it is not uncommon to run several computers at the same time, with a video project open and displayed in separate monitors. Thus each monitor might show a project containing just two videos. The analyst can move from one monitor to another while examining the same time of a particular test only from different camera views.

This was the case when we analysed our testing to extract the value of the lateral acceleration displayed on the iPhone. It was not necessary to view all the videos on the same screen therefore a second Premiere project was created using the videos from just two cameras. The display of the face of the iPhone (View from GoPro24) was used along with the view of the left side of the test vehicle (View from GoPro23). This minimized processing demand on the computer and allowed for quick repositioning along the video tracks.

The XSensor iPhone application was used to display the iPhone's lateral acceleration values. The developer of the program recommends that the recording of the data be started with the push of a button and the data is then stored in a file which can be sent to an e-mail address for further processing. Unfortunately, this recording is "blind" in that it can be difficult to match a specific acceleration value to the vehicle's precise location on the road. In our testing we were interested in knowing, as close as possible, what specific feature of the road caused the specific lateral acceleration reading. Therefore, in this initial testing we opted to ignore the developer's recommendation and simply documented the raw readings that were recorded by the camera that was pointed at the face of the iPhone.

Reading the raw data from the iPhone's display can be problematic in that the numbers are displayed rapidly and sometimes they appear jumbled up or smeared. Sometimes an after-image of a previous display still exists while the image of a new number is displayed over top of the old value. Sometimes it becomes difficult to decipher what the smudged values are and what actual reading should be the correct value. While this method has these limitations they are not insurmountable if a pre-planned procedure is consistently followed.

The video recorded by our cameras was at a 30 frame-per-second frame rate. We opted to document readings at every 5 frames such that there would be 6 documentations every second.

The procedure we followed was to advance to the 4th frame and observe what reading was visible there. Sometimes that reading was smudged and difficult to decipher. We then advanced to the 5th frame and attempted to decipher the reading. Knowing the value at the 4th frame sometimes helped us recognize that its after-image was imprinted onto the 5th frame and so this helped in those cases where it was difficult to distinguish the value in the 5th frame. When the images were blurred in both the 4th and 5th frames we then advanced to the 6th frame and recorded that value if its image was clear. These values were documented in an Excel spreadsheet. So, although our attempt was to use the value contained in the 5th frame, when that was not possible we recorded the value of either the 4th or 6th frame depending on where the value was most legible.

One might conclude that there is a certain error created when the recorded value is not taken from precisely the same frame each time. However, at a speed of 90 km/h the test vehicle travels about 25 metres per second therefore the distance travelled between two frames is just over 4 metres, which is not a long distance. When this error is spread over several drive-through tests its effect is smoothed out and is not a significant problem. The benefit in using this method is that one can be more certain of the precise portion of the road that caused the reading. That certainty is less so when the data is taken from a "blind" file and the analyst has to guess at the precise portion of the road that caused a reading. There may still be ways to overcome some of these issues but for the present testing this was the procedure that we decided to use.

As an example, the charts on Pages 11 and 12 show the data that was documented from the test vehicle's first and second northward runs through Curve 1. The chart on Page 13 shows the data when the results of the two tests are combined and averaged.

Looking at the data from Run 1 on Page 11, one can see that there were 140 points of data documented. The documentation started when the test vehicle passed by the curve warning sign. So we can see there are about 60 data points on the approach to the curve where the values are generally and slightly above zero. This is understandable because there is a cross-slope of about 2 percent that is designed into a travel lane so that rain will flow off the road surface. If there are six data points in a second then these 60 data points represent a time of about 10 seconds. But also note that, at the cruise speed of 90 km/h, or 25 metres per second, this segment of 60 data points represents a distance of about 250 metres. So, even without taking a measurement of the location of the curve warning sign, we can estimate that it is located about 250 metres prior to the curve.

Continuing with Run 1, we can also see that there is an arc of negative accelerations registered from about observation 60 to about observation 100, or about 40 data points. On Page 14 and at the top of Page 15 we have displayed three photos of the curve. Note that the first portion of the S-curve is a curve to the right so these negative accelerations are generated when the vehicle's motion is in a right curve. Again, using the previous methods of calculation we estimate that the 40 data points represent a time of about 6.7 seconds and a distance of about 167 metres. So again, even without measuring the curve, we can estimate that the length of the first portion of the S-curve is 167 metres.

Continuing with Run 1, we can see how the vehicle transitions from the right curve to the left curve by looking at the steepness of the data line and whether there are any data points, for a short distance, that are close to zero. It can be seen that the data line is quite steep as it moves from about -.18 g at about observation 90 through to about .18 g near observation 107, and there is no appreciable discontinuity in the steepness of that line, so it tells us there is no distance separating the right and left curves but that one curve proceeds directly into the other.

We then look at the second portion of the S-curve which is a left curve that exists from about observation 100 to the end of the documentations at observation 140. So, again, this tells us that the 40 data points for the left curve are number as the 40 data points for the right curve and therefore the two curves are about the same length of about 167 metres. So the total length of the S-curve is about 334 metres and it took our test vehicle about 13.4 seconds to pass through it.

The magnitude of the lateral acceleration values tell us whether there are any unusual features of the curve that might cause problems. For example, traffic engineers like to have the lateral acceleration reach a maximum of about 0.15 g. If testing reveals that it is higher then alternatives could be pursued such as reducing the speed limit so that the lateral acceleration is reduced. In our example of Run 1, we can see that the maximum acceleration generally stays below the 0.15 g value, even though the test vehicle is travelling at 10 km/h above the speed limit, so the design of curve and speed limit assigned to it appears reasonable.

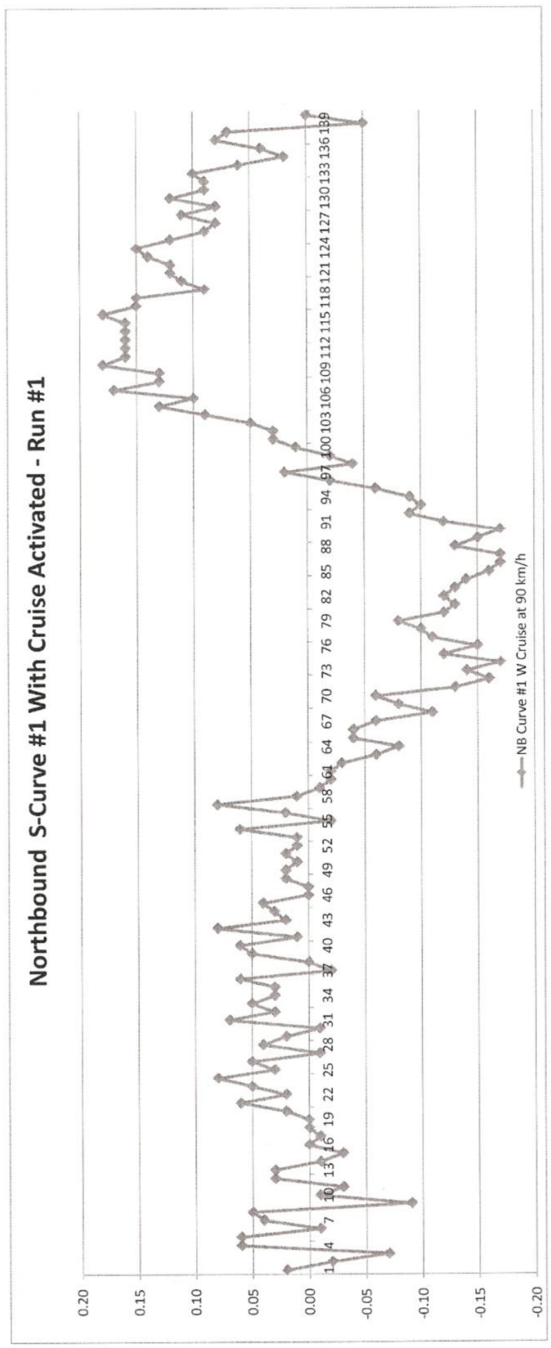

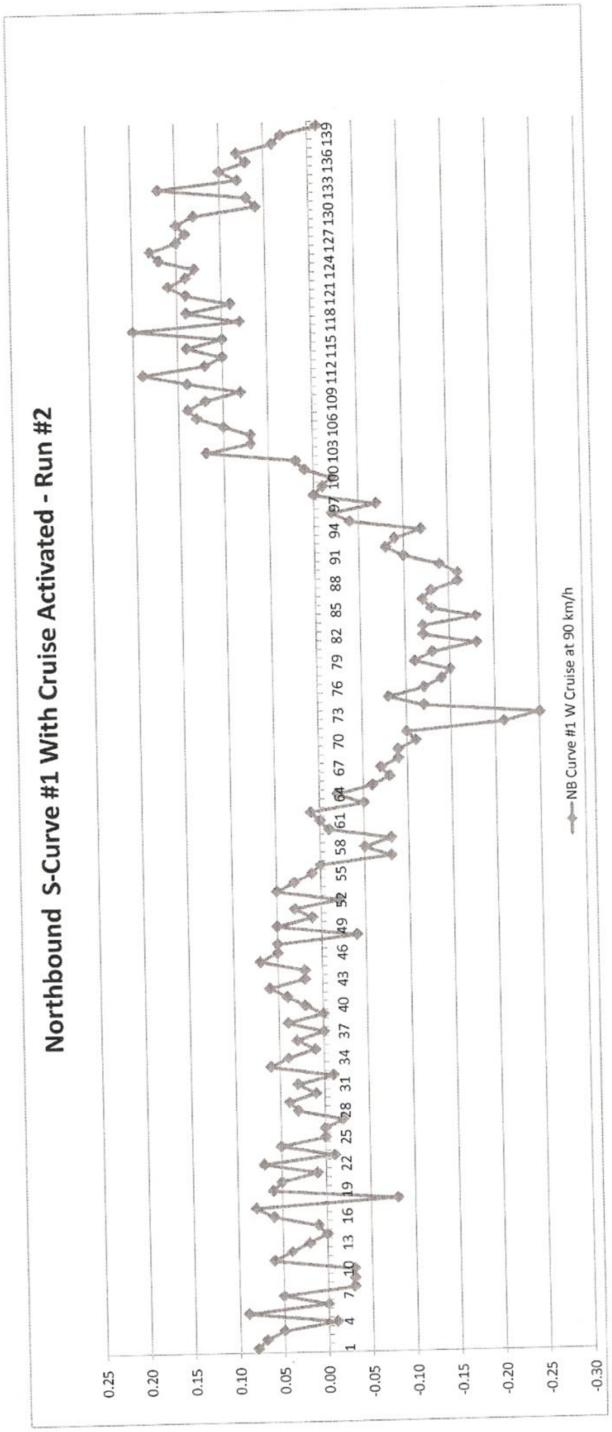
Northbound S-Curve #1 With Cruise Activated - Run #2
NB Curve #1 W Cruise at 90 km/h

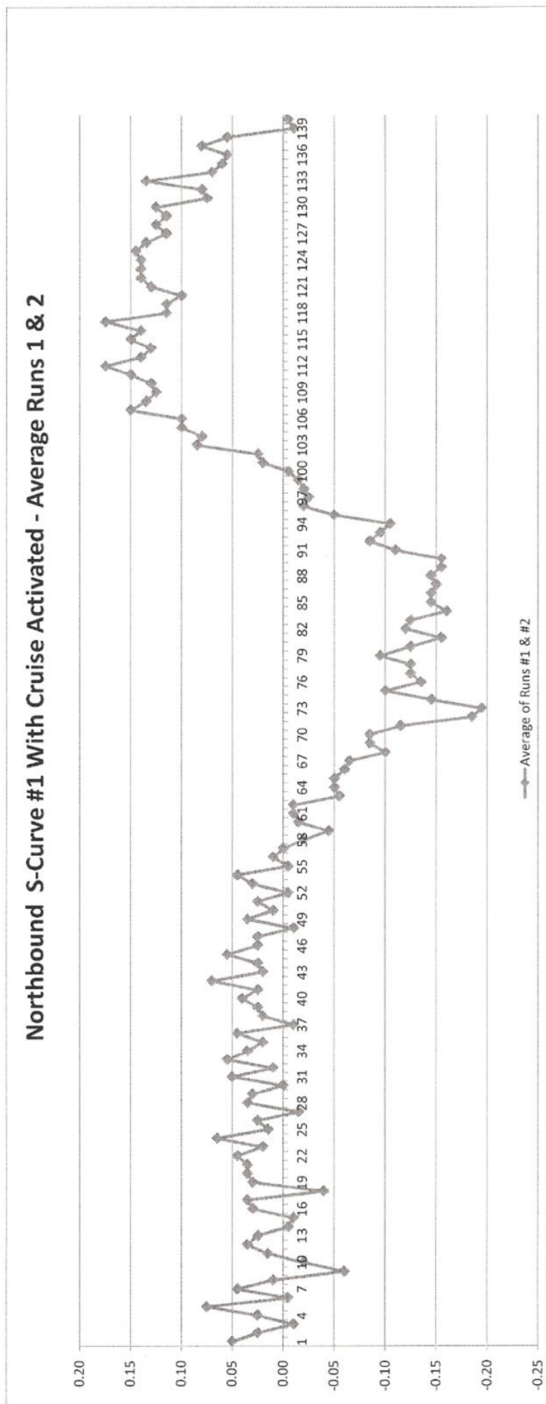
Northbound S-Curve #1 With Cruise Activated - Average Runs 1 & 2
0.20
0.15
0.10
0.05
0.00
-0.05
-0.10
-0.15
-0.20
-0.25
1 4 7 10 13 16 19 22 25 28 31 34 37 40 43 46 49 52 55 58 61 64 67 70 73 76 79 82 85 88 91 94 97 100 103 106 109 112 115 118 121 124 127 130 133 136 139
Average of Runs #1 & #2

But let us remember that these are the results from just a single test. The chart shows how the acceleration values jump up and down and we cannot be certain where they might fall if additional tests were performed. We can look at the chart for Run 2 on Page 12 to see if there are any obvious differences compared to Run 1. We can make a further comparison by looking at the average of the two runs in the chart on Page 13. Although some possibilities could exist, in our opinion, there is still too much variance caused by the idiosyncratic nature of each run to conclude that there is any concern about the design of this curve. We would recommend that the data from at least five runs be combined and averaged in a chart like this before one should begin to believe in the accuracy of that data.

So let us review the merits of this procedure. Multiple video cameras allow for a very good documentation of the features of the curve. The camera pointing forward through the windshield can be adjusted to match the approximate view provided to the driver (Alternatively an additional helmet camera could be mounted to the driver's head). The drive-through test using a set cruise speed takes away much of the variability in the lateral accelerations that would normally hide the true effects caused by the curve geometry. By consistently driving in the middle of the lane the test driver also reduces the effects on the lateral acceleration caused by changing the lateral position of the vehicle within the lane. After several runs are performed in this fashion, a plot is generated that approximates the lateral accelerations that are generated by the curve...and now the fun begins.

After we understand how the roadway affects the vehicle we can then bring in a variety of drivers to perform their various motions through the curve and we can study how these drivers cause differences in the data. How would an experienced driver travel through this curve? How about an inexperienced driver travel through the curve? How about a driver who is forced into a conversation with an experimenter who is seated in the vehicle with the driver? There are numerous possibilities of study. And this is an important issue because, as new technologies have reduced the number of fatalities in various collision configurations, vehicle loss-of-control becomes even more prominent in the collision statistics.

Loss of control leads to those head-on collisions just like it leads to rollover collisions or impacts with trees and poles. While we treat each of these collision types as separate entities many of them actually originate because a driver has lost directional control of a vehicle. We know that many such loss-of-control collisions occur on rural highways and often where curves are involved. So this research and testing is of considerable relevance.

In the next article in this series we will focus on more of the results from our testing and we will compare each of the seven S-curves to each other. We guarantee that, if nothing else, the results will be eye-catching. So don't be distracted, and pay attention to your driving.

Gorski Consulting
London, Ontario, Canada