WHITE PAPER

Autonomous Vehicle Radar:

IMPROVING RADAR PERFORMANCE WITH SIMULATION

As the world draws near to the reality of fully autonomous automobiles and transport vehicles, there is a great deal of focus on the development of artificial intelligence (AI), machine learning and rapid automated decision-making. These solutions are aimed at duplicating — and improving upon — a human's ability to perceive the world around them, and at efficiently and safely guiding a vehicle through a journey.

A critical aspect of autonomous vehicle control is the continued evolution of vehicle sensors — the eyes and ears of the control system that perceive the operational characteristics of the vehicle and the environment around us. While the AI and decision-making systems must plan the vehicle trajectory and response to the environment, the sensors must feed the control systems executing those algorithms with accurate data on the current and developing state of the vehicle's surroundings. The development of accurate algorithms depends in part upon the availability of accurate sensor data.

By **Shawn Carpenter**,
Product Manager, High Frequency/Electronics Business Unit (EBU), Ansys

Four major classes of vehicular sensors provide the lion's share of environment sensory data for an autonomous vehicle. Visual spectrum cameras provide data like that of our own eyes — motion sensing, depth perception, and, of course, color for sensing things like stop signs, emergency lights and stoplights. Laser-ranging devices (LIDARs) are employed at the low infrared spectrum to provide centimeter-accurate positioning data for objects around the vehicle that are directly visible to the sensor. Ultrasound sensors provide proximity indications for objects near the car and are useful for automatic parking capabilities to sense curb locations. And of course, radio frequency ranging sensors (radar) are employed at millimeter-wave frequencies for long-range object and obstacle detection, as well as for tracking the velocity and direction of the various actors in the environment around the vehicle.

Three major classes of radar systems are typically employed in automotive active safety systems: short-range radar (SRR), medium-range radar (MRR) and long-range radar (LRR). LRR systems are generally applied as forward-looking sensors and provide adaptive cruise control (ACC) and early collision detection functions. LRR systems at 77 Ghz have ranges that typically scan from 0.5 m to 250 m in front of the vehicle, within a relatively narrow field of view. MRR systems are typically used to watch the corners of the vehicle, perform blind spot detection, observe other-vehicle lane cross-over and avoid side/corner collisions. Typical MRR systems operate in the 79 GHz band to avoid interference with LRR systems. They observe their environment to a range of 50–60m, with higher measurement accuracy than LRR systems. SRR systems are typically applied to collision-proximity warning and safety, as well as limited parking assist features.

Of the four systems described above, most autonomous vehicle developers are looking to a combination of radar, LIDAR and visual camera systems for continuous observation of the environment in all directions. Compared to visual camera and LIDAR systems, radar is a more mature, less costly technology, with superior capabilities in inclement weather situations. Radar sensors are already incorporated into the safety systems of many existing passenger vehicles. In addition, radar can "see around" vehicles on the road, to monitor the movement or braking of vehicles that cannot be seen by LIDAR or cameras because of visual blockage. LIDAR systems, while yet too costly to be practical for standard passenger vehicles, show promise for very precise measurements of the distance to the first blocking object in any direction. Visual camera sensors will be essential for autonomous control, since their ability

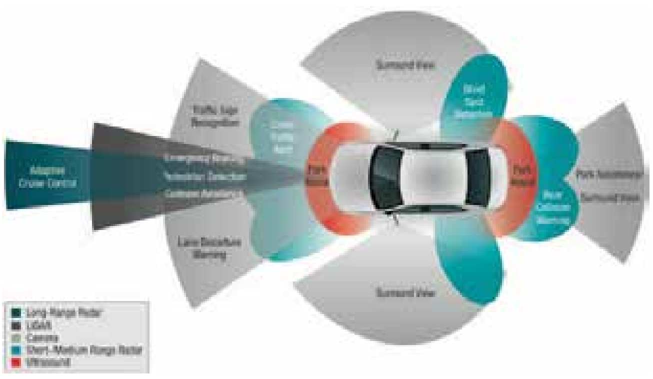

Radar, camera, LIDAR and ultrasound systems in automotive active safety systems (from https://medium.com/frontier-tech/the-next-seat-belt-60e980c3ea8b) to differentiate colors and shapes are important to such functions as traffic light monitoring, visual lane observation and brake light detection — important cues that human drivers detect with their eyes.

Today's automotive radars are incorporating advanced technology that 20 years ago could only be found in advanced research in aerospace and defense laboratories. Radio frequency integrated circuit (RFIC) technology improvements based upon new RF-CMOS and SiGe IC processes are leading to higher chip-level integration, smaller packages and sensors, fewer parts, lower power consumption, higher performance and — most important for this commercial industry — dramatically lower cost. Improvements in antenna designs and materials have led to high-performance antennas exhibiting important features like high gain and sidelobe level control. Developing current and future mm-wave automotive sensors at 24 GHz, 77–79 GHz, and in the coming 122–138 GHz bands, requires judicious use of modeling and simulation to meet aggressive development schedules and to achieve challenging performance requirements.

Radar sensor developers utilize high frequency electromagnetic field solvers to prototype, design and optimize the performance of radar antennas and subarrays. The ability to virtually prototype new antenna topologies and the effects of the packaging and sensor mounting structures on the antenna performance are crucial to achieving successful design of a single prototype. The old, traditional way of designing antennas — designing, building and testing physical prototypes — is too costly. OEMs (automotive vehicle manufacturers) are on an annual product renewal schedule and sensor manufacturers are pushed to meet these schedules with updates and improvements with each model year. To meet new vehicle model safety system performance goals, sensors are frequently improved, updated, or replaced with new-generation sensors. Meeting tight development and validation deadlines demands the creative use of modeling and simulation approaches that enable the build and test of a single working prototype.

/ Rapid development of radar sensors

High-performance radar design starts with the antenna; the antenna is the interface between the sensor and the world. Ideally, these antenna systems must concentrate energy on only the cone coverage angle that the sensor must observe. In LRR antenna systems, this observation angle needs to be quite narrow, and should be able to discriminate against energy arriving at the antenna outside that angle. In the case of the MRR and SRR antenna systems, the energy must be accepted from a wide range of angles to watch the actors over an entire 90-degree sector or more, for potential safety scenarios. Obviously, the antennas must be efficient in radiating — energy should not be dissipated in the antennas themselves, in the sensor package materials or due to poor mismatch with the transmit power amplifiers.

High frequency modeling and simulation present tremendous opportunities in the development of radar sensors. In order of the design process, these opportunities include:

1. **Prototyping and "tuning" antenna topologies quickly, without requiring fabrication.** The best topology for a sensor may not be immediately ap-parent. The ability to virtually test and tune candidate antenna topolo-gies in minutes makes it possible to arrive at a candidate design that best meets performance requirements, without taking excessive time to build and test many prototypes.
2. **Testing antenna variants quickly and painlessly to understand their behavior under a variety of structural and environmental conditions.** Depending on how an antenna system radiates, it may experience significant interaction with the sensor packaging and housing, mounting brackets, gaskets or other structures near the antenna elements. Virtual prototyping presents the opportunity to redesign the structure around the antennas related to the sensor package which might otherwise harm the antenna response through electromagnetic coupling or proximity effects.
3. **Optimizing element and array designs automatically and quickly.** Once a sensor antenna design topology and package are selected, an antenna may be combined with other antennas to form multichannel antenna arrays, and the antenna designs can be automatically tuned to achieve the best possible satisfaction of a combination of performance criteria. For instance, a specific radiation pattern beam width may be required, with minimal elevation sidelobes, while presenting a good electrical match to the transmit power amplifier.
4. **Building a single prototype to test at the end.** The aim of the modeling and simulation process is to remove the time, materials and cost associated with designing, building, measuring and debugging multiple physical prototypes. Engineering the design with modeling and simulation reduces work that formerly required 8–9 months to a few days. Developing the complex sensors required to satisfy the performance needs of tomorrow's active safety systems requires an accelerated design process based upon virtual prototyping.

Let's take an example. Suppose we would like to develop a 77 GHz auto-motive radar sensor that is based upon a two-sided printed circuit board (PCB) fabrication technique. A topology based upon a popular slotted waveguide structure is proposed. Usually, waveguide structures are hollow metal pipes. Slots can be cut on the walls of one side of this rectangular pipe and if they are properly designed, they will radiate efficiently and in a manner that can constrain the energy to any cone angle we wish (a process called "beamforming"). In this case, the broad sides of the waveguide will be formed using the metal cladding of the top and bottom of the PCB and the narrower walls will be realized by using rows of vias that are drilled through the circuit board and metalized to connect the top and bottom metal planes. This technique forms a topology called a "substrate integrated waveguide (SIW)."

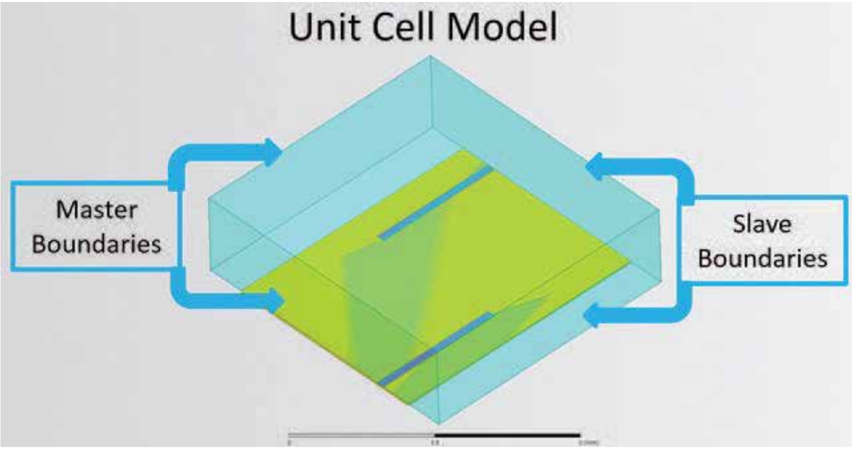

Ansys HFSS starts the design of an extended antenna array as a unit cell. The unit cell approach can be used to quickly optimize proper dimensions for each element.

Slots can be etched into the waveguide walls to radiate energy. Multiple rows of SIW slot radiators can be set up and power-fed in parallel to constrain the radiation in two dimensions. The wider the physical antenna in each dimension, the narrower the resulting radiated power angle becomes. With Ansys HFSS, it is possible to start with a unit cell of the design and let the software optimize the locations and dimensions of the slots. This unit cell design approach applies to any antenna topology that might be considered.

Once the design of the unit radiation cell is optimized for the desired frequency, a full-size array can be laid out quickly and automatically. Simulations can then be run to determine the minimum number of cells required to achieve the goals of spatial radiation coverage and the efficiency with which the array radiates power. The software can automate the process of repeating the unit cell and building up the complete antenna array structure.

When an array design has been synthesized to satisfy performance requirements, the fabrication details (vias, metal thicknesses, structures to couple power into the waveguides, etc.) can be added to simulate the full realism of the materials and manufacturing processes. A design of experiments (DoE) can be run against the expected fabrication process tolerances to assess the manufacturing yield of arrays built in this way. An initial design of the array with vias, PCB filler, and transitions is shown below, with a simulated far-field radiation pattern when all the array elements are fed with power.

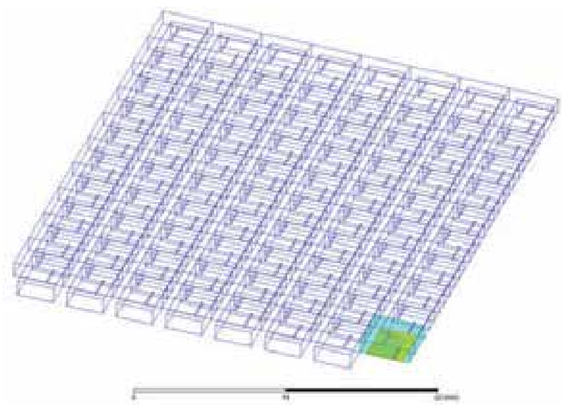

Once the initial design of the sensor antenna system is completed, the effects of the packaging and housing can be investigated. The sensor will need to be enclosed in a watertight package to keep out rain, ice, dust and other materials that will cause deterioration of the electronics or degrade their performance. The presence of metal in the packaging can create electromagnetic coupling to the array that might degrade its ability to radiate to specification. Proximity effects of radome and other nonmetallic packaging can also have an impact. These impacts can be studied, determined and even leveraged in a simulation model prior to building a physical prototype.

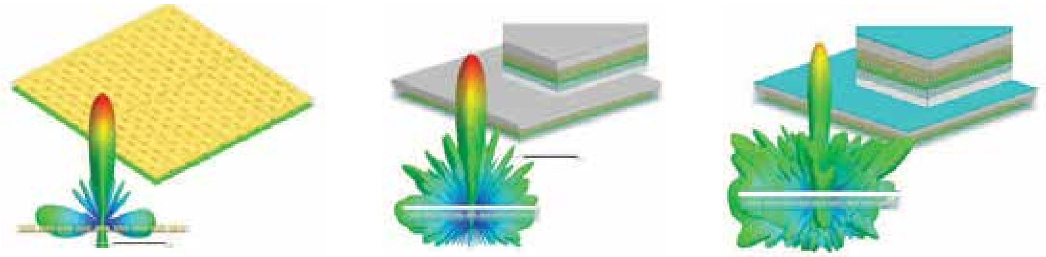

Bare antenna array design and simulated radiation pattern (top); antenna array model integrated with radome and packaging and simulated radiation pattern (middle; and antenna package with a thin layer of water over the radome and the simulated radiation performance of the array (bottom).

For instance, a 1 mm radome is proposed to cover the array, with a metal package backing behind the PCB to provide mechanical rigidity and stability to the sensor design. These packaging features can be incorporated into the simulation model and their presence is shown to reduce the array's main beam gain by nearly 1 dB (which will reduce the maximum range capability of our radar), and raises the first sidelobe level by nearly 8 dB (which will increase the amount of noise coming from off-beam directions that must be processed by the radar signal processing systems).

Environmental effects can also be considered on the packaged sensor's performance in modeling and simulation. In this case, a thin layer of water or ice (0.1 mm) is simulated over the top of the radar package. Simulation shows that the presence of this thickness of water has minimal effect on the main beam gain, but increases the sidelobe level by another 4 dB. Thicker layers of water, ice or other environmental materials can be simulated to understand and quantify the impact of environmental conditions to the sensor's performance. In some cases, this information can be used to go back and build appropriate margins into the original array design so that environmental conditions can be overcome.

/ Integrating the Radar with the Vehicle

Once a sensor design or prototype is developed, it must be evaluated as installed on a vehicle. Many radar sensors are mounted either behind a bumper or in the vehicle fascia. The proximity effects of the vehicle design will affect the performance of the radar — particularly the antenna's ability to control the focus of radar energy. The vehicle manufacturer develops bumper and fascia designs to be both aerodynamic and ergonomically pleasing to their buyers. The unique features of a body shape that meet aesthetic goals can negatively impact the performance of a radar sensor integrated into it or hidden behind it.

In past vehicle generations, the effects of radar-to-fascia and radar-to-bumper interaction effects were evaluated through cooperation between the sensor manufacturer and the vehicle manufacturer. For instance, evaluating a front-end LRR would require the vehicle manufacturer to fabricate a prototype of a fascia and bumper to be mated with a radar sensor prototype from the sensor developer. The installed prototype sensor and vehicle front-end would then be taken to a special antenna test chamber, where the antennas would be energized and measured. If the impact of the fascia and bumper design was found to impair the radar operation, one of three courses of action needed to be taken:

- The sensor required improvement to overcome the installation effects.
- The vehicle front-end required changes to enable better installed radar sensor performance.
- The radar sensor location required changing behind the bumper/fascia to enable more acceptable sensor performance.

In each case, valuable time was invested in developing the prototypes that required retooling and redesign. In the case of a vehicle front-end redesign, there is a potential impact of body shape changes to the overall aesthetics of the vehicle. This process required significant time, resources and materials.

Modeling and simulation reduces this process from as long as 9 months to a matter of days. A digital twin of the radar sensor can be quickly evaluated on a digital twin of the vehicle design, providing a complete virtual prototype of the integrated sensor and vehicle front-end. No sheet metal is cut, painted or primed, no plastic injection is done and no prototype sensor is physically built. The cooperation required between vehicle manufacturer and sensor developer is reduced to an efficient exchange of CAD models. The virtual sensor model is placed into the virtual vehicle front-end design, and modeling and simulation applied to evaluate the installed performance of the sensor in the desired location. If that location turns out to be sub-optimal, the virtual sensor can be quickly moved in the integrated design and re-evaluated. Cooperative fascia/bumper-interaction modeling is quickly becoming a standard practice between vehicle and sensor manufacturers, delivering functional designs on schedule with significant cost reduction.

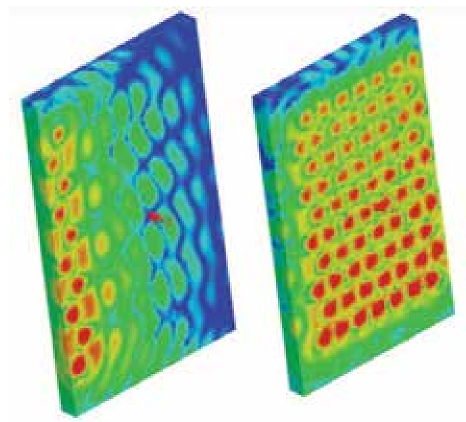

HFSS Simulated near-fields surrounding sensor for Tx channel (left) and for Rx channel (right) form the basis for excitation to the HFSS SBR+ solver.

Modeling the fascia interaction for radar sensors at 77 GHz becomes a rather large electromagnetic problem. At 77 GHz, the integrated fascia, bumper and sensor represents a problem that is on the order of 500 electrical wavelengths across. If more of the vehicle behind the fascia and bumper is modeled to capture multiple reflection effects, the problem size grows even larger. Without significant high-performance computing (HPC) resources, this problem is too large to tackle using a volume meshing electromagnetic analysis approach. A new class of high frequency EM field solvers known as asymptotic field solvers that use a ray-tracing approach is appropriate at this size to model sensor-fascia interactions.

Ansys provides a shooting and bouncing ray (SBR) solver, called HFSS SBR+ which can take the pre-computed solution for the isolated sensor system and model its interaction to the much larger fascia and bumper using ray tracing methods. The simulation process is hybridized, which means that the result of the precise finite-element simulation for the isolated sensor system can be easily applied as an excitation to the SBR+ simulation for interaction modeling. For each radiating antenna system, the finite-element simulation yields a capture of near-fields in the region immediately surrounding the sensor.

Consider the SIW radar sensor array designed here. The left-most column of radiators can be used as a transmit subarray to send the radar pulses. The balance of the columns can be used together as a larger receive subarray. For the transmit subarray, captured near-fields would appear as shown in the graphic on the left, with the strongest fields (shown in red and yellow) appearing in the region immediately above the transmit antenna. These near-fields are used as surfaces with a continuum of ray sources for the SBR simulation. The near-fields corresponding to the balance of the receive subarray are shown in the graphic on the right, and are used to define the radiation/collection capability of the receive subarray system. Both models are taken to the HFSS SBR+ simulation and used as simulation sources for the antenna systems.

The result of the hybridized FEM-SBR analysis is a simulated installed radar antenna response that shows the radar engineer how each radar antenna subarray will illuminate the road or environment when it is installed into the proposed fascia-bumper design.

The result for the radar receive antenna subarray is shown in the graphic group on the left for the isolated sensor (computed by the HFSS finite-element method solver), and on the right for the installed sensor simulation in the proposed vehicle fascia (computed by

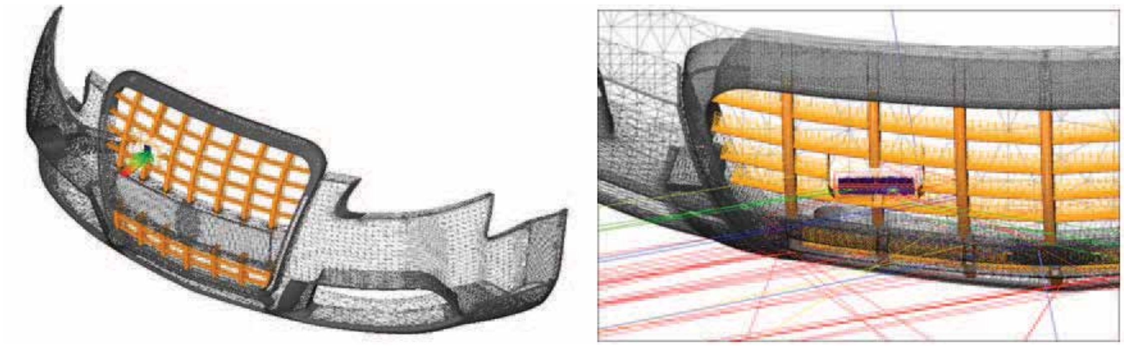

Radar sensor array model is installed in a proposed automobile fascia (left) and the HFSS SBR+ shooting and bouncing rays EM field solver is applied to model the installation interactions. The HFSS finite-element simulation for the radar sensor antenna system is shown in the proper installation location, and a subset of rays employed by the HFSS SBR+ simulation is shown at an exit angle of 80 degrees (right).

HFSS SBR+). The simulation shows an expected reduction in peak radiation gain, along with a noticeable rise in sidelobe levels.

The transmit channel subarray is computed similarly. The simulated interaction between the radar sensor and the fascia also indicate peak gain reduction and increased sidelobe levels. This technique can also be extended to multichannel radars to observe errors in the far-field phase for direction of arrival processing.

When the interaction modeling has been completed for a sensor and a proposed automobile fascia design, the far-field results can be used in yet another level of simulation — full road environment interaction modeling.

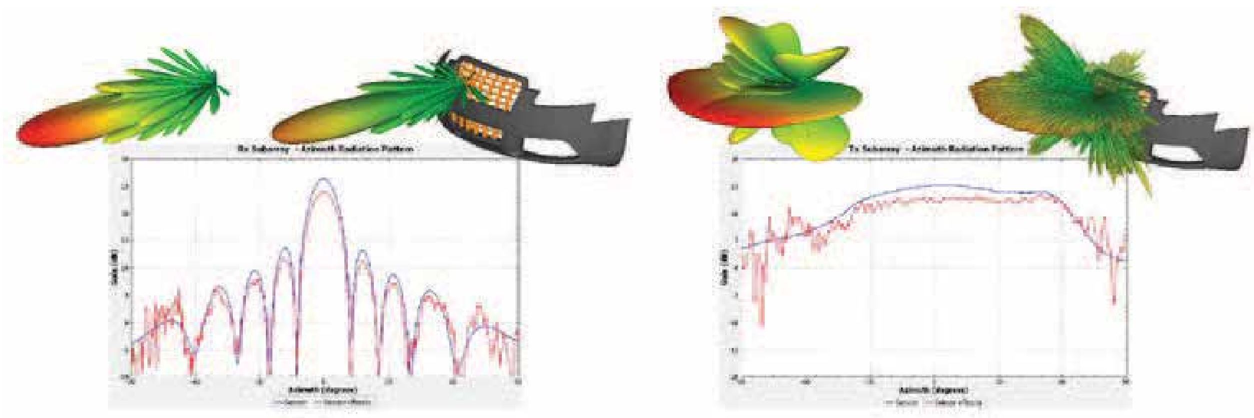

Receive channel subarray radiation pattern (left) and transmit channel radiation pattern (right) showing radiation patterns for the module in isolation and as installed to include fascia and bumper interaction.

/ Virtual Road Testing for the Radar

A revolution in transportation is fast approaching with the expected deployment of semi-autonomous and fully autonomous vehicles (AVs). AV developers are very conscious of the safety and protection of passengers against any traffic hazards that may appear. Playing a central role in these safety systems are the current and future integrated radar systems. The process of designing safe AVs involves thoroughly testing sensor systems together with vehicle control systems and algorithms to validate safe operation. To date, high-fidelity testing of sensors and vehicular control systems has required that closed-loop operation be performed through driving millions of test miles. The sensor data, control algorithm outputs and vehicle dynamics are recorded during test drives of dozens of vehicles, and usually involves at least two passengers per vehicle. After the data is recorded, it must be processed, problems identified, updates applied to the vehicle control algorithms and the process repeated. This process requires months to build up a database of test cases and to optimize the control systems, and can only be performed using currently available vehicles, sensors and control hardware. In addition, only test cases that do not involve actual collision or danger scenarios are (intentionally) tested.

Today, most AV developers are engaged in moving this process to the domain of the digital twin. In modeling and simulation, this testing can be performed for any conceivable scenario — even those where loss of life might be a potential downside. Models for vehicle dynamics could be used for a vehicle that is still in the concept phase. Control algorithms can be quickly updated and evaluated without the time constant of uploading to a vehicular control system. Sensor systems beyond those that are currently commercially available could be tested, or new sensor concepts tested. And, best of all, scenarios in modeling and simulation can be conducted at a much higher rate than in the real prototype world.

Modeling the electromagnetic performance of an automotive radar system in a real, high-fidelity world has, to this point, proven to be a major challenge. Previous approaches to radar sensor modeling have used simple point-source statistical scattering models that assume all the radar scattering comes from a single point. Further, this approach assumes no target interaction with the environment or with other actors on the road. Full physics modeling of the radar with the environment around the vehicle is desirable for synthesis of high-fidelity radar data.

The AV control system requires accurate data on the environment around the vehicle including:

- A list of actors around the vehicle, both stationary and moving.
- Information on each actor, including:
    - The distance (range) to the actor.
    - The direction vector to the actor.
    - The velocity of the actor relative to the radar host (ego vehicle).
    - An assessment of the target size (signal strength).
- An assessment of radar scattering returns that constitute a common actor, since a single actor may create multiple returns that appear in consecutive radar range bins.
- Radar-generated images of the environment around the vehicle (coming in future radar sensors).

A challenge posed by using full physics modeling and simulation for radar sensors is the sheer size of the analysis problem. At 77 GHz, the range of a typical long-range radar (LRR) is on the order of 250 m. Assuming that the radar operates with a 20 degree wide cone angle, that means that the radar is intended to "watch" an area that occupies over 1.4 million electrical wavelengths. This constitutes a rather large high-frequency EM problem. Several additional system-level requirements add up to make full physics modeling and simulation of radar-environment interactions a challenge:

- Radar updates to the central control system occur at a rate of between five and 30 times a second in most current radar systems.
- Typical LRR and MRR radars today operate with up to three transmit antennas and up to eight receive antennas. A modeling and simulation approach must therefore produce up to 24 channels of simulated radar data. (Future radar systems will require many more channels.)
- The range (distance) resolution of most MRR radars today is much smaller than the length of a typical automobile or truck. Many MRR sys-tems have a range resolution near 0.25 m, whereas an automobile may be 5-7m. Other actors like guardrails may be over 100 meters long. This is a challenge because the actor may exhibit strong scattering in some, most or possibly all of the range bins that the actor spans, resulting in an indication of multiple targets in space.
- Many actors in the environment have velocity or acceleration characteristics that must be considered. This is important for the radar's ability to yield velocity data and usually requires many sequential radar measurements of the environment to track actor velocities. A wider velocity tracking requirement means more measurements must be taken. Higher velocity measurement accuracy (velocity resolution) also leads to a need for more radar measurements to be taken against the environment.

While these considerations pose challenges to high-fidelity EM modeling of radar–environment interaction, they're not insurmountable. An appropriate application of the SBR technique can provide full-physics simulation of such problems with reasonable efficiency — both in terms of computer resources and modeling time.

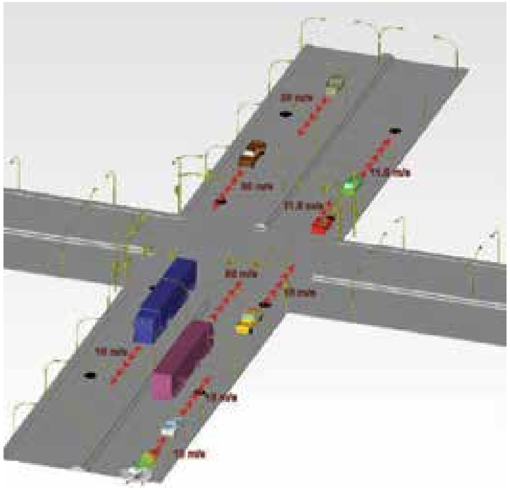

Busy intersection environment geometry; the velocity of each moving actor in the scene is shown.

Consider the traffic scenario below, which is representative of a suburban intersection scene. This scene is about 120 m by 120m, and has many of the common actors that would be expected in a realistic street scene — streetlights, traffic lights, curbs, manhole covers and an assortment of real-istic vehicles with different velocities. Actor velocities are indicated in the figure for this scene. The radar model with fascia interaction effects from the previous section is placed into this scene (lower left) and positioned in the right-hand lane at a height of 0.5 m above the roadway — consistent with its placement on the ego vehicle. Both the ego vehicle and the actors move as dictated by their self-velocities, and during the movement the radar must provide updates regarding the targets that it observes, as well as extract velocity information about each target.

The radar model broadcasts EM fields throughout the environment through a shooting and bouncing ray ray-trace approach. A ray trace used by the simulator is shown in the second figure, illustrating the how the simulator spreads energy to the scattering features in the environment. Each ray carries an amount of energy that is weighted by the transmit subarray's far-field radiation pattern and serves to impress high-frequency physical optics (PO) currents onto the full environment. In a second processing step, the solver integrates the radiation of all of those PO currents back at the receive subarray, weighting the reception by the receive subarray's receive pattern.

Many automotive radar systems are based upon sending short signal pulses and measuring the time required for the pulse to bounce off features in the environment and return an echo. A single pulse can be used to determine the distance from the radar to the echo source. The precision to which range to an object can be determined (the range accuracy rating) is directly related to the bandwidth of the radar signal and signal processing after the pulse is received and detected. In general, radars with wider bandwidth signals can locate reflection sources with higher precision than those with narrower bandwidth and will need to process more range points to cover a specified total range with a smaller granularity. Therefore, wider band radars provide higher accuracy, but require more costly RF and signal processing hardware.

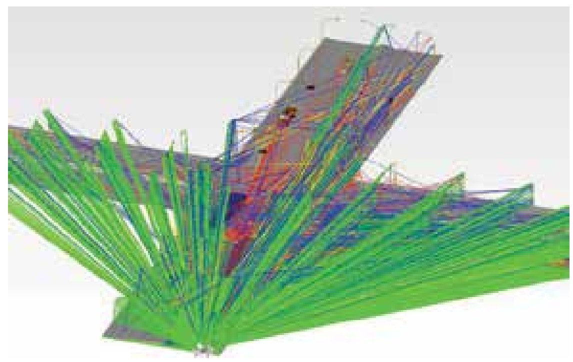

Shooting and bouncing rays traced from a radar transmit channel throughout the environment. Multiple colors correspond to ordinal reflection for each ray track pictured.

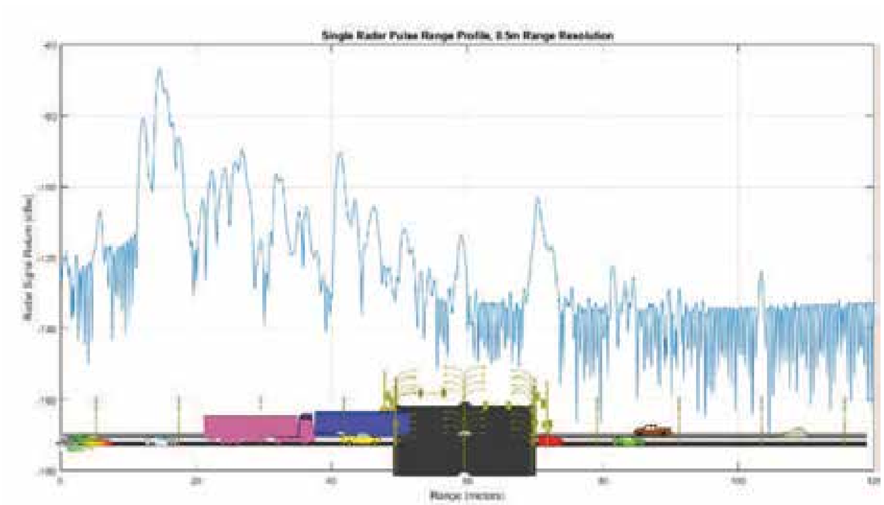

Range profile for a single radar pulse with 300 MHz bandwidth. Radar is visible in overlay at lower left.

HFSS SBR+ can be used to synthetically reproduce the signals obtained by a high-fidelity radar model. In addition, any specified bandwidth may be applied to the simulation, enabling the engineer to test new waveforms that may not yet be available from sensor suppliers. For the simulation environment shown below, a radar pulse is synthesized that will provide a range accuracy of 0.5 m, which implies a radar pulse with a bandwidth of 300 MHz. The simulator traces out the dispersed and reflected energy back to the receive subarray, yielding raw radar receive data in the form of complex voltages at the receive subarray. This data can be quickly collected and transformed into the time domain, to yield a range profile, as shown below.

The range profile shows time of flight (or distance of flight) for all radar echoes received by the radar system in response to the environment model. In this case, the range profile shows very strong radar returns for some light posts, the facing surfaces of several of the vehicles and some reflections that travel between the vehicles. The signals are stronger from closer targets than from more distant targets, but even targets well down the road are detected. In addition, since a waveform was used that provides a range accuracy (resolution) of 0.5 m, several of the actors exhibit multiple radar returns from different scattering features. For instance, the tractor trailer provides returns that are distributed across at least 30 range cells, which means that the radar will detect multiple target echoes which are part of the same vehicle. Radar signal processing systems will need to intelligently group distributed target returns that belong to the same actor in the environment, or the vehicle control system will be overwhelmed with too many targets to track.

One way in which radar systems can determine which returns belong to the same vehicle is to process the possible Doppler-shift of the signal of each return, which occurs when signals bounce off surfaces with a velocity that is different from the observation domain. Radar signals from targets that have the same velocity in consecutive range bins could be considered to be from the same target. Velocity measurement of targets is performed routinely by automotive radars in Doppler processing, but this capability requires more than a single pulse. In Doppler processing, the returns from multiple pulses are observed by the radar controller over an interval of time. During this time, the radar sends and receives hundreds of pulses very rapidly, integrating the received signals and particularly watching the relative electrical phase of the received signals. If enough pulses are processed and stored, the history of the movements of actors can be processed to make an accurate estimate of the velocity of each target. If targets can be assessed in terms of both range and velocity, it is easier to decide which returns come from the same target. In addition, the vehicle control system can begin to make some estimates about where actors in the environment will be at some point in the future — a feature very important to safety systems.

For the system under investigation here, a velocity accuracy of 0.5 m/s is desired, together with the ability to resolve speeds over the interval of -100 m/s to +100 m/s, where speed is measured relative to the velocity of the radar. To achieve the required velocity accuracy and velocity range requirements, the radar will need to track pulses over a total of about 9.8 milliseconds (sometimes called a radar frame), and will need to process at least 200 pulses. In simulation, that means that the environment simulation will need to be repeated at least 200 times, advancing the moving actors and ego vehicle (radar) an appropriate distance between pulses as dictated by their individual velocities. Once the pulses within this radar frame are received, the radar signal processing hardware can process the data to obtain both range and Doppler velocity information for every return from the environment.

To simultaneously examine both the range and velocity of actors in the environment, radar systems engineers will often use Range-Doppler maps. This data presentation provides the range to the environment returns on one axis and the extracted velocity of the targets on the other. The colors shown in the map are keyed to the strength of the signal that comes back in each range and velocity bin.

In this case, the radar has a self-velocity of 10 m/s. In the radar's frame of reference, that means that all stationary objects in the environment will appear to have a closing velocity of 10 m/s, which registers as a +10 m/s velocity on the Range-Doppler map. Signal returns with positive velocities (greater than zero) indicate objects coming toward the radar; this is known as "closing velocity." Conversely, signal returns with negative velocities (less than zero) indicate objects moving away from the radar. Negative velocities are known as "opening velocities." For example, the car immediately in front of the radar is traveling 5 m/s faster than the radar and is the closest large reflecting actor in the scene. This corresponds to the strong return observed starting at a range of 10 m and showing an opening velocity of -5 m/s.

The distributed nature of large actors like tractor trailers is also clearly observed in the Range-Doppler map. Note the large return that starts at a range of about 20 m, with a velocity of -20 m/s. This return is from the tractor trailer in the adjacent lane that is moving at a velocity of 30 m/s. The return from this large vehicle is distributed through several range bins starting at 20 m in a vertical line extending to around 37 m. Since the returns in this line have the same velocity (they extend in a vertical line), it's safe to say that they belong to the same target and the radar processor may combine them together and flag them as a single large target with an opening velocity of -20 m/s.

Radar simulations can be computed as the environment changes. A typical automotive radar sensor provides updates to the vehicle control and safety systems at a rate of between 5 and 30 per second. Simulating this changing environment is a natural extension of the single-frame simulation. A complete simulation of the movement of vehicles through this environment has been performed, and an animation of the scene and the Range-Doppler map can be viewed at https://youtu.be/v2sJKa3vjEg. In the animation, the progressive range separation of the targets provides intuition into the radar's ability to discriminate between targets based on range and Doppler signatures.

This simulation can be placed within a complete AV simulation loop, controlling the simulated environment and waveform selection at each radar frame (or each pulse). The radar sensor simulation can pass either raw radar data or target indication data back to the simulated vehicle control system or active safety system.

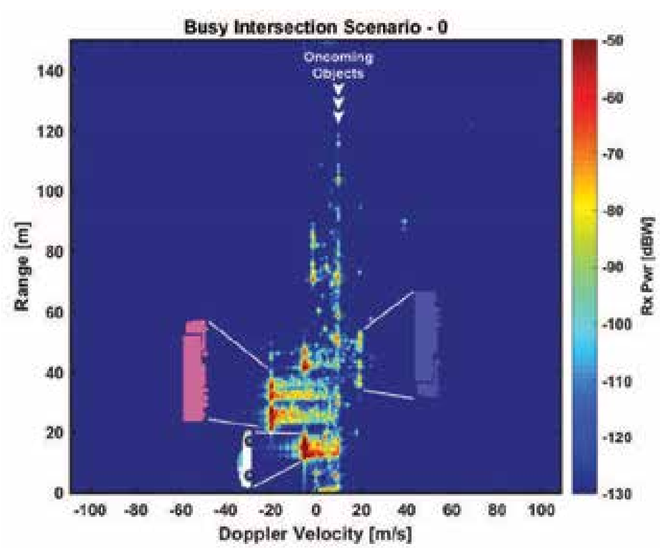

Range-Doppler map for radar system over a radar frame of 200 consecutive 300-MHz pulses.

Complete Modeling and Simulation Work Flows

Ansys has the tools and expertise to assist radar sensor developers, automotive OEMs, active safety systems developers and autonomous vehicle control systems developers. Whether the challenge is designing the radar sensor module, studying its installed performance on the vehicle, or gaining insight into what the radar reports for moving and stationary targets on a full, dynamic road scene, Ansys can help. With over 25 years of expertise in microwave and mm-wave electromagnetic (EM) modeling and simulation, Ansys can help you to apply best-in-class tools to solve any high frequency problem from MHz to THz frequencies.

Ansys also fields world-class modeling and simulation tools for mechanical analysis, fluid analysis, thermal analysis, systems analysis, functional safety and display systems and more, providing a full suite of tools aimed at solving complete vehicles in high-fidelity environments. Ansys is moving autonomous driving and active safety systems development to the digital twin, where development can be performed safely, more quickly and with considerably less cost than a serial design-build-test loop. In the digital twin universe, the design of multiple stages of a vehicle can be accomplished in parallel, using the latest digital twin of each subsystem available.

Let Ansys show you how to take schedule and cost out of the development of autonomously driven vehicles and active safety systems.

ANSYS, Inc.
Southpointe
2600 Ansys Drive
Canonsburg, PA 15317
U.S.A.
724.746.3304
ansysinfo@ansys.com