

Design of a Driving Simulation Sound Engine

David A. Heitbrink, Steve Cable

The University of Iowa, NADS
2401 Oakdale Blvd.
Iowa City, IA 52242
(319) 335-4785
(319) 335-4658 fax
Contact email: dheitbri@nads-sc.uiowa.edu

An integral part of any interactive simulation is sound. Traditionally this has been done with a high cost proprietary costume hardware/software solution. Recently we decided upgrade our sound system, and decided to go with a system based on off the shelf PC sound hardware. PC based sound cards have advanced rapidly over the last 10 years. They now offer 3D sound spatialization, reverberation, and many other important functions. Also, a large number of API's have come out to support sound engine development. We choose FMOD by Firelight Technologies. This allowed us to focus on accurate synthesis of the main components of interior car noise (tire/road, engine, wind). We mostly used a wave table approach. Given this, we also decided to record new high quality engine and road/tire sounds. Over all we were able to put together an effective sound engine, in a relatively short period of time.

David Heitbrink is a Software Engineer at the National Advanced Driving Simulator. His primary responsibilities include maintaining the scenario-authoring tool (ISAT), development and maintenance of the current audio subsystem, and other software development tasks. Mr. Heitbrink received his Bachelors of Science from the University of Toledo in Computer Science and Engineering. He continued to receive is Master of Science in Engineering at the University of Toledo in the spring of 2005. Mr. Heitbrink is originally from Cincinnati and lived there until he moved to Toledo upon entering the University of Toledo.

Stephen Cable was the Senior Systems Administrator and Audio Engineer at NADS, and has recently taken a position elsewhere. He has ten years of audio experience and thirteen years of experience in systems administration. He has logged well over eight thousand hours as the lead recording/mastering engineer for hundreds of musical performers. His academic and professional career has focused on positional audio, digital and analog recording techniques, waveform analysis and platform integration, as well as network and systems programming.

Design of a Driving Simulation Sound Engine

1. Introduction

One of the most overlooked aspects of driving simulation is sound. Audio queues affect speed judgment, urgency, operator performance, alertness, and fatigue. As we drive we use many audio queues for making decisions, such as using the pitch of the engine sound to decide to change gears. In a simulation out of place and inaccurate sounds can also distract a participant and negatively affect results.

At the National Advanced Driving Simulator we have had, since the start of operations, a functional sound engine capable of simulating road, engine, and wind noises. However after some time the original proprietary hardware/software solution had become unreliable and difficult to support. Given these problems, we decided to build a new sound engine.

In the past, sound engines have mostly been based on proprietary DSP based solutions. Consumer PC based hardware simply did not have the real-time signal processing capabilities to support a 3D sound engine. However over the last decade the capabilities of consumer sound hardware have greatly increased. Today, 3D sound has become a standard feature of most computer video games; this has led to a large number of 3D sound API's.

Given the current capabilities of PC based hardware, and the plethora of 3D sound APIs, we decided to design a new sound engine in-house using PC based hardware. In terms of software APIs we choose to go with FMOD, by Australian based Firelight Technologies (www.fmod.org). FMOD provides functionality to handle 3D sound placement and Doppler shifts. Also FMOD provides a node based DSP engine that provides such important functions as pitch shifting, white noise and sine wave generation, signal filtering, and supports geometry based wave tracing to calculate reverb, obstruction and occlusion.

2. Sound Synthesis

In-car noises come primarily from three separate sources: road/tire noises, wind, and engine/power train. Most of the inputs into our sound engine are based on accurate vehicle dynamics simulation [7]. For example the engine model uses RPM and torque values from an engine simulation. Without a good physical simulation of the driving environment it would be impossible to create a good audio simulation.

Most audio simulators use a sample based synthesis approach, also known as a wave table approach. This technique uses a collection of sound samples that are manipulated and/or mixed in some manner to simulate some sound. An integral part of our simulation is the use of graphs. The graphs are simply a lookup table that matches input values to output values. Points in between defined points on the graph are linearly interpolated. The graph files are loaded at run-time; this allows us to change the characteristics of our sound engine without rebuilding the sound engine.

One important thing to keep in mind is the audio synthesis process will only be as effective as the speakers used to play the audio. Since most of the interior car noise is in the low frequency range [5], it is important to have a good low frequency response in any speaker system used. Also, it is important to properly set the speaker equalization, delays and other relevant settings to insure proper playback, M. Blommer [6] discusses this in great detail. We were able to reduce the number of discrete components used in the solution set for by M. Blommer by mounting our audio engine in the cab, and directly connecting it to a Crown CTs 8200USP/CN amp, that is also mounted in the cab. The Crown amp contains several DSP units that were able to provide equalization and hardware delays.

2.1 Road Noises

Road noises generally make up the majority of in-car noise. Road/tire noises are complex phenomena. Sources of road/tire noise include: tread impact, air pumping, tangential motions, and adhesion. Tread impact is where the tread impacts the surface of the road. The tangential motions are sounds created by the tire slipping across the surface of the road, and adhesion is where the tire treads sticks to the road and is lifted off. The air pumping sounds are created by air being compressed and released from the grooves in the tire [2]. Altogether, these make a difficult phenomenon to simulate.

Our road noise simulation uses a sample-based synthesis approach. Each tire is modeled (one sound source per tire). The samples we use were taken from real drives taken at 30 mph, and 60 mph. The 30 and 60 mph samples are cross-faded depending on speed. The sounds were recorded close to the source (tires) during a real world drive. A low pass filter was used to isolate the tire/road sounds. This approach gave us pretty good isolation of the road noise, and is simpler than the approach given by D. Riemann [8], where a speed dependant FIR filter is applied to a recording from a car coasting from maximum speed to standstill. In addition, our approach allows for us to record much longer samples at a much more consistent speeds. In addition each sound sample is played with a randomly staggered starting point as to ensure all the samples are not playing in phase.

Each road type has a sample set, a pitch modifier, and amplitude modifier. Multiple road types use the same sample sets, but simply have different pitch and amplitude modifiers. This allows us to use the same sample sets for similar road surface types.

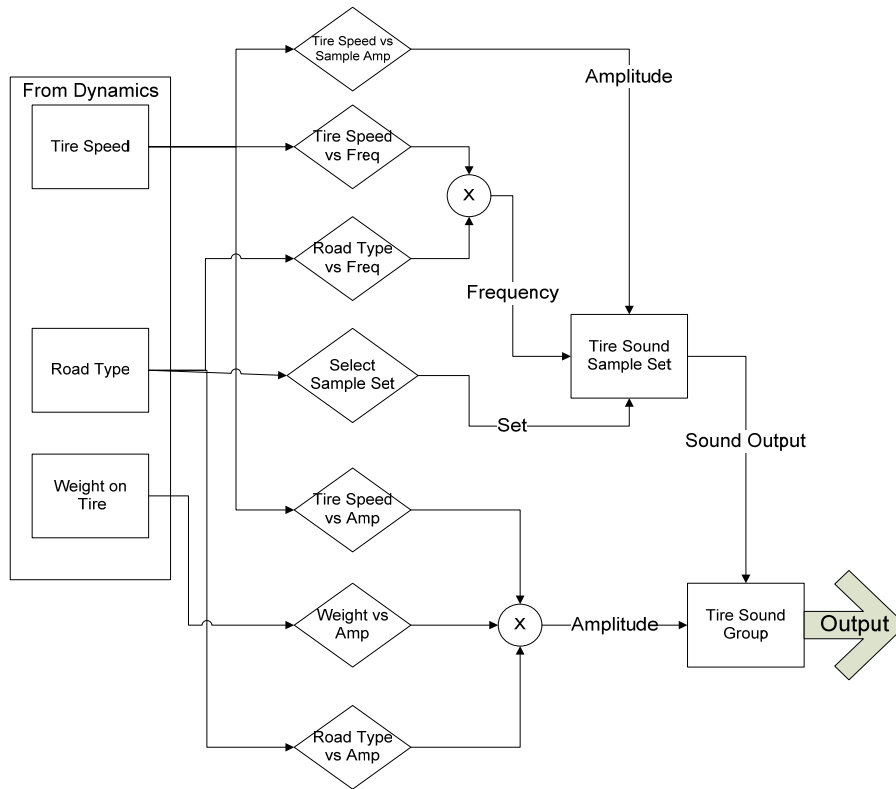


Figure 1 Overview of Road/Tire Noise Generation

The amplitude is set by multiplying the output from a speed vs. amplitude chart and a weight on tires vs. amplitude chart. Additionally, a slight pitch shift based on speed is applied to each sample based on speed.

Tire squeal sounds are implemented by taking the tire slip percentage and looking up the value on a tire slip versus amplitude graph. No frequency shift is performed for the tire slip sound.

2.2 Engine Sounds

Engine sounds are generally not as loud as road noises when driving, although participants tend to notice engine sounds more. Engine noises tend to be more interactive. Also in many instances such as operating a heavy truck, the driver tends to use audio queues for deciding when to change gears. It is important that engine model perform fairly accurately, and it is important that the engine model actually match the vehicle that is being simulated. If the participant is driving a heavy truck it should actually sound like a heavy truck.

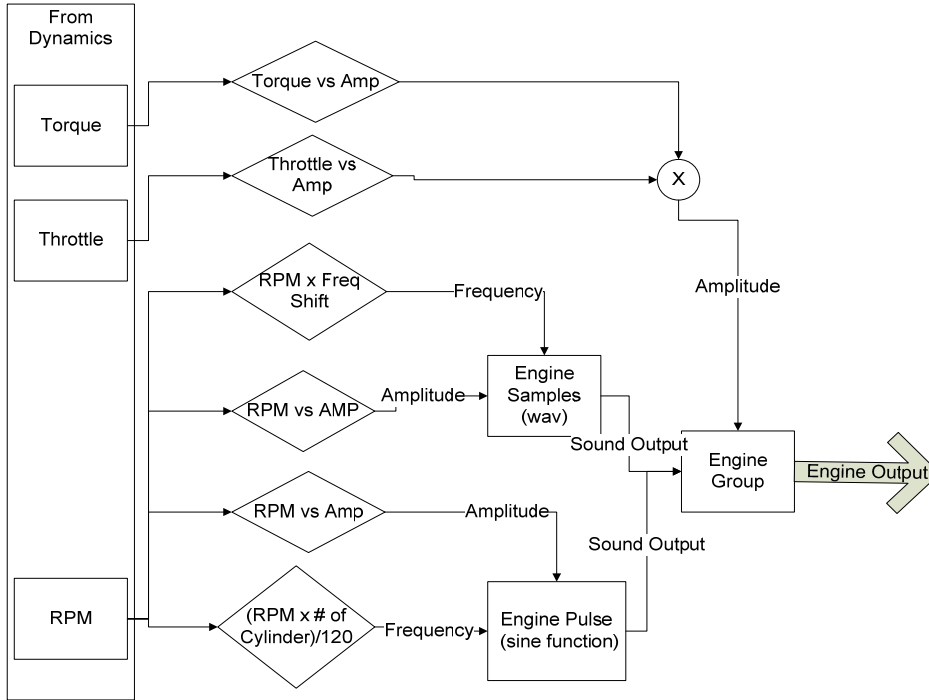


Figure 2 Overview of Engine Sound Generation

The engine sound synthesis mostly uses a wave table technique; although, a pure tone that matches the pulse frequency of the engine is also added. Our samples were recorded at set RPM levels. During playback, the samples are cross-faded depending on their own specific RPM level. In other words, if the current RPM level from the engine model is outputting 1650 RPM, samples taken at 1500, and 2000 would be played at 70% and at 30% respectively. Each sound sample has its own chart that measures RPM input vs. sound output. The sum of the engine outputs from each sample should sum to 100%. Each individual sample would also be frequency shifted by a factor of:

$$\frac{\text{Engine rpm level}}{\text{rpm level of recording}}$$

Recordings taken from the manifold and exhaust were used. The manifold samples tend to contain more high frequency components, while the exhaust sound tends to have more of a lower frequency sound. The mixture between manifold and engine sounds was done “by ear” for each cab.

As we can see in Figure 2, a pure tone (sine wave) using the pulse frequency is added. The pulse frequency for a 4 cycle piston engine the pulse frequency is equal to:

$$f = (\text{number of cylinders}) * \frac{(\text{revolutions per minute})}{120}$$

The amplitude of the pure tone is determined through an RPM versus amplitude chart. The pure tone is used to give an additional tone where the mixture of samples is found

lacking. For instance some of the older sound sample sets had a lot of the low frequency sound cut out. So for lower RPM levels a large pulse frequency tone was added in, and as RPM levels increase, the pulse tone amplitude decreases. The pure tone was also a member of the engine group so its amplitude would be further modified by modifications made to the group.

Overall the pitch of the engine sound is affected only by the RPM level, and should not affect the amplitude of the engine sound. The amplitude of the engine is controlled by a combination of throttle position and the torque output. Both the throttle position and engine torque are compared to a lookup tables for a multiplier. If the torque output is not available a window function is used to filter the throttle position to prevent the amplitude from increasing too fast.

2.3 Passing vehicle noises

Passing vehicles sounds are composed of three components: aerodynamics noises, engine/power train noises, and road noises. These sounds also have a Doppler shift applied to them. Fortunately, FMOD has functionality for handling the Doppler Effect. We only had to only provide the speed, direction, and position of the listener and sound source and FMOD would perform the necessary frequency adjustments.

For simplification, a much simpler model for the road and engine sounds is used. For the engine sound, each vehicle is given an engine sound and an adjustment value for the engine volume. The engine samples were taken from vehicle exhaust and generally match the vehicle type. The amplitude of the engine sound is adjusted using the speed of the vehicle.

Tire sounds are handled similar to engine sounds. A simple generic road sound is used passing car sound. The volume of the road sound component is adjusted by the size and speed of the vehicle.

Aerodynamic sounds produced by automobiles are a very complicated subject. They are complicated fluid dynamic events that do not occur by themselves; as on-coming cars pass each other their turbulent wakes interact. Modeling these systems in real-time would not be practical. Given these limitations, we decided to use a 10 second pre-recorded sample from our wind sound simulation taken at 35 mph using our Jeep Cherokee model.

Aerodynamic sounds are adjusted for the speed and size of the vehicle. The speed for the wind sound is calculated as the relative speed of the vehicle from the participant's vehicle. A "sound cone" is used; it is a directional sound source supported by FMOD. The programmer can specify a degree and width that the sound is projected from the sound source. The sound cone is placed slightly ahead of the passing car and projected towards the passing car. This gives the effect of the interaction of the turbulent wake of the passing vehicle with the participant vehicle.

2.4 Wind Sounds

Wind sounds come from aero-acoustic interaction of the vehicle with the atmosphere. In a typical passenger car, these sounds tend to be transmitted into the cab from the windshield, passenger windows, and from the front A-pillar. Automotive designers tend to spend a lot of time trying to reduce the amount of wind noise a car produces as well as ways to reduce how much of the noise is transmitted into the cab. Automotive designers tend not only worry about the amount of noise that is produced, but also the harshness. A highly irregular wind noise that frequently changes pitch and amplitude, caused by say buffeting, can cause a much greater level of annoyance to the driver than a noise with a consistent sound.

Originally we planned on doing a sample based synthesis. We decided to give up this approach due to the difficulty of isolating wind sounds. Typically, recordings of in-car wind sounds are done within a wind tunnel. Unfortunately such facilities were not available to us. We had considered using a shotgun microphone pointed at the windshield or side windows during a real world drive with a passenger car.

We gave up this idea after we discovered we could create a reasonable sound with randomly generated noises. For the random noise we simply took a white noise generator and applied a low-pass filter with at around 600Hz with the Q value set at 1. The Q value affects the steepness of the roll-off slope. The higher the Q value the steeper the slope. FMOD's low-pass filter supports Q values from 1 to 20. The frequency of the low pass filter is gradually increased with speed of the vehicle, up to a maximum of about 700Hz. Also a random value between -50 to 50 is added to the low-pass filter at the rate of 120 Hz. This randomization helps to simulate turbulence, and harshness.

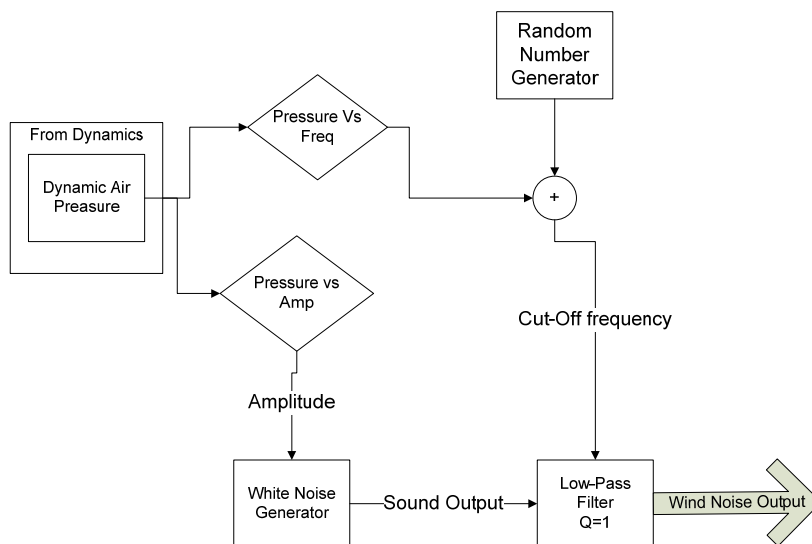


Figure 3 Overview of Wind Noise Generation

The amplitude of the wind sound is increased at an exponential rate with the speed of the vehicle. Through wind tunnel testing with a typical passenger car, Peric [1] found that in-car wind sounds generally increase by about 4db per 20 km/hr or approximately 12.5 mph. They also found that a car traveling at 120 km/hr (74.5 mph) would see a sound pressure level of about 73 db (A). We have generally tried to match that value.

3. Sound Recording Process

Once the wave table technique was chosen, creation of a new sound set became integral in the NADS audio engine development. Since the end product of any sound reproduction system can only be as good as the original recording, emphasis was placed on creating the highest quality samples possible.

The samples that would be required for the new system broke down naturally into two main areas of interest: cab-specific sounds (such as engine and power train) and external sounds (such as road noise, autonomous vehicles, and environmental audio). While the cab-specific recordings would be tailored to specific vehicle types, external sounds would be generic recordings that could be used in any scenario.

We tended to use stereo miking, as close-miking in mono can result not only in an unnatural-sounding magnification of detail, but also in tonal imbalances such as overly bright tones, sibilance, boominess.

3.1 Cab-Specific Sounds

The majority of the cab-specific audio recording was targeted at capturing engine sounds. Using two pairs of microphones and a portable digital recorder, samples were taken at the exhaust pipe and engine manifold of specific target vehicles. These two sample sets would be combined at playback to create a realistic reproduction of engine noise.

Ideally, engine samples would be recorded using a vehicle mounted on a load cell dynamometer. A commercial dynamometer was not available for our initial recording sessions, so all engine sounds were recorded using a stationary vehicle in neutral, with load to be induced artificially at playback. RPM levels for individual cabs were captured at 500 RPM intervals, from idle through 3500 RPM, using two recording “nodes” (each comprising of a discrete left and right channel) and a NAGRA-DII portable digital recorder.

At each recording node, a pair of matched Neumann KM 184 microphones were set up in an ORTF configuration. ORTF combines the characteristics of volume difference (provided as sound arrives on- and off-axis at two cardioid microphones spread to a 110 angle) with timing difference (provided as sound arrives at the two microphones spaced 17 cm apart). Since the cardioid polar pattern of each microphone rejects off-axis sound,

less of the ambient environmental characteristics are picked up, thus creating an accurate representation of isolated sound in less-than-ideal surroundings.

The mixing and mastering phases involved transferring all four channels via an AES/EBU link to an audio workstation. Using an off-the-shelf audio editing program, the left and right channels were combined and normalized into two sets of looped, stereo WAV files; one set for manifold sounds and another set of corresponding exhaust sounds. As necessary, equalization was applied to compensate for anomalies.




Figure 4 The Signal Path for Recording Engine Sounds




Figure 2 An Illustration of the ORTF Stereo Recording Technique

3.2 External Sounds

External sounds proved to be very straightforward to record. Depending on the source and circumstance, multiple techniques were used.

The large majority of environmental recordings (such as ambient sounds) were recorded with various omnidirectional microphones, since specific positioning and orientation information were not required. Recorded channels were transferred to the DAW, where multiple channels were mixed, bounced down to a single track, looped, normalized, and saved as a stereo 48 kHz WAV file.

Specific external sounds (such as tire/surface interaction, passing vehicles, pedestrians, etc) were captured, again utilizing two KM 184 microphones in an ORTF configuration to maintain a good stereo image. Similar to the recording of engine sounds, specific external audio was recorded on two discrete monaural channels, looped (if necessary), normalized, and mixed into a single stereo 48 kHz WAV file.

4. Results

At this point most of our testing/validation consisted of spectral band matching with real world recordings and subjective jury evaluations. In the future we plan on taking more “reference” recordings of real world drives to help with validation, and further tuning our audio system. Unfortunately we could not do a good comparison with our old system as it had a hardware failure shortly after our new system was completed.

We were able to a relatively good set of reference recordings for a heavy truck from a recent trip to a local trucking company. As we can see in spectral analysis in Figure 5, we have been able to closely match the spectral output from a real world recording, although we do have a deficit in the 20Hz spectral range. Also, looking at Figure 6 we were able to produce a smooth acceleration from 0 to 2180 RPM. The frequency on the left hand does level out for a short period of time; this does correspond to the RPM level flattening out for a short period from the dynamics model. Since truck drivers tend to use engine tone when deciding when to shift, it is important that we do match the tone pretty closely.




Figure 5 Heavy Truck Engine at 1800 RPM, Right Simulated Model, Left Real Life Recording

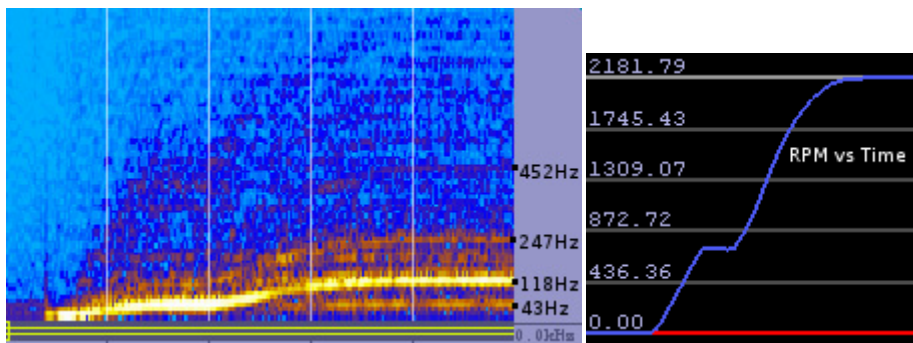


Figure 6 Heavy Truck Model, Engine Acceleration Throttle is at 30% (Left Spectrogram, Right Engine RPM vs Time)

5. Conclusion

In conclusion, we were able to construct a reasonable accurate sound engine. We found that we were able to reproduce engine, road, wind and passing vehicles sounds that matched to real world events. In the future we would like to utilize FMOD's geometric engine so we can handle such events as driving in a tunnel. We would also like to add road seams and other common surface defects. One advantage to our new system is much more configurable, for instance we can change balance between source types (road, engine, etc) during execution. Overall we were able to produce a reasonably effective solution in a short period of time.

References

1. Peric, C., Watkins, S., and Lindqvist, E. *Wind turbulence effects on aerodynamic noise with relevance to road vehicle interior noise* Journal of wind Engineering and Industrial Aerodynamics 69-71 (1997) pages 423-435
 2. Bernahard, R., and Wayson, R., *An introduction to Tire/Pavement noise of Asphalt Pavement* <http://www.asphaltalliance.com/library.asp?MENU=555>
 3. Fox, C. *Utilizing Omnidirectional Microphone Arrays for Surround Field Recording*. The proceedings of the 19th Annual AES International Conference, Schloss Elmau, Germany (2001).
 4. Watkinson, J., *The Art of Sound Reproduction*. Woburn, Massachussets: Focal Press (1998).
 5. Higaki, T., Ahuja, K., and Funk, R., *Dissection of Automobile Interior Noise Spectrum with Emphasis on the Infrasonic Region* Proceedings of the 11th AIAA/CEAS Aeroacoustics Conference, Monterey, California (2005)
 6. Blommer, M., and Greenberg, J., *Realistic 3D Sound Simulation in the VIRTEX driving simulator* DSC North America 2003 Proceedings, Dearborn, Michigan (2003)
 7. Salaani, M., Heydinger, G., and Grygier P., *Heavy Tractor Trailer Vehicle Dynamics Modeling for the National Advanced Driving Simulator* SAE Paper 2003-01-0965 (2003)
 8. Riemann, D, and Sottek, R., *Interactive Auralization of Powertrain Sounds Using Measured and Simulated Excitation* SAE Paper 2007-01-2214
-