Mitsubishi Electric Research Laboratories

30 YEARS

1991 — 2021

INNOVATION | IMAGINATION | INSPIRATION

MITSUBISHI ELECTRIC

"At MERL, it was exciting, engaging, tough, sleepless, and fulfilling all at the same time."

— Ashok Veeraraghavan, Ph.D.
Associate Professor, ECE Department
Rice University
At MERL 2008-2011

"MERL is more like a family where everyone is supportive, friendly, and always ready to help in your research without many expectations. It was a fantastic learning experience, and I am glad to have worked at MERL for 8 years when I was really shaping my research career."

— Srikumar Ramalingam, Ph.D.
Associate Professor, School of Computing
University of Utah
At MERL 2008-2016

Production:
Richard C. Waters

Photos:
Fred Igo

Design:
Frances Scanlon, Scanlon Design

Introduction

Celebrating 30 years of innovation, Mitsubishi Electric Research Laboratories (MERL) invites you to find out more about us.

MERL conducts application-motivated basic research and advanced development in Physical Modeling & Simulation, Signal Processing, Control, Optimization, and Artificial Intelligence. We are an open laboratory, strongly involved with the world research community by publishing our work and collaborating with numerous interns and universities.

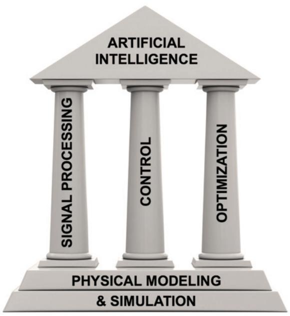

With 60 researchers, MERL is small enough to be flexible and agile, while gaining leverage from our large global parent Mitsubishi Electric. We turn our technical achievements into business successes by partnering with the tens of thousands of researchers and engineers in Mitsubishi Electric's operations around the world.

The information on the following pages paints a picture of MERL's past and present, while pointing the way toward our future.

For additional information about MERL, visit www.MERL.com.

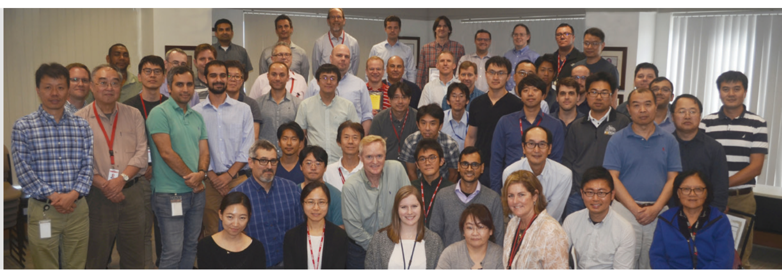

MERL History

One history of MERL is organizational. In 1991, Mitsubishi Electric's Corporate Research and Development organization (CR&D) opened MERL in Cambridge, Mass., under the leadership of Dr. Tohei Nitta and Dr. Laszlo (Les) Belady. From the beginning, the focus of MERL

Dr. Tohei Nitta

Dr. Les Belady

Dr. Richard Waters

has been on long-range research. However, MERL has evolved in various ways over the years, particularly due to the incorporation of two other labs.

Seven years before the founding of MERL, Mitsubishi Electric's Computer Business Unit founded a lab called Horizon Research, Inc. (HRI) in Waltham, Mass., to design mid-range computer hardware. The Computer Business Unit left that business, and HRI was transferred to CR&D in the early 1990s, where it transitioned into an advanced development software lab.

In 1993, Mitsubishi Electric's Audio/Visual Business Unit founded a lab called the Advanced TV Lab (ATL) in Murray Hill, N.J. ATL's first project was collaboration with Bell Labs on the development of the first chipset capable of decoding U.S. HDTV signals.

When difficult economic times became prolonged in Japan, ATL was also transferred to CR&D.

In 1995, CR&D decided to combine HRI and ATL together into a new organization called Mitsubishi Electric Information Technology Center America (ITA). The next year, MERL was merged into ITA as well. An important goal of the formation of ITA was fostering collaboration between CR&D's labs in the U.S., but geography trumped good intentions and little collaboration ensued.

In 1999, Dr. Richard C. Waters became the head of MERL and proceeded to confront the root of the collaboration problems. In 2000, the staff in Waltham moved to Cambridge, and ITA returned to the name MERL. In 2003, the staff in Murray Hill also moved to Cambridge. This led to a structure of two labs in Cambridge: one focused on long-range research and one on research & advanced development.

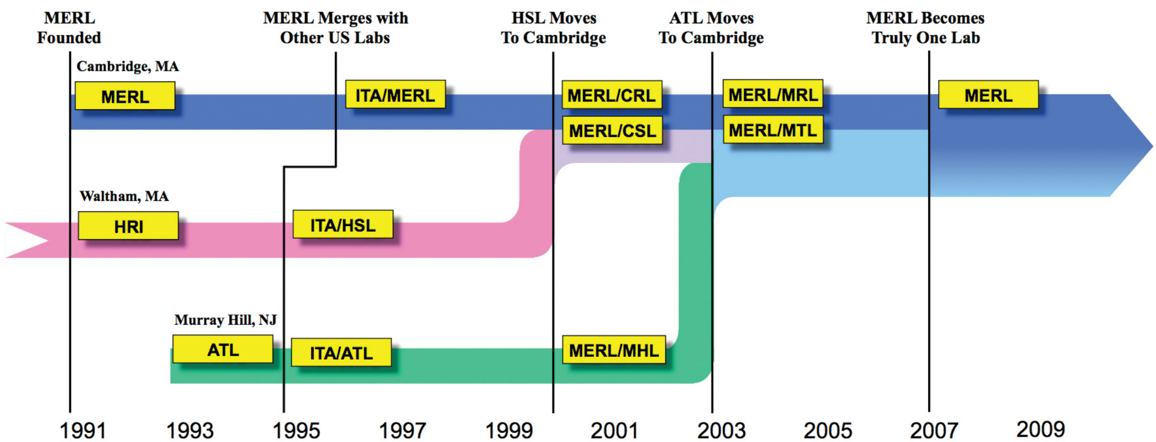

Once everyone was in the same place, strong collaboration became the norm; however, there was still a psychological division between the two labs. In 2007, MERL was reorganized into a single lab, with the unified goal of long-range research, but retaining advanced development capabilities—a structure that continues to this day.

Interaction Models

A second history of MERL is philosophical. The initial model of interaction between Mitsubishi Electric and MERL was similar to interactions with a university. MERL was to pioneer new technology in areas outside Mitsubishi Electric's current business, and it was the responsibility of Mitsubishi Electric to pick up and use the new technologies.

It should come as no real surprise that with this approach, MERL's contribution to Mitsubishi Electric's business suffered from the problems concomitant with university/industry relations in general. It is very difficult for any company to adopt new technology without a lot of support from the inventors of the technology and particularly so when the technology is outside of the company's area of business.

In the years around 2000, the interaction model changed. MERL began focusing on pioneering new technology in areas related to Mitsubishi Electric's business and increasing collaboration with our research colleagues in Japan, so that we could play an active role in transitioning our technology into business. This has proven to be a much more successful approach, leading to many significant impacts on Mitsubishi Electric's business.

Results

A third and most important history of MERL is what we have produced. The following pages summarize areas of MERL research, interspersed with short descriptions of major results in those areas. The first four areas are ones where we are no longer active, but where important results were created. The rest are areas MERL is currently engaged in.

Research Areas and Results

The following pages summarize areas of research at MERL, interspersed with short descriptions of significant research results in those areas. The first 19 pages describe work in areas that MERL was previously active in but is not very active in at the moment. The remaining pages describe work (some going back many years) in the areas where MERL is currently most active.

"At MERL, I was in a unique position to publish high-quality research and at the same time see many of my ideas spawn new product families."

— Ramesh Raskar, Ph.D.
Associate Professor, Media Lab
Massachusetts Institute of Technology
At MERL 2000-2008

"MERL is such an amazing place where researchers with various expertise work closely together on research projects, which I believe is the most effective way for engineering research in modern times where problems are interdisciplinary."

— Lei Zhou Ph.D.
Assistant Professor, Department of Mechanical Engineering
University of Texas at Austin
At MERL 2019-2020

Computer Graphics



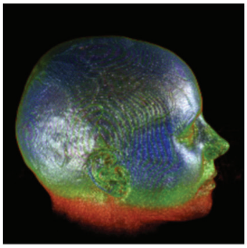
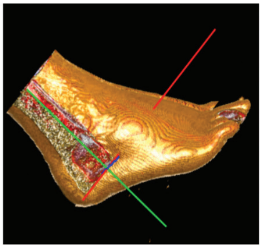

Although Computer Graphics is not presently an area of research at MERL, it was a major area, particularly from the middle 1990s into the early 2000s.

The dominant representation for a 3D graphics model is a collection of 2D polygons positioned in 3D space. This is a good representation for many tasks, but not ideal in every situation. MERL researchers did pioneering work on the direct rendering of 3D raster-scanned, volumetric data and invented an entirely new representation: Asymmetrically-Sampled Distance Fields (ADFs)—a parametric representation of the distance between each point in 3D space to the nearest surface.

MERL developed a number of innovative applications combining cameras & projectors and using graphic-rendering techniques to pre-distort images sent to the projectors, in order to create innovative effects. In addition, we were pioneers in image resizing and computational photography, which captures image data in innovative ways that can be extensively manipulated later.

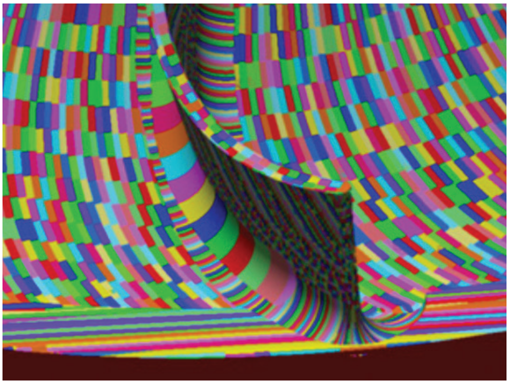

Real-Time Volume Graphics

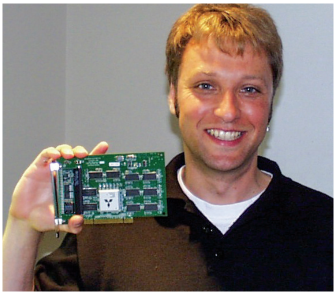

Demonstrated at the SIGGRAPH conference in 1999, VolumePro was the first PC graphics card for real-time visualization of 3D voxel data from scientific and medical instruments such as computed tomography (CT) and magnetic resonance imaging (MRI) machines. At the time, it was the fastest commercially available visualization solution and won several industry awards.

VolumePro could render a 256 x 256 x 256 block of voxels at 30 frames per second. By directly rendering from the captured data, it avoided the approximations required when converting such data into a polygon representation to be rendered by standard graphics processors. By operating at 30 frames per second, it allowed the viewer to manipulate the image in real time, which had previously been impossible on commercial systems.

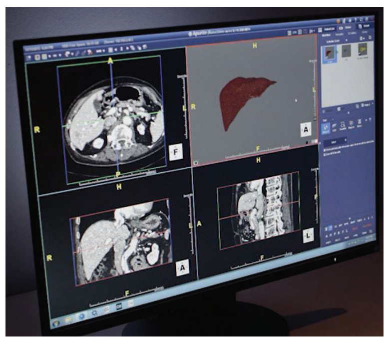

Initially commercialized by a company spun out from MERL, the technology was subsequently acquired by TeraRecon and was used by customers in medicine, biology, and engineering for many years.

Computational Projection

Projectors decouple the generation of an image from the surface it is displayed on, introducing room for innovation but causing potential alignment problems. MERL research combined projectors and cameras with computer vision and graphics algorithms to transcend the problems and enable novel effects.

A camera-equipped projector can automatically calibrate itself so that the image cast fits a screen without precise positioning of the projector.

Large immersive displays, such as those on hemispherical screens, can be created at low cost by using a camera and multiple projectors that automatically calibrate themselves to create a seamless result.

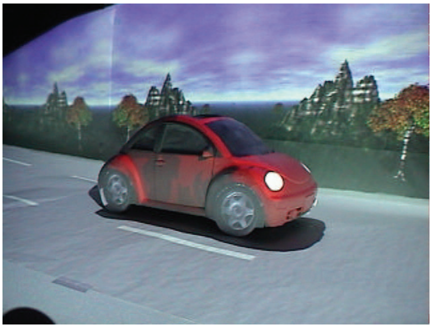

This approach can be extended to more complex shapes to change their appearance and/or make them appear to be moving with so-called "Shader Lamps."

MERL researchers demonstrated novel interaction between people and mobile camera/projector pairs, such as having the projector point out things in view to a user.

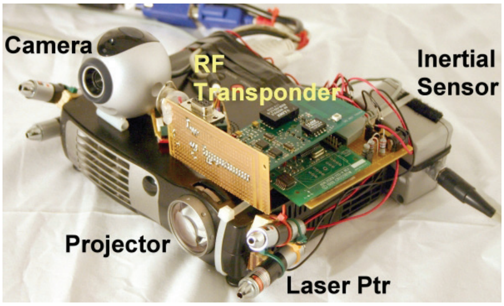

Computational Photography

A traditional digital camera differs from a film camera only in the fact that the output is digital. Computational photography takes a step further—instead of just capturing a directly viewable image, it captures more complex data that may not be directly viewable, but with proper post-processing can be manipulated in more useful and interesting ways.

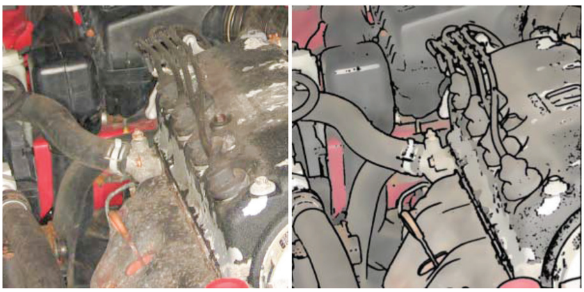

MERL was a leader in computational photography. In early work we created the “multi-flash camera,” which took multiple images using lights in different positions to make the extraction of depth edges easy. Our “flutter shutter” camera took a single image with a shutter that opens and closes multiple times to make the removal of motion blur from an image possible. Our “coded aperture” camera took images through programmable masks that make it possible to refocus an image during post-processing.

Today’s cutting-edge commercial cameras make extensive use of computational photography principles to make photography more forgiving and precise.

Seam Carving

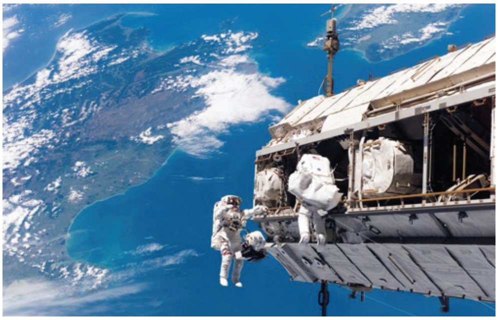

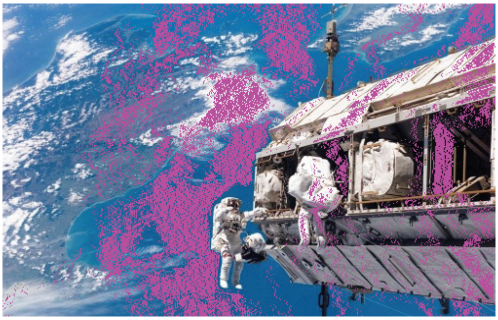

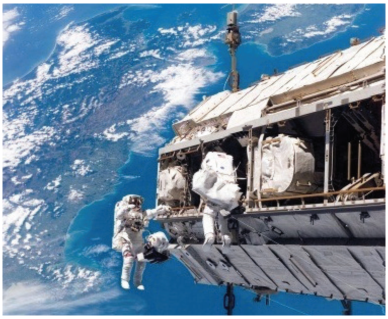

When a picture does not have the right aspect ratio to fit in the space available, there are two basic options: cropping it or distorting it. Seam Carving introduced a new option: recomposing the picture. Seam carving finds the image-spanning chain of pixels that is least important. This chain can be removed to make the image one pixel narrower with minimal change in the appearance of the image. Alternatively, similar new pixels can be inserted beside the chain to make the image one pixel wider.

Because seam carving is a simple and fast dynamic-programming algorithm with surprisingly good results, it was one of the first algorithms from MERL to go viral on the Internet. We quickly followed on with a more flexible algorithm enabling video carving. Image Darting then introduced significant quality improvements by finding the best set of pixels (connected or not) to remove so that the remaining parts of the image can be pieced together to produce a realistic-looking image. Fifteen years on, these algorithms remain popular in photo-editing software.

Adaptively Sampled Distance Fields

Polygonal graphics representations explicitly define the surface of an object. Adaptively Sampled Distance Fields (ADFs) implicitly describe a surface via parametric mathematical functions defining the distance between each point in space and the nearest part of the surface. ADFs use a detail-directed adaptive partitioning of space and can represent very high levels of detail, using much less memory than polygonal approaches.

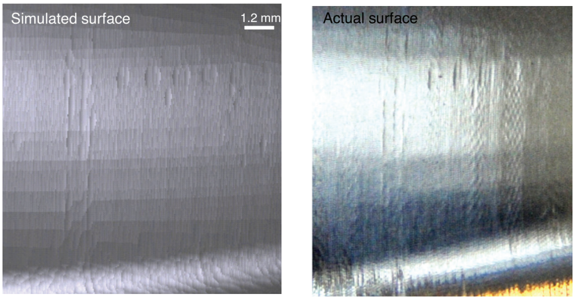

In 3D, ADFs have been applied to Computer Numerical Control (CNC) milling to create extremely accurate simulations of complex parts that can reveal minute details of a milled surface to help diagnose manufacturing problems.

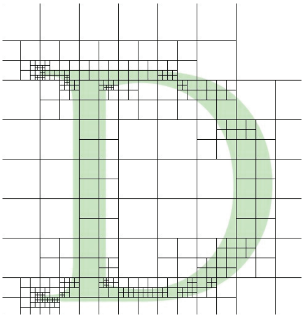

In 2D, ADFs have been used to compactly represent fonts that can be scaled to arbitrarily large sizes and rendered clearly even at very small sizes. They require less memory than other approaches and can produce high-quality rendering across all sorts of displays, from computer screens to hand-held devices. MERL's font-rendering technology has been used by billions of people in a wide range of devices and font-rendering systems.

Human-Machine Interfaces

HMI research per se is not a major focus of MERL's research today, however, several MERL projects in past years had a strong influence on the HMI research community and ultimately on the innovations appearing in current human-machine interfaces.

MERL's DiamondTouch table was the world's first reliable multi-user, multi-touch surface. MERL combined interactive tables with wall displays and personal devices to develop novel interaction methods for collaborative spaces.

MERL developed a distributed platform supporting multi-user interaction in 3D virtual worlds that pioneered key concepts used in today's large multi-user online environments.

MERL developed an influential platform called Collagen that provided interactive support for users, based on knowledge of the various tasks a user might perform.

DiamondTouch

The DiamondTouch table was a front-projected, interactive display that allowed multiple users to sit face to face and work together on the same screen. The key difference between DiamondTouch and other multi-touch interfaces was that multiple users could touch the table at the same time and DiamondTouch could keep track of who was touching where. Even today, most other multi-touch interfaces cannot tell the difference between two users touching in one place each and one user touching in two places.

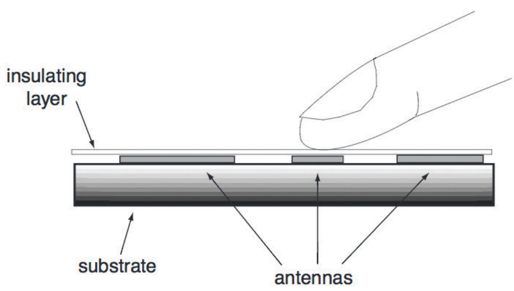

The surface of the table contains two perpendicular transmitter antenna arrays. When users touch the table, they complete circuits between these arrays and receivers in the chairs they are sitting on via capacitive coupling. Because each user has a different receiver, DiamondTouch can determine which user is touching which parts of the antenna arrays.

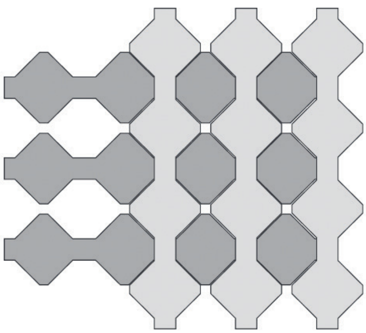

Diamond Park

Demonstrated at COMDEX in 1995, Diamond Park was a multi-user environment supporting multiple geographically separated participants interacting in a 3D virtual world. Users could interact using an exercise bicycle interface, in which case they appeared as bicyclists in the Park, or via a computer terminal, in which case they appeared as unicyclists. The Park also included computer-controlled autonomous agents.

Diamond Park was implemented using MERL's Scalable Platform for Large Interactive Environments (SPLINE), which could support large numbers of users in a large environment with no centralized bottlenecks. SPLINE was the first virtual reality substrate capable of supporting real-time spoken interaction between the participants. More importantly, it was the first to support arbitrary modification and extension of the environment during operation. Both of these features are essential parts of the massive online gaming environments of today.

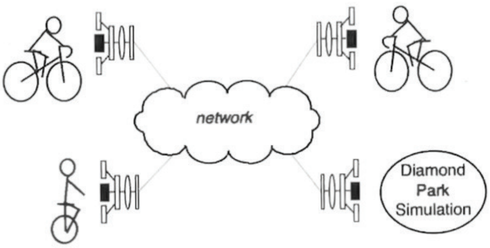

Collaborative Agents

As high-tech home products have become increasingly complex and interconnected, many people have become overwhelmed and utilize only a small fraction of these products' capabilities. MERL introduced a new approach to addressing this problem.

Our solution was to provide an intelligent collaborative agent that supports a simple, unified, and conversational interface to all the appliances in a household. Furthermore, this agent has a model of the typical tasks that the user might want to do and can interactively guide the user through the necessary steps.

This work resulted in a software platform called Collagen. It was used to construct a number of prototype collaborative agents, including DiamondHelp, which was a finalist in two design competitions in 2005: the INDEX Award in Copenhagen, Denmark, and 3AD (the 3rd Int. Conf. on Appliance Design) in Bristol, UK. Collagen has been influential on subsequent dialogue-processing research.

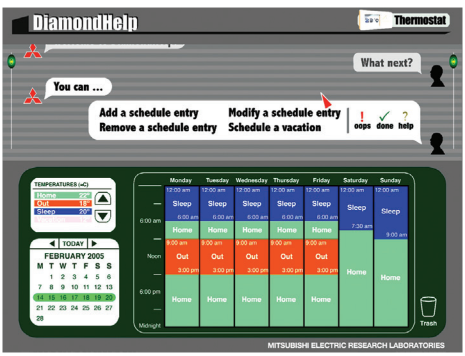

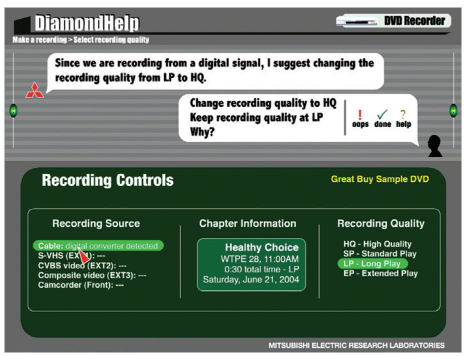

Digital Video

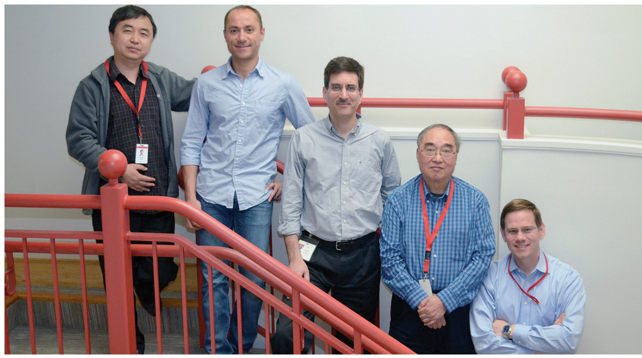

The transition from analog to digital television in the U.S. triggered the beginning of digital video work at MERL. In the mid-1990s, we contributed to the design of the first HDTV receiver chip set and began participating in the development of MPEG video-coding standards.

In the early 2000s, we developed leading algorithms and architectures for video transcoding, which led to a PC-based product that simultaneously transcoded multiple MPEG-2 streams from surveillance cameras into compact, low-resolution MPEG-4 streams, delivering them to mobile devices. On the analysis side, we contributed extensively to the standardization of MPEG-7 descriptors and developed novel algorithms to support video highlights playback in DVD recorders.

MERL made contributions to the H.265 High Efficiency Video Coding (HEVC) standard and drove the development of extensions that improve the coding of screen-captured content. In addition, MERL led the development of several 3D and multi-view video-coding standards including the Multiview Video Coding amendment to the H.264 Advanced Video Coding (AVC) standard, which was been commercially deployed for 3D Blu-ray Discs.

HDTV Receivers

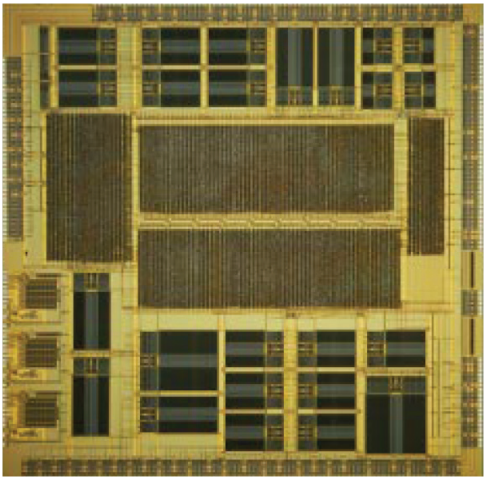

In 1998, Mitsubishi Electric and Lucent Technologies' Bell Labs jointly completed the development of the first integrated chipset meeting the North American standard for receiving and displaying High-Definition Television (HDTV) signals. The chipset supported demodulation, demultiplexing, audio & video decoding, and display-processing functions. MERL contributed to key parts of the chipset, with particular focus on demodulation and video decoding.

For a 2005 second-generation HDTV chipset, MERL developed high-quality, highly efficient down-decoding technology for converting a high-definition bit-stream into a standard-definition one. This allowed the chipset to support backward compatibility with existing TVs at low cost, which helped propel a successful transition from analog to digital television broadcasting in the U.S.

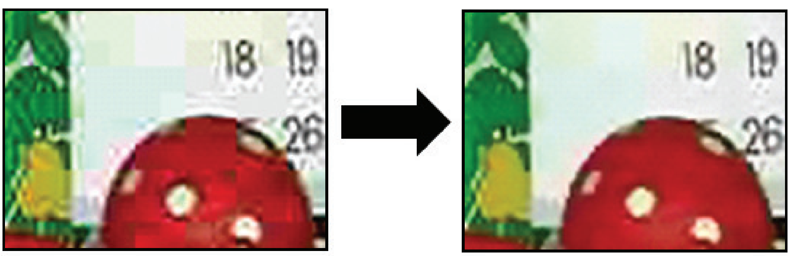

Video Highlights Playback

In 2005, Mitsubishi Electric began selling a DVD recorder with automated highlights-playback technology from MERL. The product feature was the first of its kind, received critical acclaim in the Japanese press, and enjoyed a significant commercial success in the Japanese market.

Specifically designed for sports programming, the detection of highlight scenes was enabled through analysis of the audio signal, including the timbre of the announcer's voice and the characteristics of audience sounds such as cheering and applause. The importance level for different segments of the recorded video was displayed as a graph, and the user could control the playback level, skipping the less exciting parts of the recording.

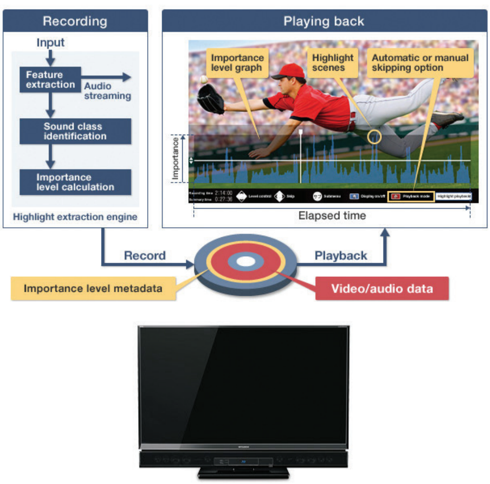

Digital Communication

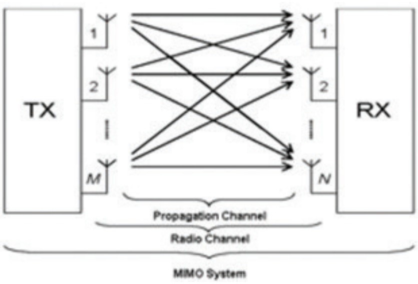

In the early 2000s, MERL focused on standardization and developments of new emerging technologies. Research topics included antenna selection, channel equalization, efficient channel state estimation, and Orthogonal Frequency Division Multiplexing (OFDM). Our antenna selection technology for Multi-Input Multi-Output (MIMO) systems was included in the 3GPP LTE release in 2008 and the IEEE 802.11n high speed WiFi standard in 2009.

We were a world leader in the development of Ultra-Wide Band (UWB) technology for short-range, high-definition video transmission (shown at CES 2005) and Impulse Radio UWB for accurate indoor localization. We also addressed system-level issues such as ad hoc mesh networking, wireless channel access and multi-hop routing for smart meter and sensor networks. In 2008, we shifted our focus to ultra-reliable wireless for machine-to-machine (M2M) communication that achieved very high reliability (3 orders of magnitude better than WiFi) with fixed low latency (<10ms).

MERL expanded our research around 2010 into optical devices and optical communication systems. We developed new algorithms to compensate for impairments introduced by optical fibers, novel high-dimensional modulation formats, and a new multi-subcarrier transceiver technology that for the first time in the world achieved a capacity of 1 Tbps using a single optical receiver.

Mesh and Zigbee Networking

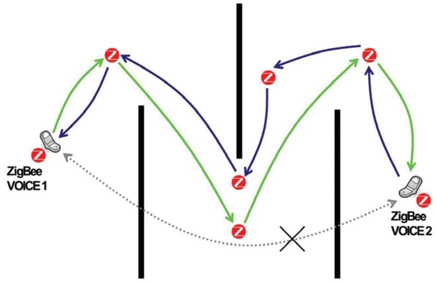

MERL's mesh networking research focused on lightweight, low-complexity network formation and routing algorithms allowing simple low-cost devices to efficiently form robust mesh networks. Such networks are the backbone of what is now called the Internet of Things (IoT).

Mitsubishi Electric was a founding member of the Zigbee® Alliance in 2003. This group developed one of the first standards for ad hoc self-organizing mesh networks. MERL researchers were key contributors to, and editors of, the first version of the protocol. Subsequently, MERL developed methods to control the communication delay as information traverses over the mesh, which enabled applications such as the streaming of voice and images over Zigbee® networks.

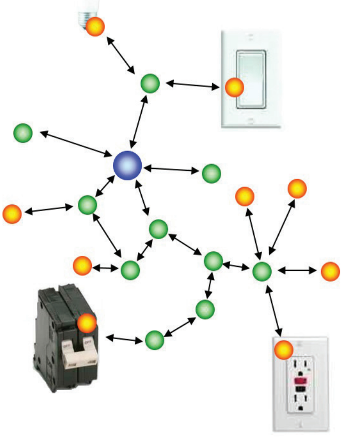

Over the course of several years, interest in IoT led to the rapid growth of applications and standards optimized to support these new applications. MERL researched routing methods that: conserve energy when nodes are battery powered, ensure redundant paths for messages carrying critical information and schedule duty cycles to ensure lower latency. Applications included smart meters for power, gas & water utilities, sensor networks for agriculture, and smart homes/buildings.

More recently MERL's focus shifted to ensuring the coexistence of multiple mesh networks and developing applications that make use of the data provided by these networks.

UWB Impulse Radio

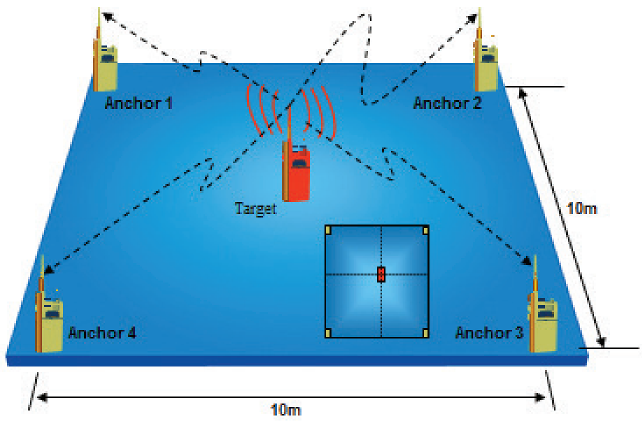

Ultra-Wide-Band (UWB) systems can simultaneously provide robust communication and accurate self-localization capabilities, making them ideal for wireless sensor networks and location-aware applications such as surveillance, infrastructure monitoring, and healthcare.

In 2007, MERL demonstrated the world's first UWB prototype using impulse radio technology to provide not only short-range wireless communication but also high-precision ranging. The system was compliant with the IEEE 802.15.4a standard and provided 110 kbps data rate with flexibility to scale up to 27.24 Mbps. The ranging used time-of-arrival sensing and a two-way ranging protocol to determine the distance between two communicating UWB nodes with an accuracy of 15 cm. Using the distance measurements from three reference nodes, the position of a UWB node in the network could be obtained and tracked with better than 50 cm accuracy.

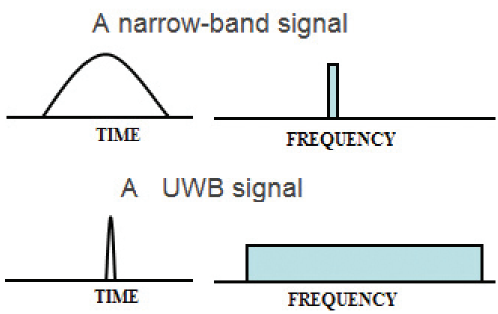

Ultra-Reliable Wireless

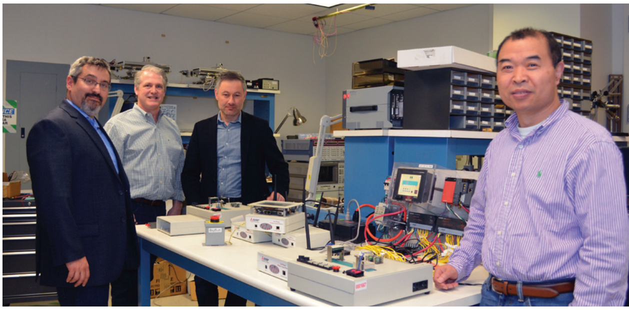

From 2008 to 2011, MERL researchers developed a proprietary wireless communication system, termed Ultra-Reliable Wireless (URW). This system achieves a message-loss rate of less than 10^{-8}, which is comparable to what can be achieved with a wired system. Through the use of a deterministic channel-access protocol, URW also ensures a delay of less than 10 milliseconds.

MERL worked on a range of novel wireless transmission techniques optimized for machine-to-machine (M2M) communication. By adapting ideas from multi-antenna cellular concepts to inexpensive, low-power M2M devices, we achieved performance approaching theoretical limits. Additionally, MERL has developed machine learning-based methods to detect and classify interference sources using commercially available WiFi receivers, enabling the deployment of interference avoidance methods on WiFi-based M2M devices.

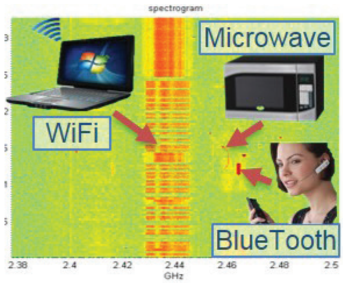

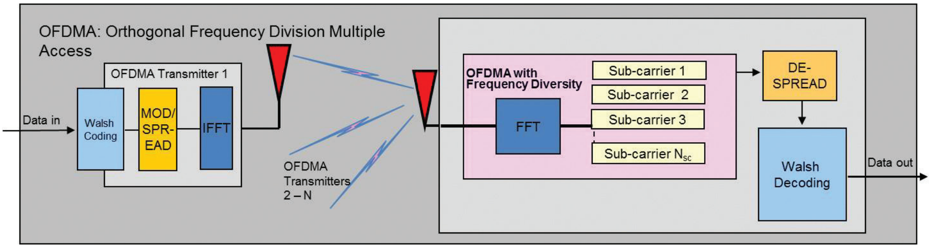

Optical Transmission & Error Correction

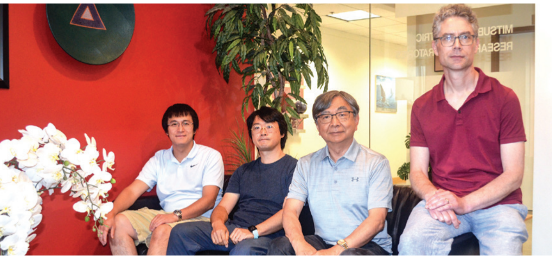

Data transmission using optical fibers forms the core of the Internet. To meet the continuing growing need for higher capacity, MERL researchers developed novel transmission schemes, signal processing methods, and error correction technologies.

The high intensity of the laser light constrained in the core of an optical fiber causes nonlinearity by changing the refractive index of the glass itself. MERL developed multi-dimensional modulation methods that optimize performance in the optical fiber in the presence of this very challenging problem.

Another method to improve transmission performance uses probabilistic constellation shaping, where the signal statistics are modified to match the ideal case more closely. MERL's research reduced the complexity of doing this allowing more practical implementation and partial compensation for nonlinearity.

MERL researchers have also been at the forefront of error correction codes for optical transmission, with work on Low-density parity-check (LDPC) codes leading to world-leading error correction performance. Our recent work on irregular polar codes has shown new ways to outperform existing codes with reduced complexity and power consumption while still maintaining high performance.

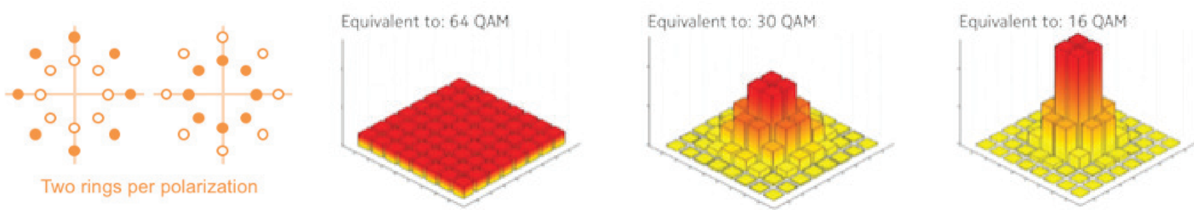

1 Tbps Optical Receiver

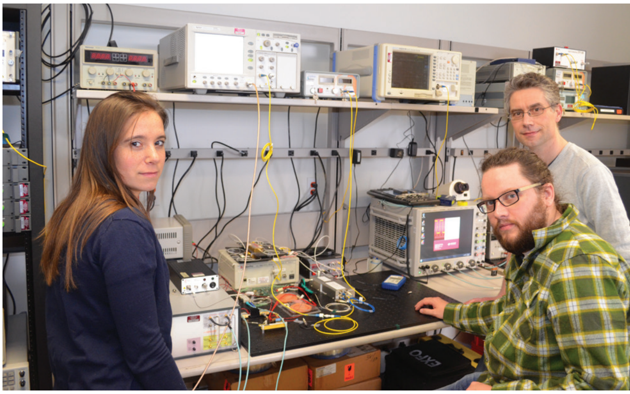

Optical fibers carry much of the world's data traffic, and the capacity requirements continue to grow by more than 20% annually to support connectivity among people and machines.

MERL researchers developed the world's first transceiver technology capable of 1 Terabits per second using a single optical receiver. It used a novel multi-subcarrier transceiver technology that experimentally achieved a spectral efficiency of 9.2 b/s/Hz, which was a world record at the time for a single receiver. A key feature of MERL's approach was the use of a pilot signal to enable high-quality compensation for signal impairment over a fiber. MERL's technology was compatible with previously installed optical fiber systems, allowing for a smooth upgrade path.

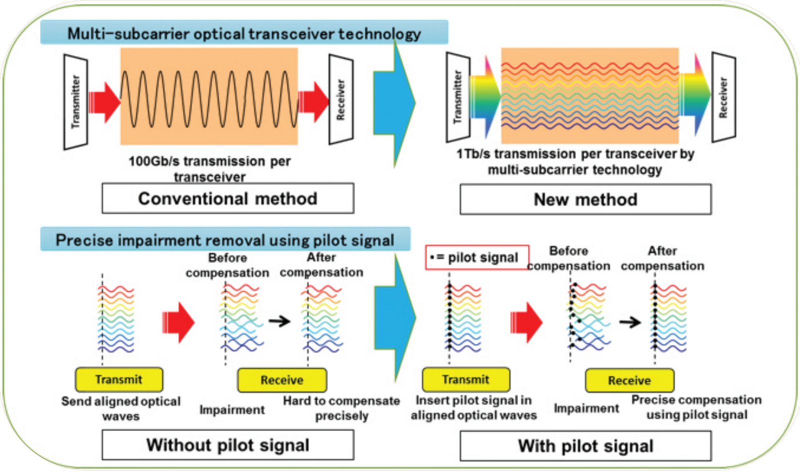

Model-Based Design

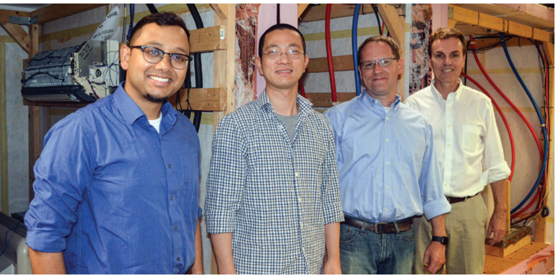

Products such as heating, ventilation & air conditioning (HVAC) systems, and robotic manipulators are heterogeneous, multi-physical systems-of-systems that are organized as a hierarchy of subsystems and components. Dynamic interactions among these systems demand new strategies and tools for modeling, analysis, and computer simulation to improve system performance and increase the efficiency of the design process.

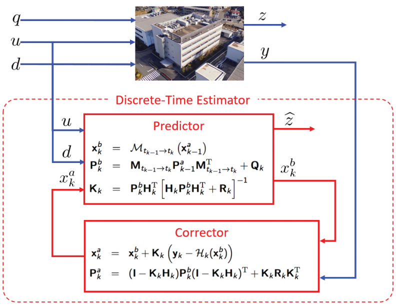

MERL researchers use the component-oriented, equation-based modeling language Modelica as a platform to solve complex dynamic and control problems, primarily for thermo-fluid and mechatronic applications.

Dynamic building models can provide valuable information about unmeasurable variables by combining the models with observed data. State estimation methods can then be used to learn thermal load behavior for use in advanced control or performance monitoring applications.

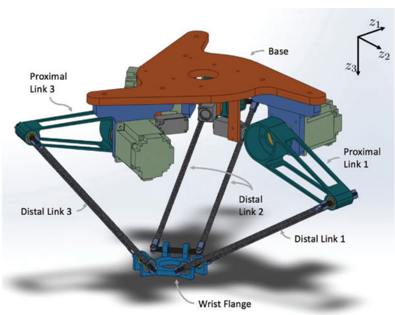

The delta-robot shown in the accompanying figure has properties, such as low mass, that make it well suited for robotic assembly operations. This requires development of new algorithms for force and impedance control. Modelica enables efficient dynamic modeling of the robot, its assembly task, and the nonlinear algorithms used to control it. This is useful for computer simulation, experimental validation, and control.

Heat Exchanger Path Optimization

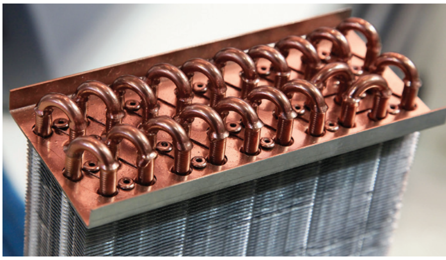

Tube-fin heat exchangers that transfer thermal energy between a fluid and air are ubiquitous in thermal management and energy conversion systems, ranging from HVAC systems to automotive applications to power plants. Designing the best path of fluid flow through a heat exchanger requires detailed simulation of the physical properties of the device, and the solution of a large mixed-integer optimization problem under numerous constraints that must be satisfied to realize solutions that are compatible with manufacturing requirements. Together, this is a prototypical model-based design problem.

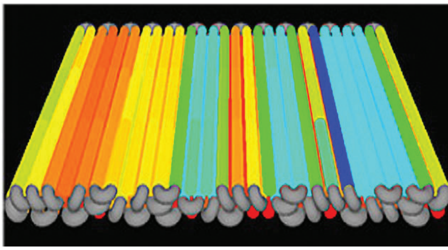

MERL used its expertise in both optimization and thermo-fluid systems to develop a decision tree formulation of the problem in the thermo-fluid domain, which enables the algorithm to capture the numerous geometric and manufacturing constraints in a parsimonious and efficient manner. This formulation was implemented as a custom branch-and-bound-based tool, which uses problem-specific information to improve the solution time over more general linear programming solvers (like CPLEX) by more than 20x, and reduces the fluid path design time from weeks to hours.

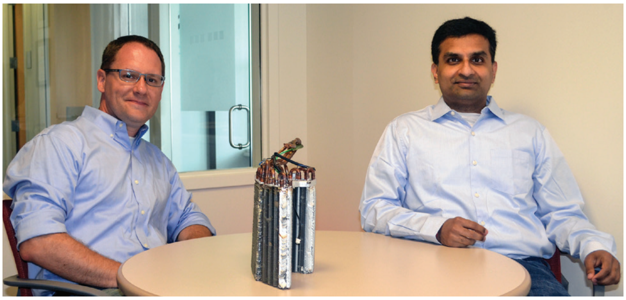

Data-Driven Optimization

Physics-based models of multi-physical systems are powerful because they have good extrapolative properties and their parameters often have physical interpretations, but these characteristics are often accompanied by nonlinear behavior and significant computational complexity. MERL is developing machine-learning methods to optimize these models in a robust and efficient manner so that their behavior corresponds to data obtained from field-installed equipment.

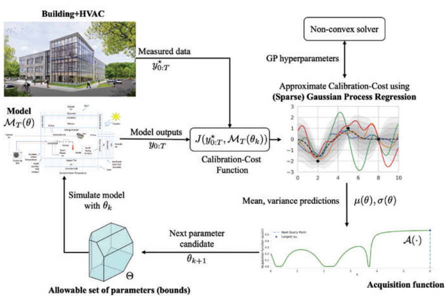

Fundamental research into Bayesian optimization has yielded new calibration methods for high-dimensional models of dynamically interacting HVAC equipment in buildings. MERL has developed new methods that can identify parameter values and uncertainty bounds in both the equipment and building models that meet standard accuracy requirements, using a number of iterations that is an order of magnitude smaller than typically required.

MERL research into the data-driven optimization of HVAC systems also seeks to utilize the information provided by simulations in a computationally- and data-efficient manner. Recent work in failure-robust Bayesian optimization methods has resulted in approaches that are able to use the information provided by incomplete simulations to cut the calibration time in half, while new batch Bayesian optimization methods that use attentive neural processes can parallelize the calibration process and further reduce the time required.

Modeling & Design for Electric Motors

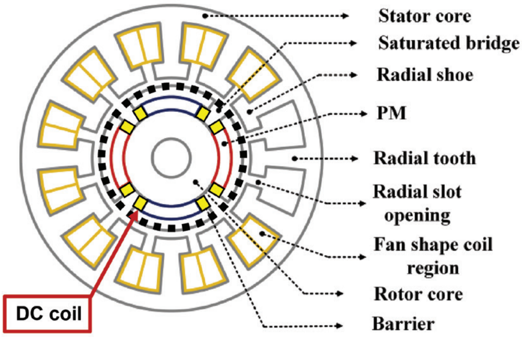

The modern motor design process depends heavily on finite-element analysis (FEA). Although FEA is becoming well-developed and can have excellent accuracy, it is computationally intensive and time-consuming, and therefore not suitable for large-scale motor design optimization. An alternative is analytically based magnetic field modeling techniques, which enable reduced use of high-fidelity FEA simulations and speed up the design analysis and optimization process. We have developed a subdomain-based method for the calculation of the instantaneous magnetic field distribution for interior permanent magnet synchronous motors. This method has accuracy close to FEA simulations while using a small fraction of the calculation time.

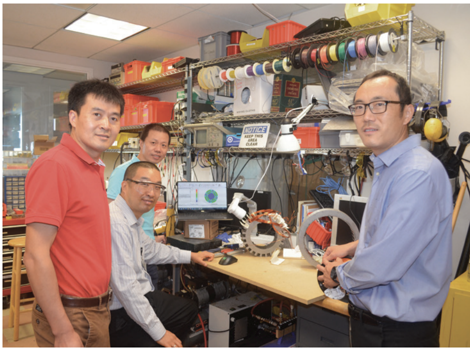

Our physics-based methods also facilitate detailed harmonics analysis and help us understand the mechanisms of new flux modulation motors, such as Vernier motors. We have proposed a direct-drive motor combining both radial and axial flux Vernier permanent magnet motor components. Initial simulation results show over a 1.5x torque improvement compared with a benchmark motor of the same size.

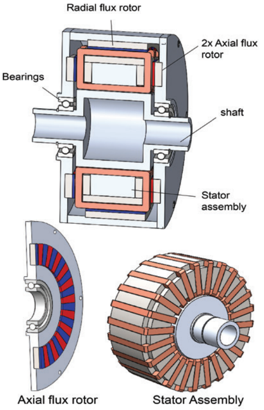

Fault Detection for Electric Motors

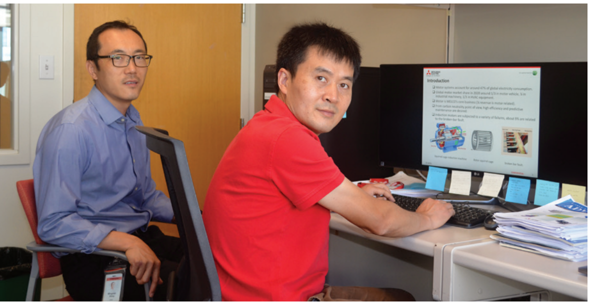

Electric motors are widely employed in a variety of industry applications and electrified transportation systems. These machines may operate under unfavorable conditions (such as high ambient temperature, high moisture or overload) which can eventually result in motor malfunctions that lead to high maintenance costs, severe financial losses, and safety hazards. We have been developing technologies for motor condition monitoring to detect faults at an early stage and permit maintenance scheduling.

Using an accurate thermal circuit model we developed for permanent magnet motors, we designed observers that use limited measurement data and can monitor the temperature of permanent magnets to avoid overheating and irreversible demagnetization. We have been developing quantitative models for bearing faults of induction motors to distinguish different types of faults and quantify their severity level using motor current signature analysis without additional sensors. We have also been evaluating the application of machine learning and deep learning for motor fault detection and classification.

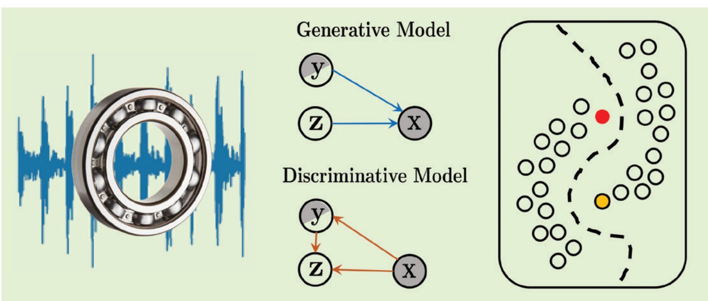

Devices and Circuits

We explore various device technologies and device & circuit architectures to dramatically improve RF device and radio system performance to achieve higher efficiency, high linearity, and flexible operation. We develop novel photonic integrated circuits to improve performance and reduce cost in optical communications applications. We combine physics modeling and simulation with modern machine learning techniques to allow rapid design of devices that obtain target specifications for metamaterials and photonic and solid-state semiconductor devices and circuits. These design methods are used by Mitsubishi Electric's Semiconductor & Device production facilities in Japan to fabricate state-of-the-art components that are key parts of today's high technology services world-wide.

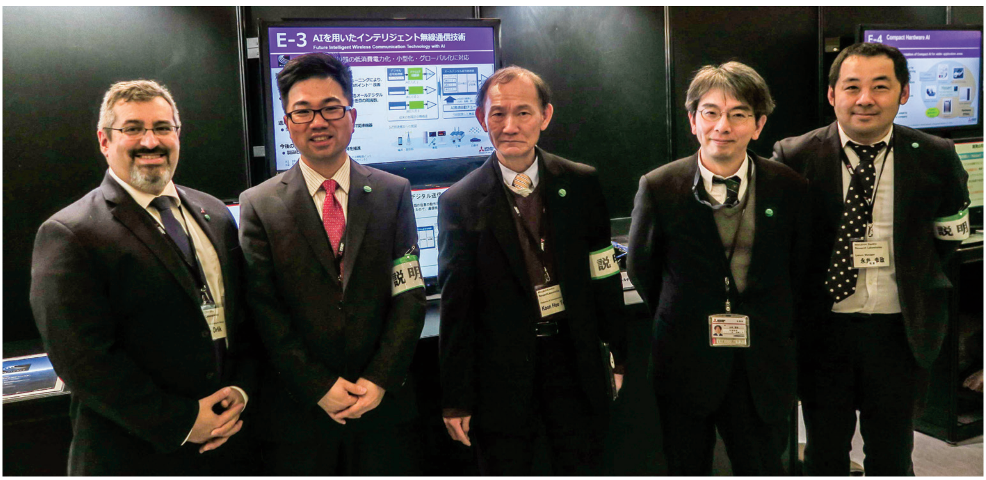

GaN Devices and Power Amplifiers

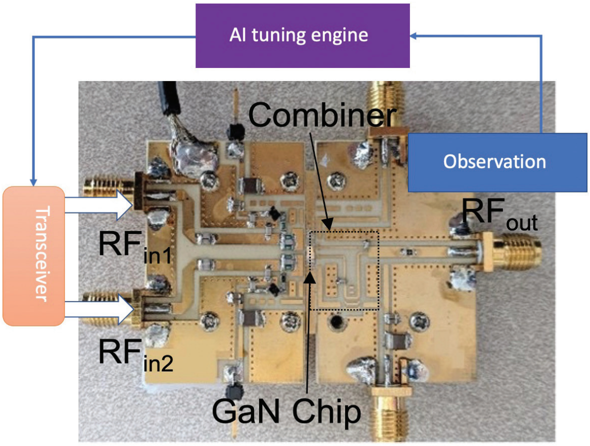

Since 2011 MERL has been investigating Gallium Nitride (GaN) devices to take advantage of their high-power density and high frequency/switching-speed capabilities. We pioneered the application of machine learning techniques for operating condition optimization of advanced RF power amplifiers.

We explored 3D multi-channel architectures for GaN devices to push the boundaries of their performance and to broaden their applications. We leveraged the properties of piezoelectric and polarized doping of GaN to develop 3D device structure designs for applications in RF and quantum computing electronics. Preliminary simulation work indicates that for applications in RF power amplifiers, noticeable improvement in GaN device parameters such as power density, frequency, and linearity can be obtained.

We developed several cutting-edge power amplifier circuits for 4G and 5G wireless infrastructure applications, including wideband class-J at 2GHz and AI-enabled digital Doherty amplifiers with significantly enhanced performance compared with conventional designs in terms of bandwidth, efficiency, and linearity.

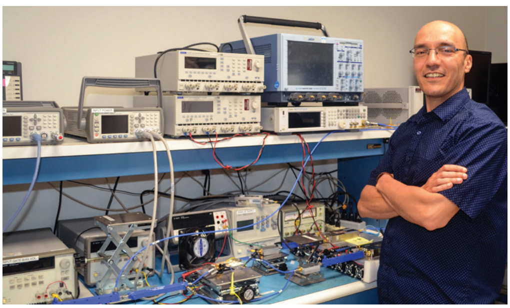

All-Digital Transmitter and Array

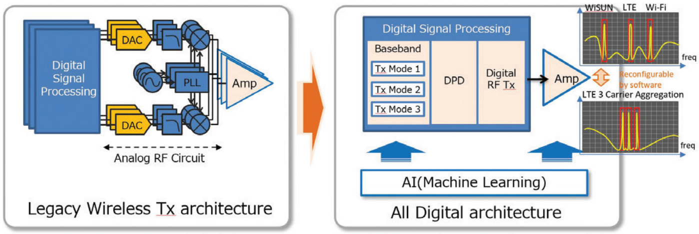

In 2018, MERL announced that we had developed the world's first wireless communication technology using an advanced All-Digital Transmitter that can use various wireless communication access-mode standards while transmitting up to three arbitrary operating frequency bands simultaneously on a single digital circuit in real-time. This novel technology is expected to significantly reduce the power consumption, improve the miniaturization, and support the globalization of multiband IoT equipment for multiple communication standards and regulations.

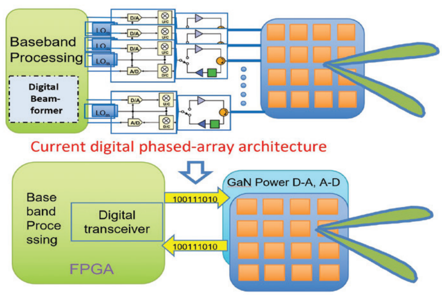

Proposed all-digital phased-array architecture

More recently, MERL has developed an All-Digital Phased Array technology which further extends our all-digital transmitter technology towards future directional communications and radar applications. Our FPGA-based prototype demonstrated 50% fractional bandwidth centered at 2.5GHz.

DNN-Aided Optical Device Design

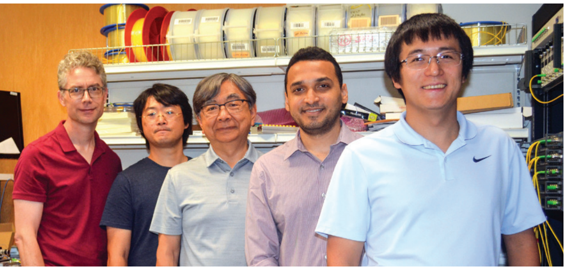

Nanophotonic devices are very compact optical devices with a size of several square microns, very low loss, and compatible with high density photonic integrated circuits. The applications include optical communications, sensing, and optical signal processing. However, the challenge is to optimize a structure which has hundreds, and often tens of thousands, of parameters.

In 2011, we started developing novel optical device design techniques with advanced multi-parameter optimization methods. And since 2017, we have been working on nanophotonic device design assisted by machine learning. We applied MERL's state-of-the-art Deep Neural Network (DNN) training algorithm (adversarial conditional variational autoencoder with cycle consistency) to nanophotonic device design.

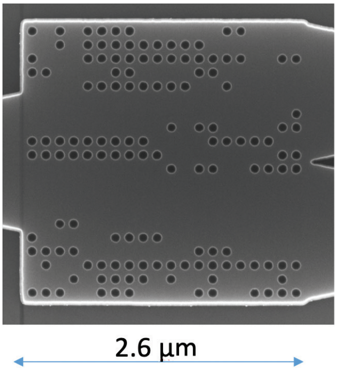

In many cases, we need to design many devices with slightly different target performance (such as splitting ratios in the case of power splitters, or target peak/valley wavelengths in the case of wavelength splitters). The deep neural network, once trained from a small number of training data samples, can generate devices with arbitrary target performance. This is much more efficient than optimizing devices for each target performance from scratch. Prototype devices were fabricated using an external foundry, and performance testing was conducted using facilities at MERL. Results have shown the DNN-assisted methodology provides the ability to design a variety of devices while meeting or exceeding performance targets.

Computational Sensing

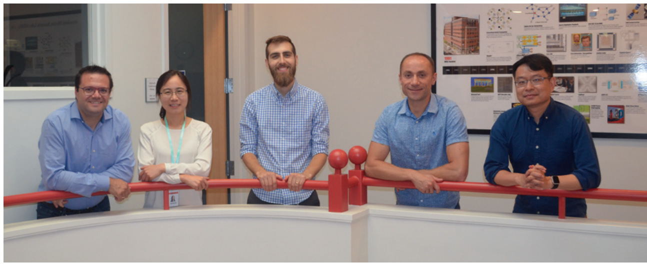

Sensing and signal acquisition are indispensable technologies in our increasingly digital world. At MERL, we exploit the ever-increasing availability of inexpensive computational power to overhaul the signal acquisition paradigm and significantly improve our sensors and their capabilities.

We aim to fundamentally understand and develop models able to describe the signals we desire to sense, their behavior and propagation in the environment, and the operation and properties of the sensing systems. Using this knowledge, we develop reconstruction algorithms to recover signals or the information of interest, with significantly improved fidelity and robustness.

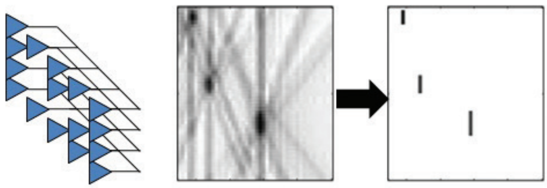

We have applied our methods to advance the state of the art in radar, lidar, ultrasonic, acoustic, and optical. We have successfully demonstrated significant improvements in applications that include high-accuracy low-cost active depth sensing, fusion of signals from multiple sensors with different modalities, radar and ultrasonic imaging, sensing of ionosphere density, automotive radar, transportation security, underground imaging and infrastructure monitoring, airflow sensing, THz imaging, and sensing using Wi-Fi signals.

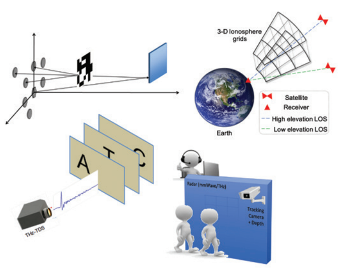

Compressive Synthetic Aperture Radar

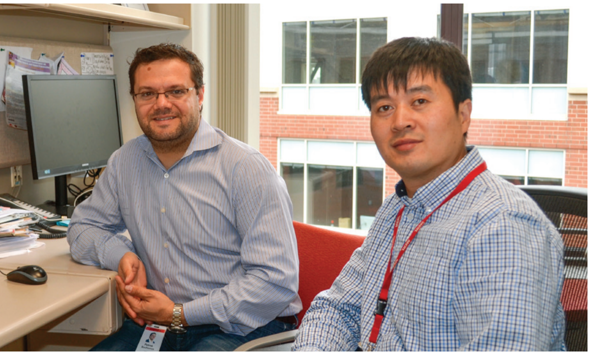

MERL pioneered several new approaches to synthetic aperture imaging based on compressive-sensing techniques. These methods have garnered substantial attention by the research community and have the potential to revolutionize the future of remote sensing systems.

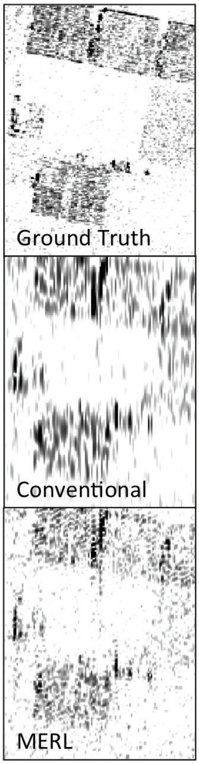

Conventional radar imaging systems transmit pulses and receive them at a uniform rate, which dictates a trade-off between resolution and imaging area. MERL's approach substantially improves this trade-off through non-uniform designs.

One of our first designs demonstrated substantially increased azimuth resolution using a high-rate non-uniform pulsing scheme, allowing interference between pulses. The same design can be configured to double the observation area instead. Rather than alter the pulsing rate, follow-up designs utilized non-uniform beam-steering patterns to achieve similar improvements in imaging resolution or observation area compared to conventional acquisition modes.

MERL's work in this area was recognized with the 2014 Symposium Prize Paper Award by the IEEE Geoscience & Remote Sensing Society.

Infrastructure Monitoring

Monitoring the structural integrity of public infrastructure such as roads, bridges, and tunnels is becoming increasingly important and costly as these structures age. MERL is developing radar-based imaging technology able to see inside and monitor the condition of aging infrastructure.

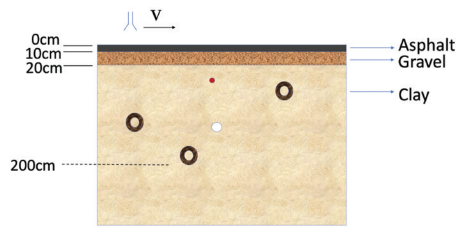

Radar reflections from inside complex materials contain information about the interactions between the pulses and the parts of the material through which they propagate, and, therefore, on the structure of that material. The challenge is to extract this information from the reflected pulses and use it to determine the internal structure of the material.

MERL has developed fundamental methods for exploiting the secondary scattering of the incident pulse inside the structure as it interacts with different parts of the material as a source of more information about the structure of the material and we use it to improve image quality. Our algorithms can produce high resolution images of internal structures even when the dielectric contrast ratio is very large, a common failure case with existing approaches, and a common scenario in infrastructure monitoring applications.

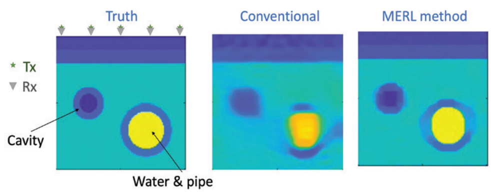

Terahertz Sensing

Terahertz (THz) frequencies are a relatively new, underexplored modality for sensing applications. They allow very high-resolution, coherent time-of-flight imaging, while penetrating certain kinds of non-conducting materials. Industrial applications include contactless elevator/train positioning encoder, material/coating inspection, non-destructive evaluation, air quality monitoring, and security scanners.

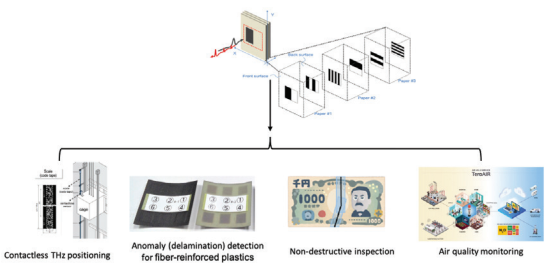

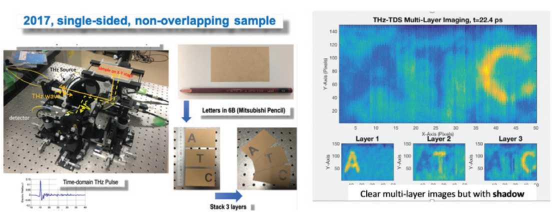

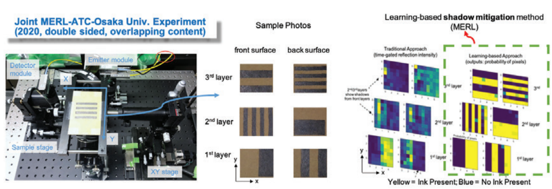

MERL has developed learning-based THz multi-layer content extraction methods to address challenges such as: 1) depth variations from one pixel to another due to either the irregular sample surface or the vibration from the mechanical scanning process; 2) shadowing effects caused by non-uniform penetrating illumination from front layers to deep layers; and 3) limited capability to recognize content in the back surface of each layer.

Optical Tomography & Depth Sensing

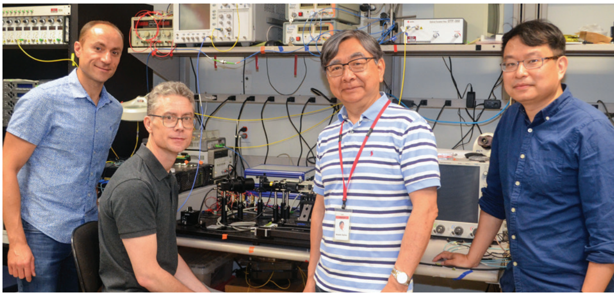

Optical sensing methods, such as optical coherence tomography (OCT), have been extensively used for medical imaging, but with high cost. MERL is researching system architectures and algorithms that can expand the application areas (for example, to industrial applications), while reducing cost.

The most expensive part of an OCT device is the optical source (usually a laser whose frequency is swept in a very linear fashion over a broad bandwidth, to obtain excellent resolution and accuracy). MERL has developed new signal processing algorithms that can compensate for nonlinearity in the swept frequency, allowing cheaper sources to be used, while maintaining high levels of accuracy. We have also developed new architectures using sparse stepped frequency sources with advanced reconstruction algorithms to recover an image.

Most recently, we have developed a new line-field sensing method using a cheap LED source and a CMOS camera sensor, to provide depth information over a line simultaneously (rather than at a single point) providing micron scale resolution and accuracy with ultra-low cost.

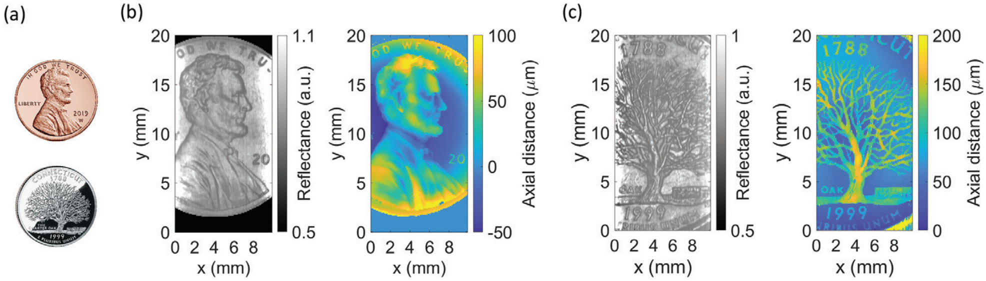

CM-Level GNSS Positioning

Global navigation satellite systems (GNSSs) are used in positioning and navigation applications worldwide. The GNSS estimation problem consists of both continuous and discrete variables and can include hundreds of measurements. MERL is researching novel real-time estimation methods and architectures that improve GNSS positioning accuracy from meters to centimeters.

By leveraging MERL fundamental research in statistical signal processing, optimization, and dynamical systems, MERL developed structure-exploiting particle filters that enable real-time centimeter-level accuracy in presence of disturbance uncertainty and an abundance of measurements. In addition, we developed sequential nonlinear mixed-integer estimation methods based on statistical linearization and optimization, suitable for low-cost processors as well as learning-based methods for spoofing attack detection and prevention.

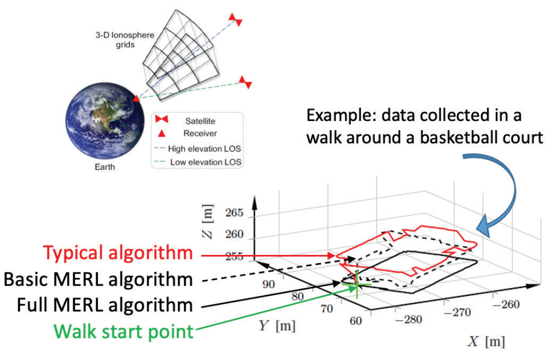

Further, MERL has developed tomographic methods for determining the distribution of electron density in the ionosphere. Given that information, one can correct for the position errors caused by differing GNSS signal delays as they pass through different parts of the ionosphere, further improving positioning accuracy.

IoT and Connectivity

The rapid increase in the bandwidth of cellular, WiFi, and Bluetooth networks coupled with an even more dramatic drop in the cost of communication hardware has enabled an ever-expanding range of communicating devices, the so-called "Internet of Things" (IoT).

Early MERL work focused on ad hoc and self-organizing networks for sensing and monitoring. We developed novel methods for the efficient encoding and transmission of short messages for highly reliable machine-to-machine (M2M) communication. We are exploring asymmetric communication and support for networks with limited computational power and battery capacity.

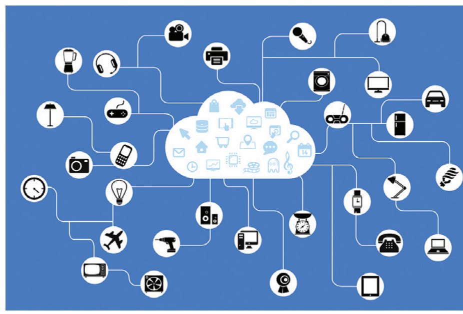

Recently, with the continued expansion and adoption of wireless systems for IoT networks, we have focused our research efforts on the study of coexistence mechanisms, particularly for smart meter applications. MERL is also developing new technologies based on machine learning and exploiting the numerous signals and data collected by IoT systems to infer the location of objects or people, count the number of persons in a room, or determine the activities of building occupants. These technologies have the potential to improve user comfort and the energy efficiency of buildings.

MERL's research includes connectivity over larger geographic scales, such as smart cities and intelligent highway systems, where we explore how the sharing of information can reduce regional traffic congestion and pollution.

Smart Meter Coexistence

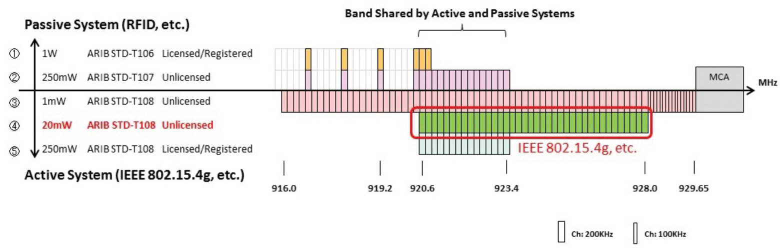

It is inevitable that the proliferation of communicating devices has led to problems of interference between them. An example of this is smart electric meters that communicate using IEEE 802.15.4g in the Sub-1 GHz frequency band where the amount of spectrum is severely constrained and many other kinds of devices are attempting to communicate using different standards such as IEEE 802.11ah.

MERL initiated and led the development of the Sub-1 GHz band smart meter coexistence standard IEEE 802.19.3, which was published in April 2021. We made several major technical contributions to the standard to support improved coexistence between IEEE 802.11ah systems and legacy meters. These include traffic prediction methods that allow meters to prevent interference before it happens and improved methods to estimate the number and type of interfering radios in a given area.

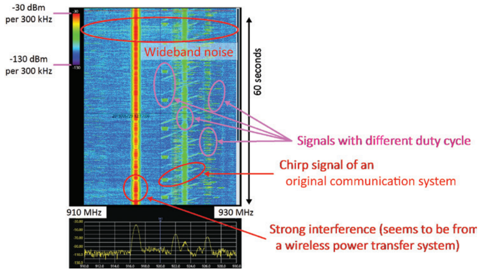

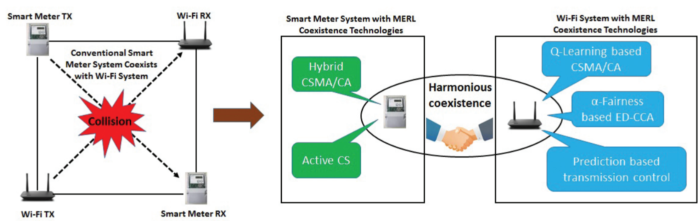

Indoor Localization & Wireless Sensing

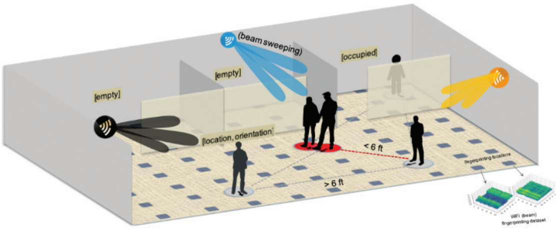

Our indoor daily life is surrounded by wireless signals (Wi-Fi, Bluetooth, LTE/5G). We can leverage these ambient wireless signals to localize people and sense their activity indoors. MERL's research exploits existing wireless standards operating in millimeter-wave (mmWave) bands for fingerprinting-based indoor localization and wireless sensing. We perform inference on the channel measurements used to optimize the radio link to determine locations of people and to recognize human actions. Our approach requires no special hardware nor any additional overhead on the network.

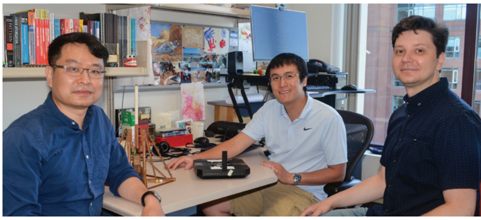

Currently we are working on fusing measurements from lower frequency (<7GHz) systems together with measurements from mmWave systems. Overall, our approach can achieve a decimeter-level accuracy for indoor localization. We are actively participating in the emerging IEEE 802.11bf (WLAN Sensing) standards body where our methods are currently under consideration for inclusion.

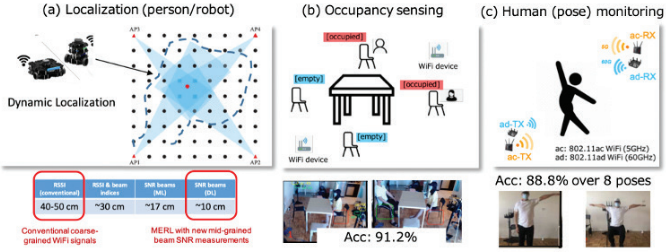

Internet of Vehicles

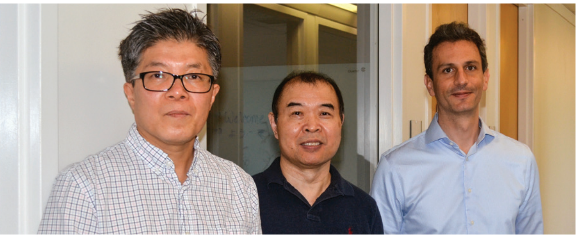

Research on the Internet of Vehicles (IoV) seeks to improve safety, traffic efficiency, fuel efficiency, and individual driving experience. However, implementing IoV systems in the real world is difficult. Cloud computing cannot realize real-time vehicle control due to long communication latencies. However, on-board control does not have sufficient data to make globally optimal decisions. MERL has developed zone-based vehicle control models, in which road sections at highway merging points or intersections are divided into zones where information about arriving vehicles is gathered and closer zones where edge nodes with more complete situational knowledge compute vehicle trajectories that are optimized for fuel efficiency and driving comfort.

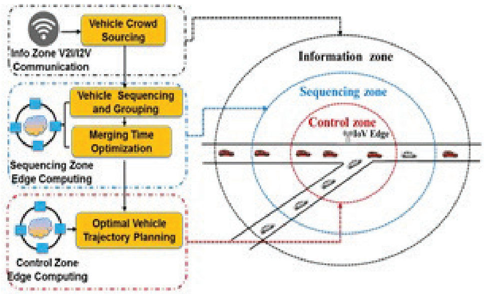

Recent efforts have developed multi-task federated learning models to address data privacy issues, communication bandwidth limitations, and data heterogeneity issues. IoV edge nodes act as learning servers and data collectors while vehicles act as learning agents. Servers and agents perform federated learning to predict travel time and determine optimal routes.

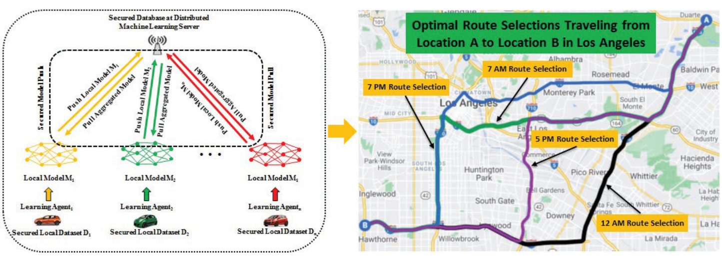

Compact AI for Constrained Environments

MERL has developed several tools that enable AI applications that are fast and efficient and can be implemented in lightweight environments, such as edge computers.

Fundamental work at MERL on low-dimensional and quantized embeddings has shown that it is possible to significantly reduce the dimensionality of the data needed for AI learning and inference tasks while still preserving the most relevant information. This saves a lot of space but more importantly greatly reduces the complexity of using the data.

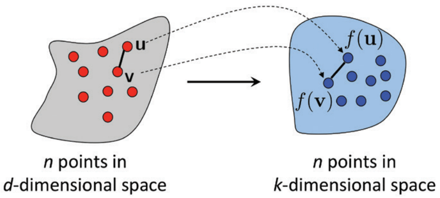

We have demonstrated that the detailed relationships between similar data points represent most of the information that is important for learning and inference and should be accurately preserved. In contrast, the detailed relationships between dissimilar data points are not as important and can be significantly distorted as long as the points remain dissimilar. This allows significant flexibility in designing a system to reduce complexity and storage cost. Furthermore, this flexibility allows for designs with provable information-theoretic privacy-preserving guarantees, important in privacy-sensitive applications.

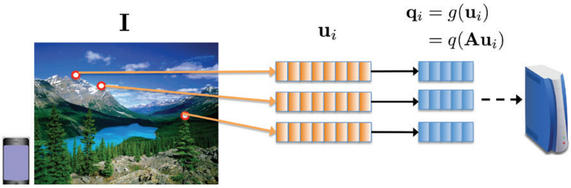

Optimization

Optimization is a fundamental technology that is needed whenever optimal decisions or settings are needed to improve a performance metric. MERL has developed advanced optimization algorithms to attack hard industrial problems such as voltage & power optimization in electrical power networks, control of autonomous vehicles & satellites, temperature set-point scheduling for air conditioners, group elevator scheduling, free-form lens design, and vehicle routing.

For optimization problems with continuous decision variables, such as run curve and feeder voltage optimization for electrified trains, we have been able to solve very large problems in a fraction of a second, even for non-convex situations.

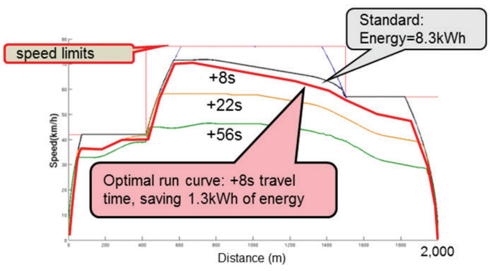

The Parallel Quadratic Programming (PQP) algorithm developed at MERL can solve convex quadratic programs with millions of variables in a fraction of a second.

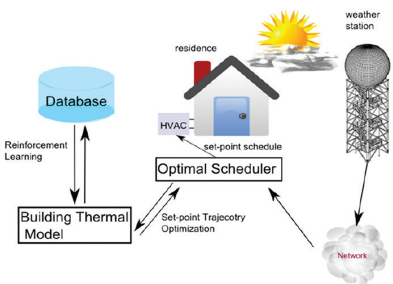

For problems with discrete variables, we have developed representation methods based on decision diagrams that can be reduced to compact Mixed Integer Programs (MIPs) and solved efficiently by modern MIP solvers. This technology is also applicable to problems such as vehicle scheduling and heat exchanger design optimization.

Efficient Optimization Algorithms

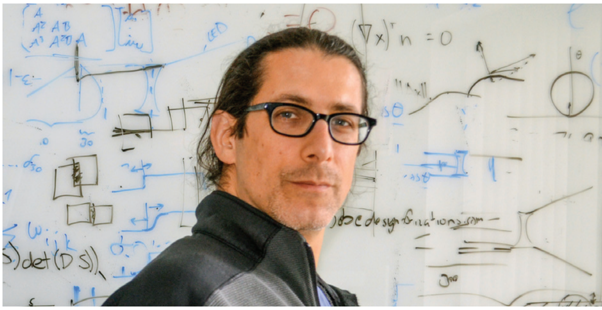

Much of MERL's research activity involves reformulating scientific and engineering problems as optimizations. We develop efficient new algorithms for canonical optimization problems, because an advance in that area can push many other technologies forward, making heretofore infeasible approaches practical.

A prime example of such an advance is Parallel Quadratic Programming (PQP), which solves a broad class of quadratic programs. PQP offers fine-grained parallelism and can solve very large problems with unusual speed. As a generic optimization procedure, PQP set the stage for new results at MERL in optimal control, radiation-therapy planning, image de-blurring, motion planning, and computer vision.

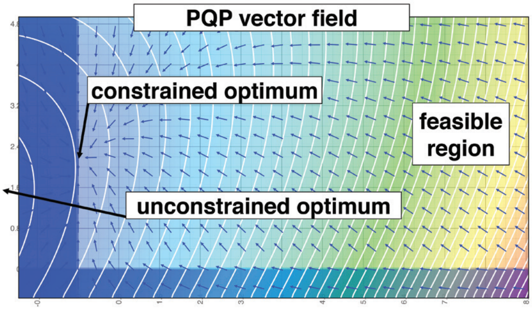

MERL has also developed efficient algorithms and theoretical bounds for other canonical problems and techniques including the first sub-cubic-time incremental singular value decomposition (SVD), bi-criterion minimum-cost path, alternating direct method of multipliers (ADMM), hyperbolic programming, and numerical preconditioning.

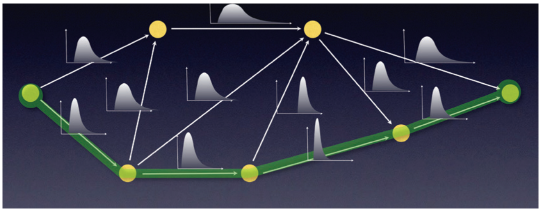

Railway Power Optimization

Spaced along an electrified railway, electric substations supply power to the overhead wires. Trains draw power from these wires to run. With regenerative breaking, trains can put power back into the overhead wires to be used by other trains. However, the ability to do this is limited in standard electric railway systems.

The key problem is that standard systems maintain all the substations at a constant voltage near the maximum voltage the wires can support. This significantly limits the amount of power that can flow from a breaking train to an accelerating one. Mitsubishi Electric developed equipment that can vary the voltage at each substation in real time. MERL developed an algorithm to optimize the voltage at the substations so that the flow of power from one train to another can be maximized. In a simulated experiment, this enabled a 5% reduction of total railway energy consumption.

Transportation Scheduling

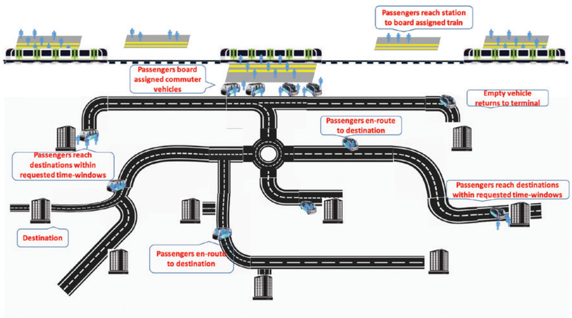

Ride-hailing services have expanded the role of shared mobility in passenger transportation systems, creating new markets and creative planning solutions for major urban centers. Coordinating the last-mile passenger transportation with a mass transit service can provide a seamless multimodal transportation experience for the user. Any system that provides passengers with predictable information on travel and waiting times in their commutes is immensely valuable.

The main objective of MERL's solution is jointly scheduling passengers on mass transit and on small ride-sharing vehicles so that long-term planning, maintenance, and environmental impact are all considered. To this end, the system balances minimizing travel time and minimizing the number of trips taken by the last-mile vehicles. MERL developed a novel optimization approach based on decision diagrams that can solve instances of real-world size (10,000 passengers spread over an hour, with 50 destinations and 600 ride-sharing vehicles) in a computation time of just 1 minute, which is orders-of-magnitude faster than other methods appearing in the literature.

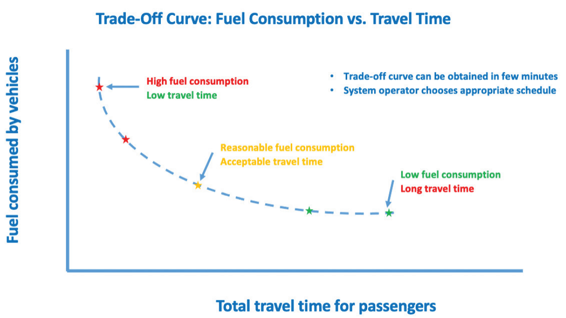

Decision Diagram Optimization

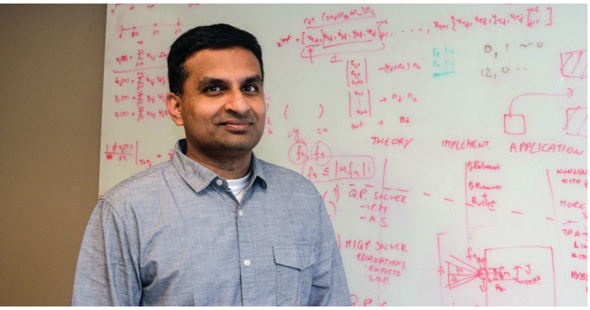

Many combinatorial optimization problems embed a sequential decision-making structure. Exploiting such a structure in the modeling can reduce the computational time for solving these problems when compared to standard modeling approaches.

The key innovation is to realize the existence of the structure and develop a model tailored to the problem that allows fast computation. MERL has successfully applied this approach to problems from three different application domains. Two domains (transportation scheduling and heat exchanger design)

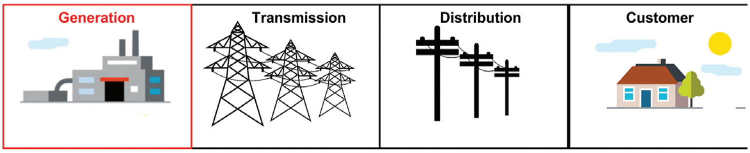

are described on other pages of this book. The third is the scheduling of thermal electrical generators (known as unit commitment).

The unit commitment problem is deciding which generator to turn on when in the face of fluctuating electricity demand and involves many constraints related to how generators operate. MERL's approach outperforms state-of-the-art algorithms and runs in reasonable time on problems involving as many as 2,800 generators.

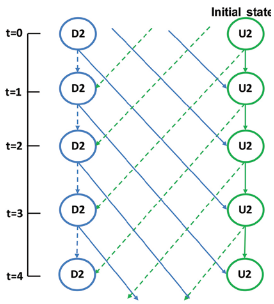

Freeform Optics

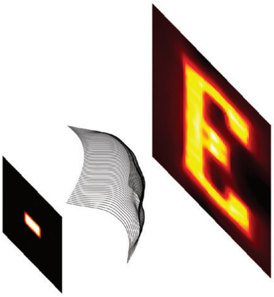

In theory, by using an appropriately shaped freeform lens or mirror one can convert a collimated light beam into almost any desired pattern of light. This can allow a single freeform optic to replace several conventional elements in an optical system, paving the way for smaller, lighter, and more efficient devices. Advances at MERL in the mathematics of optimal transport made it practical to compute what shape is appropriate.

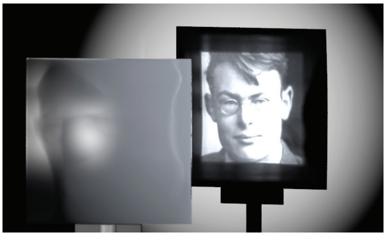

However, the theory behind our initial solution assumed that all the light originates from a single infinitesimal point light source. Until recently, it was not known how to design a freeform optic for an extended light source, and it was widely conjectured that blurriness was unavoidable. MERL research resolved both of these questions: we introduced a universal solution method for designing a freeform optic for any field of rays and demonstrated that such a freeform could generate illumination patterns with sharp features even when the light source is large relative to the optic. This makes it possible to design high-efficiency freeform illumination optics for high-power LEDs, which can be several millimeters across.

The technology enables much higher efficiency in areas such as municipal lighting, vehicular illumination, task lighting, projective signage, and laser beam shaping.

Spacecraft Trajectory Design

MERL has developed a framework for designing minimum-fuel, electric propulsion transfer orbits from geostationary transfer orbit (GTO) to geostationary Earth orbit (GEO). Electrical propulsion is more efficient than conventional, chemical propulsion but produces smaller amounts of thrust. Low-thrust transfers need to engage thrusters for long periods of time and therefore must take into account the free dynamics of space as well as the effect of non-impulsive thrust. MERL's framework consists of four steps to determine a minimum-fuel solution.

MERL has also developed a framework for the design of station-keeping on near-rectilinear halo orbits (NRHOs) around the Earth-Moon Lagrange point L2. NRHOs are in theory closed orbits, but in reality are never truly closed. For this reason, MERL has proposed a scheme where segments of NRHOs are patched together and the spacecraft tracks each segment using a conventional control approach and performs minimum-fuel transfers between segments. The resultant fuel consumption is 0.083m/s/year. This is important for the instantiation of a space gateway—the Earth-Moon L2 is the place from which the least energy is needed to leave the Earth-Moon system and, for this reason, it is desirable to have a permanent presence nearby. L2 is behind the Moon, but an NRHO around L2 can have constant line-of-sight with the Earth and this is supported by MERL's framework.

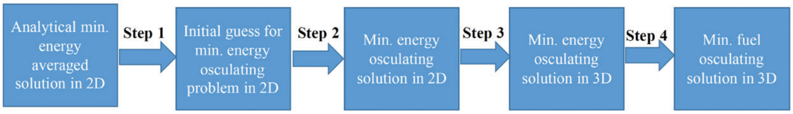

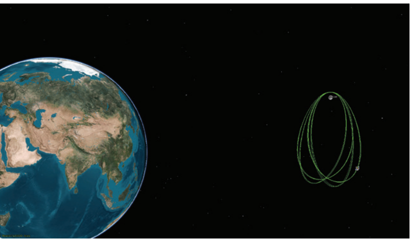

Control

Control systems are ubiquitous in today's smart devices, constantly regulating a device's behavior to achieve desired responsiveness, efficiency, and consistency. MERL is researching novel control methods and investigating their impact in relevant applications for Mitsubishi Electric Corporation and the future.

In particular, MERL is developing Model Predictive Control (MPC) for constrained multi-variable systems to provide performance guarantees even in presence of uncertainty and optimization algorithms suitable for low-cost microprocessors. MERL's research in constrained control also includes reference and command governors and the construction of control invariant sets for uncertain systems.

MERL control research spans optimal and robust control, nonlinear control, statistical estimation and control of stochastic systems, and the integration of model-based and data-driven methods for control design, all aimed at providing guarantees of performance and safety.

Key applications where MERL control methods are proving their effectiveness include automotive safety systems, electric vehicle management, HVAC systems, satellite station keeping, spacecraft rendezvous, railway systems, servomotors, laser processing machines, robotic manipulators, and elevators.

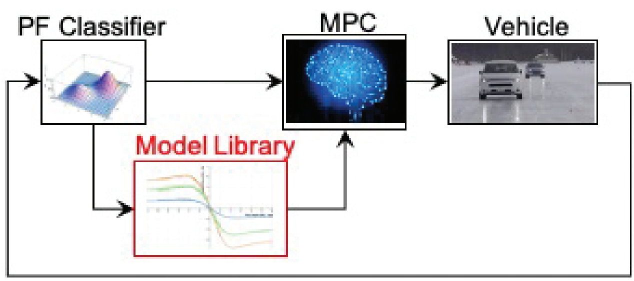

HVAC Control

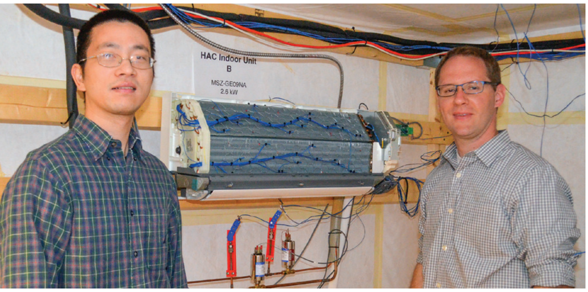

Today's Heating, Ventilation, and Air Conditioning (HVAC) systems have variable speed compressors, adjustable expansion valves, and variable speed fans, and as a result, are highly interactive, dynamic systems that couple in non-intuitive ways to the thermal and fluid dynamics of the building environment.

MERL develops next-generation robust control architectures for HVAC systems using a model-based development process based on the Modelica modeling language. MERL has made fundamental improvements in the modeling of vapor compression systems and in efficient control algorithms to make this possible.

MERL uses detailed models to analyze the dynamic behavior of systems and design robust and predictive controllers. These are experimentally validated in MERL's laboratory, which is capable of conducting transient experiments such as those shown in the accompanying figures, before being transferred to Mitsubishi Electric's research and development laboratories in Japan.

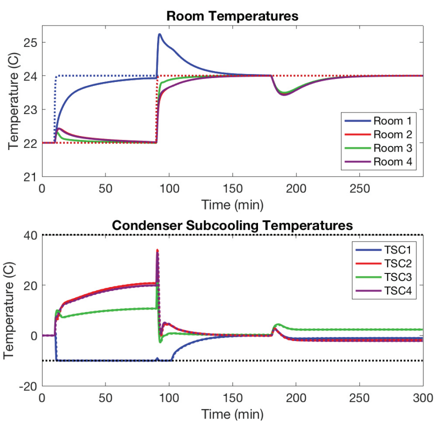

MERL's research in this area began with Room Air-Conditioner (RAC) products and our controllers are in use in more than 1 million systems sold per year. MERL has developed a new control architecture for a multi-zone Home Air-Conditioner (HAC) product that uses Model Predictive Control (MPC) to enforce critical constraints while also achieving good disturbance rejection properties.

High-Speed Laser Cutting

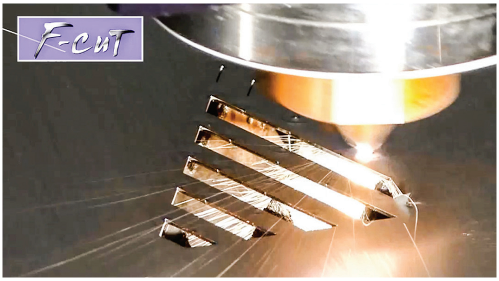

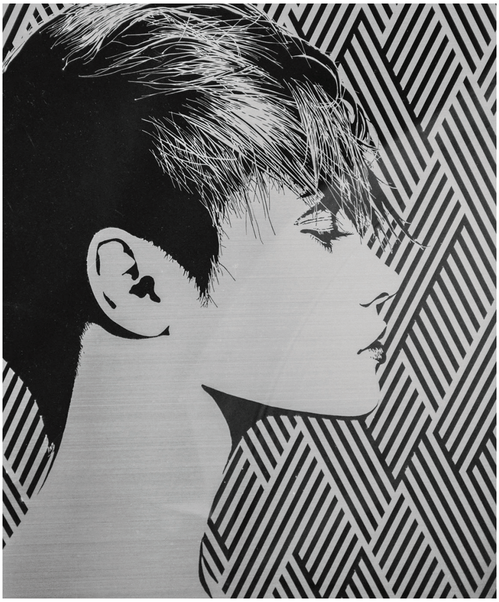

Mitsubishi Electric's laser cutting machine uses a high-power laser on a moving head to rapidly and accurately make cuts in sheet metal. It can be used to make a wide range of flat metal parts. To maximize the value of the machine, one would like to make the desired cuts as fast as possible.

The time it takes to make a set of cuts depends on the order the cuts are made in and the path the cutting head takes between cuts. At its mathematical core, determining the best cutting order is an instance of the Traveling Salesman Problem (TSP). However, it is a particularly complex TSP because the physics of the way the head can move also has to be taken into account. For example, paths between cuts must be curved to respect the acceleration properties of the head. MERL developed a fast algorithm to solve this mix of discrete TSP and continuous optimal control, which generates near-optimal paths in a much shorter time when compared to conventional approaches.

Spacecraft Guidance and Control

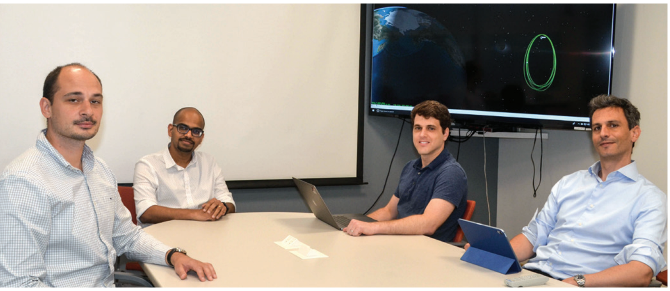

Spacecraft guidance and control (G&C) maneuvers are traditionally computed on the ground, implemented open-loop, and require ad hoc corrections against errors. As the number of spacecraft in orbit is poised to skyrocket, scalability of traditional G&C is a challenge. Additionally, as spacecraft venture further out into the solar system on increasingly ambitious scientific missions, the need for autonomous, on-board, closed-loop G&C capability is growing.

MERL has developed several advanced, low-computational-burden spacecraft G&C methods that enable new capabilities and high performance, while increasing robustness, safety, reliability, and reducing operational costs. For all-electric geostationary satellites, MERL has leveraged multi-horizon model predictive control and inner-outer feedback loop design to concurrently handle station keeping, attitude control, and momentum management. MERL's method achieves superior fuel efficiency, almost reaching the theoretical maximum level.

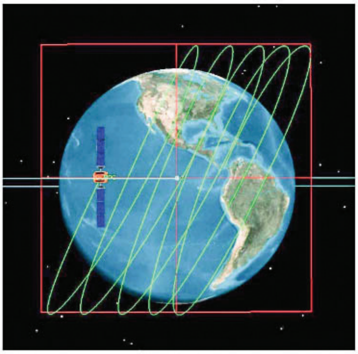

For spacecraft rendezvous on arbitrary orbits, MERL has developed control methods leveraging reachability, computational geometry, and stochastic processes that guarantee safety even in case of partial or total propulsion failures. MERL's technology enables safe rendezvous even in complex orbits such as near rectilinear halo orbits (NRHOs) around the Moon, which are fundamental for future space utilization.

Airflow Optimization and Control

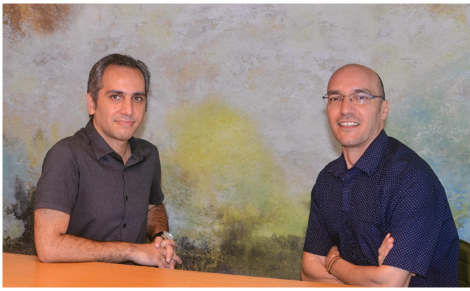

MERL does fundamental research on data-assimilation and inverse problems in the context of complex, high-dimensional turbulent thermo-fluidic systems. We develop theory, computational methods, and algorithms for PDE-constrained optimization using tools at the intersection of fluid dynamics, dynamical systems, and optimal control.

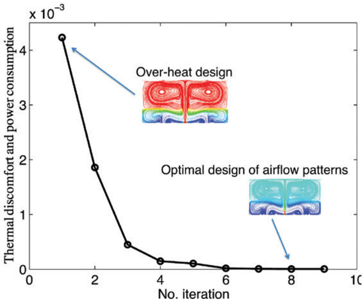

Fluid mechanics is historically a big data field and amenable to recent innovations in machine learning and dynamical systems. MERL research emphasizes both model-based and data-driven approaches for optimal control. Through years of development, our adjoint solver has become very stable, efficient, and includes various mathematical guarantees for convergence rate. Our research also spans nonlinear, robust, adaptive, and real-time estimation and control of turbulent flows with guaranteed performance using a combination of Lyapunov-based analysis and deep reinforcement learning. We exploit hybrid approaches that combine first-principles physics with data-driven methods to derive interpretable and physically consistent results.

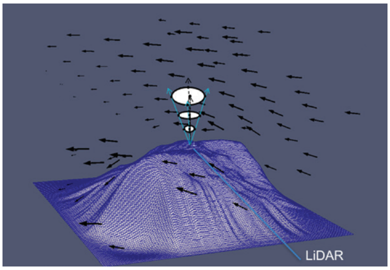

Applications of our research include advanced HVAC systems, net-zero buildings, PDE boundary control, shape optimization, optimal sensor/actuator placement, slosh suppression in fuel tanks, and LiDAR & remote sensing for wind reconstruction in the atmospheric boundary layer.

Autonomy & ADAS

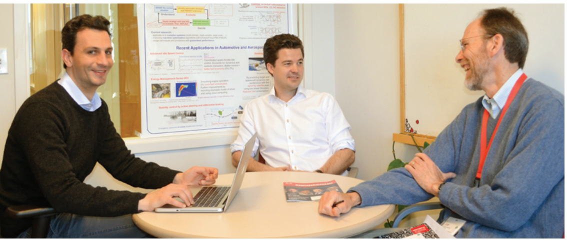

Autonomous systems execute complex tasks without human supervision, and are expected to be pervasive in the future, in vehicles, homes, and cities. At MERL we research methods for safe, efficient, and cost-effective autonomy on spacecraft, cars, drones, and robots, to ensure they make life better and coexist seamlessly with our future society.

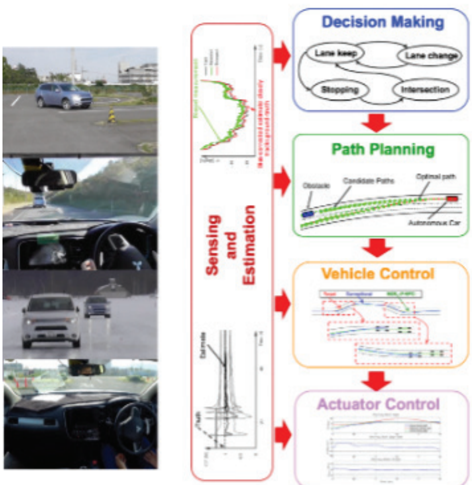

Leveraging MERL fundamental research in decision & control, dynamics, signal processing, computer vision, and machine learning, we have developed sensing and perception methods, localization methods, and guidance and control methods.

MERL planning and positioning methods based on particle filtering have been extensively demonstrated in cars in real life and extreme driving conditions. MERL compounding control methods based on model predictive control are used as the basis of Advanced Driver Assistance Systems (ADAS).

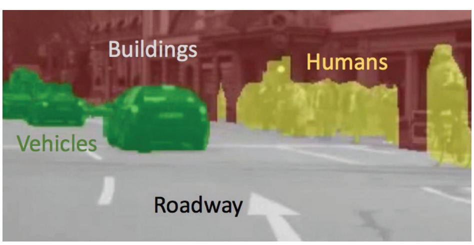

Signal Processing, machine learning, and AI research at MERL allows improved perception and understanding of the world surrounding an autonomous system. MERL work in spacecraft guidance and control supports the full potential of future generations of space systems.

Vehicle Motion Planning

Autonomous vehicles must complete complex journeys while following complex traffic laws in a rapidly changing and uncertain environment. Key enablers for autonomous driving are a real-time decision-maker that determines the next steps towards a destination, and a motion planner that generates trajectories for the steps.

MERL developed methods for decision making and motion planning for autonomous vehicles that ensure safety in changing traffic conditions, provide high performance achieving multiple driving goals, and are efficient enough to run on current automotive-grade processors.

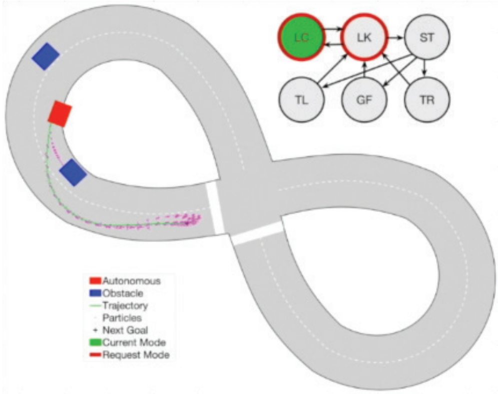

MERL's decision-maker uses reachability theory to determine what actions a vehicle can accomplish while satisfying traffic rules, vehicle dynamics, and remaining safe. These "safe feasible actions" are provided to MERL's particle-filter motion planner, which quickly determines safe, effective, drivable trajectories in traffic by a combination of statistical estimation and control techniques.

MERL developed automated parking software that can navigate complex parking lots and execute parking maneuvers with centimeter-level precision in the presence of other vehicles.

In collaboration with other Mitsubishi Electric R&D Labs all of these methods have been tested in real driving conditions, including cities and snow-covered roads.

Real-Time Optimization for ADAS

Advanced Driver Assistance Systems (ADAS) support drivers in avoiding accidents and dangerous conditions. Advanced control methods, such as Model Predictive Control (MPC), increase ADAS' effectiveness via real-time optimization. This requires solving constrained optimization problems up to hundreds of times per second in an on-board processor with limited computational resources. For this to be practical, fast solvers are essential.

By leveraging long expertise in MPC and numerical optimization, MERL has developed several methods to increase real-time optimization speed by orders of magnitude, while maintaining reliability and accuracy. Using online linearization and warm starting, we can determine needed ADAS actions by solving a single convex quadratic optimization problem at each control cycle. Over the last 10 years, MERL developed several types of methods, from gradient-based, such as PQP and ADMM, to second order methods, such as active-set and interior-point methods. Recently, MERL developed the PRESAS active-set method based on a preconditioned iterative solver, which provides high-accuracy solutions and is 10-100 times faster than alternative real-time optimization methods. PRESAS enables MPC for ADAS and has been deployed in several hundred thousand vehicles.

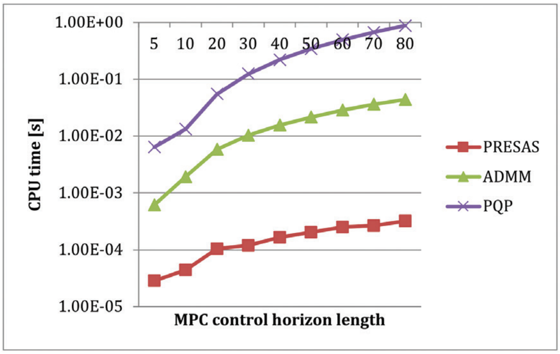

Vehicle Control and Estimation

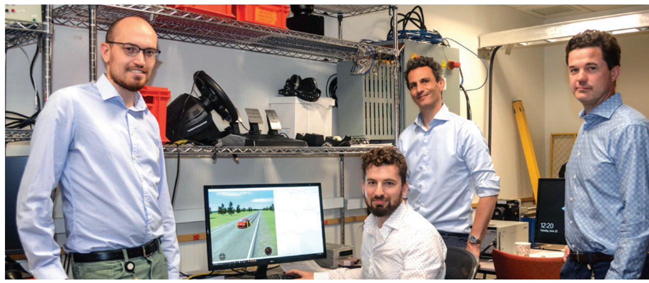

Autonomous Driving Systems (ADS) and Advanced Driver-Assistance Systems (ADAS) must operate safely under a wide range of conditions of vehicle, road, weather, and traffic. ADS and ADAS leverage estimation and control algorithms to cancel or mitigate the effect of varying conditions and uncertainty about those conditions.

MERL has developed algorithms for estimating the environment and vehicle conditions, and for controlling a vehicle in varying conditions. These include parametric and nonparametric methods that estimate uncertainties such as tire-road friction, vehicle center of gravity location, and obstacle properties, by using data from automotive-grade sensors. The uncertainty estimates are used in MERL's high-performance control algorithms, providing performance guarantees even in the presence of uncertainty and using computational resources typical of automotive-grade processors.

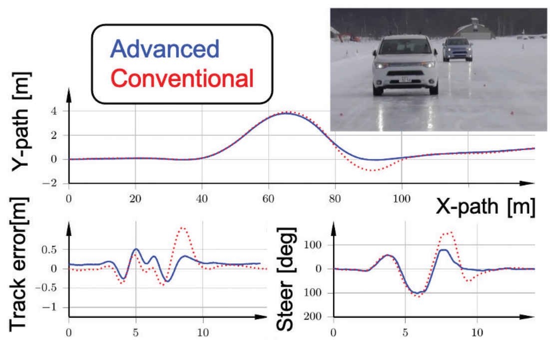

Using robust and stochastic model predictive control, MERL algorithms provide vehicle stability control that ensures proper vehicle behavior, even in the presence of varying and uncertain environmental conditions. MERL has applied these methods to Electric Vehicles (EVs), optimizing at the same time vehicle responsiveness and passenger comfort, through the coordination of multiple electric motors.

Vehicular Radar

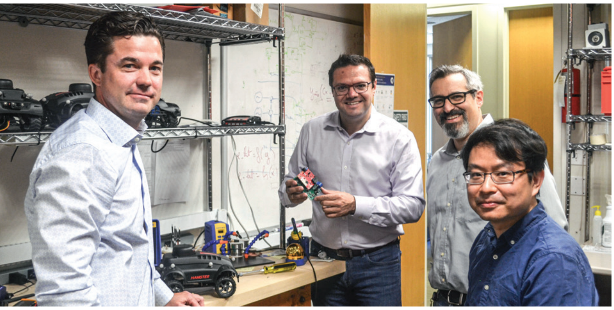

Automotive radar plays an important role in autonomous driving by providing reliable perception in all-weather conditions at an affordable cost. Modern radar sensors can resolve multiple detection points per object. Compared with conventional point object tracking, extended object tracking (EOT) with automotive radar can determine a target object's length and width in addition to its position, velocity, acceleration, heading, and turn rate.

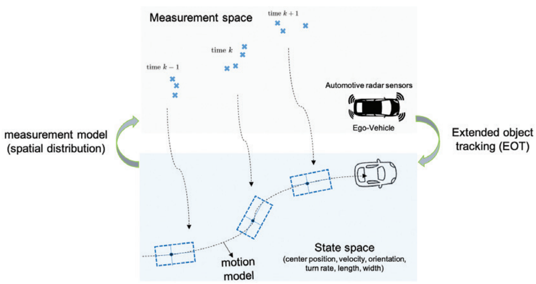

MERL's research focuses on new kinds of spatial-distribution modeling of real-world automotive radar measurements using online models & models learned offline along with online model adaptation that fine-tunes spatial models learned offline into real-time spatial models useful for onboard automotive radar sensors. We also work on state prediction & updating of an object's size, position, velocity, etc. that takes into account these new spatial distribution models.

This work enables radars that can better distinguish multiple objects (car, bicycle, pedestrian, etc.) and track them as they move. Our more recent work seeks further improvements by combining measurements from several radars, which we expect will be deployed on vehicles in the future.

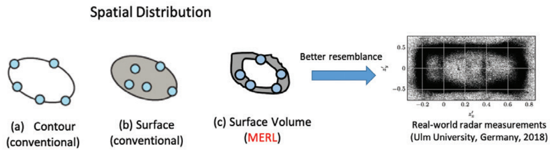

Speech and Audio

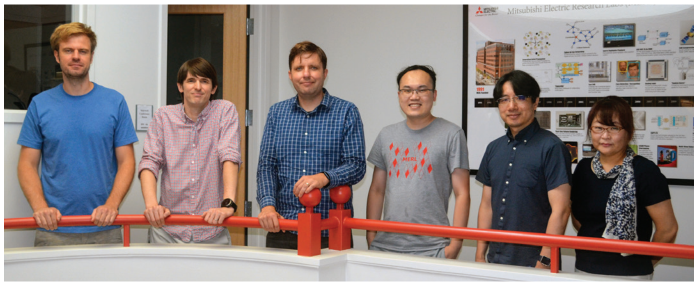

Research at MERL addresses a range of challenging speech and audio machine-perception problems involving acoustic signals, human language, and everything in between. On the signal acquisition end, we focus on source separation in both single- and multi-channel settings, sound event detection, and sound analysis. For automatic speech recognition (ASR), we have introduced leading methods for end-to-end ASR, including novel deep-learning methods for acoustic and language modeling. On the language end, our efforts focus on natural-language understanding and multi-modal techniques to realize scene-aware interaction.

A particular difficulty in this field is bridging the gap between two differently structured domains: acoustic signals and language. Our approach has been to develop novel machine-learning techniques that go beyond classical pattern-recognition frameworks. In recent years, this has enabled us to take on a series of groundbreaking research projects, achieving world-leading performance in a wide range of tasks and consistently placing in the top tier in international competitions.

SpokenQuery

Long before Siri and Cortana became mainstream user interfaces, MERL pioneered a novel approach to speech-based information retrieval called SpokenQuery. The system allows natural, free-form spoken requests for information and retrieves results with high accuracy.

Speech-based data retrieval is composed of two basic steps: recognizing what the user has said and doing database retrieval based on the user's request. However, if these steps are strictly separated, the retrieval is very unlikely to be correct when there is any error in recognizing what the user has said.

SpokenQuery extracts from the speech recognition engine not just its best guess at what the user has said, but all reasonably likely guesses; it then uses all the salient words in these guesses as keywords to search a database, avoiding both recognition and grammatical problems. SpokenQuery enhances the usability of commercial speech-recognition engines and provides a more flexible input style. It has been commercialized for automotive systems to find music and points of interest from databases.

Speech Enhancement and Separation

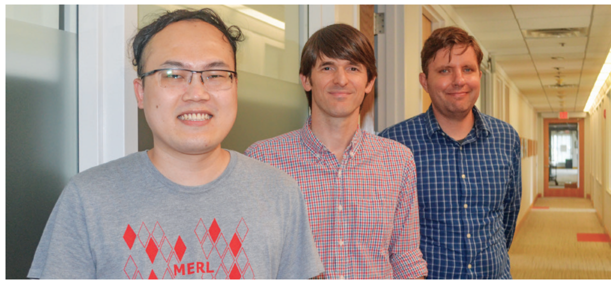

The human auditory system enables us to do an extraordinarily good job conversing above the chatter of a lively cocktail party. Endowing a computer with this ability has been the holy grail of speech processing for over 50 years.

For 20 years, MERL has pioneered advances in source separation. Combining originality with solid theoretical foundations, MERL's methods have dramatically influenced the field and solved problems that were previously out of reach. Early highlights include seminal work on non-negative matrix factorization of audio, audio-visual enhancement, bandwidth expansion, and speaker-dependent speech separation.

Recently, by harnessing the power of cutting-edge deep learning, our algorithms have been able to overcome far more challenging interference than ever before.

In particular, MERL invented the first algorithm capable of accurately separating the overlapped speech of two unknown speakers, which led to speaker-independent speech separation being a new major focus in the field. MERL has since introduced methods for separating general sounds, learning to separate sounds without ground truth signals. In addition, we have re-framed the sound separation problem in a hierarchical paradigm, working towards achieving ever more general and high-quality sound separation.

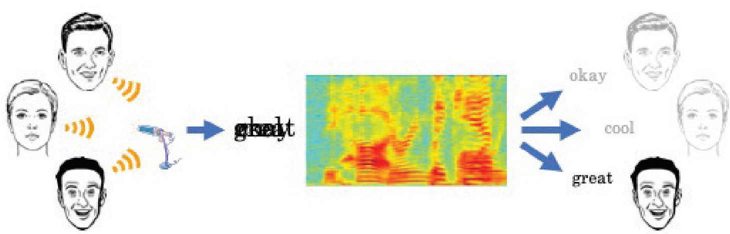

End-to-End ASR

With the advent of deep learning, Automatic Speech Recognition (ASR) has been undergoing a revolution in the last decade. MERL has made multiple pioneering contributions to the emergence of a new crop of so-called end-to-end ASR systems that replace the whole recognition pipeline with a single neural network, directly mapping the input acoustic waveform to the output word sequence. This direct mapping allows for a simpler, easier-to-deploy design needing less expert knowledge with the joint optimization of the whole system leading to better performance.

MERL's contributions have helped improve recognition accuracy thanks to better training methods and architectures for long-context systems that exploit the past acoustic and linguistic context in long recordings. They have enabled online processing in attention-based systems that were previously offline, leading to state-of-the-art performance. And they have unlocked new capabilities for ASR, with the development of the first "seamless ASR" system that can recognize speech in multiple languages spoken simultaneously by multiple speakers. MERL is now pushing ahead towards the goal of a machine capable of total transcription of an acoustic scene, from the content of the speech being uttered to a description of the surrounding sound events.

Scene-Aware Interaction

Recent advances in object recognition, video description, natural language generation, and spoken dialog technologies using deep neural networks have opened the door for the development of algorithms enabling machines to better understand their surroundings and interact with humans more naturally and intuitively. MERL has been at the forefront of this new area of research, establishing the field of audio-visual scene-aware dialog and developing fundamental technologies to realize scene-aware interaction.

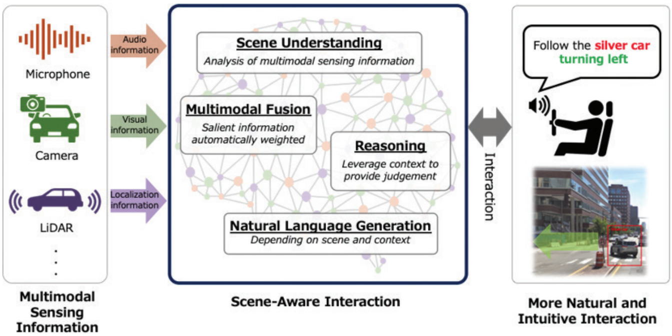

As an example application, MERL developed a car navigation system that can guide or warn a driver based on multimodal sensing of the surrounding environment, featuring multimodal fusion technology. This system was able to realize a powerful end-to-end scene-aware interaction system that demonstrated highly intuitive interaction with users in diverse situations. This concept and the related technologies are widely applicable to other domains, including robotics, surveillance, and home systems.

Computer Vision

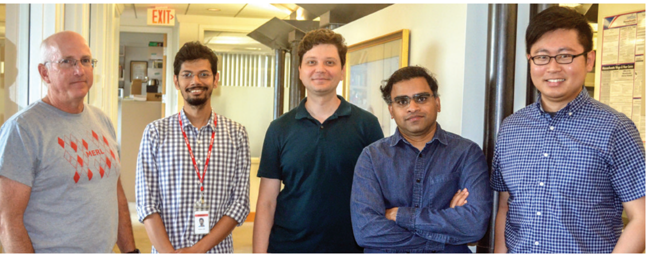

Over its history, MERL has performed research in almost every area of computer vision. As a precursor to the Wii and Kinect-style gaming devices, we developed human-action analysis on low-resolution video data. We also did early, precedent-setting work that combined camera and projector systems and pioneered new approaches that combined cameras and computation, which ultimately became known as the field of computational photography.

Work with stereo cameras led to research on 3D reconstruction and scene modeling, which grew to include providing object position and pose for robotic grasping, and the creation of 3D models of the world for the purpose of navigation and planning.

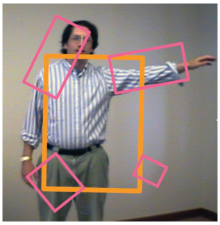

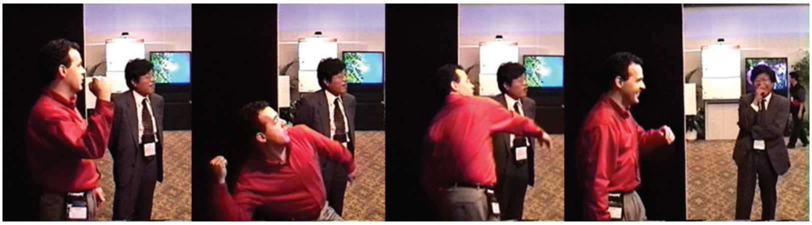

In recent years, we have developed new deep-learning algorithms for super-resolution, segmentation, noise reduction, and many variants of detection and classification. We have also combined visual data with other modalities such as acoustic data to perform tasks such as audio source separation and future prediction. Our work on graph neural network applications enables the segmentation and tracking of objects in point cloud data coming from laser radar devices for automotive applications.

Generalized Belief Propagation

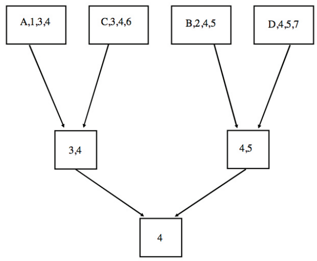

Belief Propagation is a message-passing algorithm for performing inference on graphical models. Given probabilistic information about some of the nodes in the graph and knowledge (embodied by the graph) about relationships between the nodes, it computes the most likely values of all the nodes. This is useful in a wide range of applications from Artificial Intelligence reasoning to error-correcting codes.

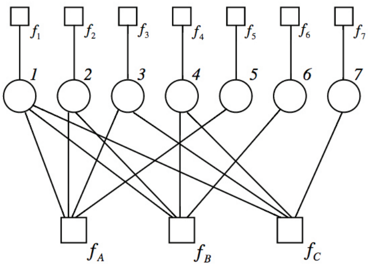

The basic Belief Propagation algorithm works well most of the time; however, while it often produces the correct result when applied to a cyclic graph, it is not guaranteed to do so. This is unfortunate since most graphs encountered in real-world situations are cyclic.

In 2000-2005 MERL researchers pioneered the development of Generalized Belief Propagation algorithms that are guaranteed to work on every graph. This work triggered the start of a new sub-field of research in the world that continues to this day.

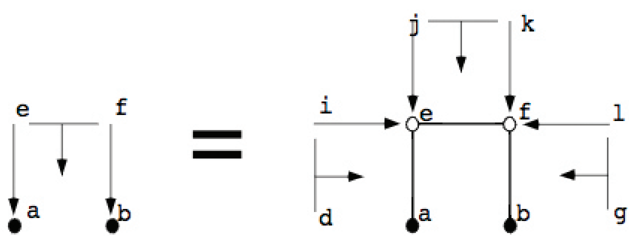

Object Detection

Viola and Jones began their seminal work on a new framework for detecting objects in images before they came to MERL. However, at MERL they extended their work in major ways and applied it to a range of tasks.

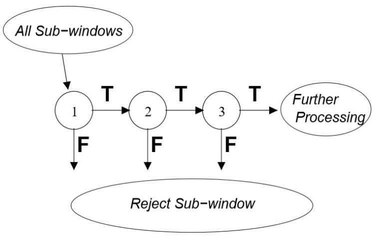

The Viola-Jones approach includes the concepts of Haar-like features, integral images, AdaBoost learning, and classifier cascades. Together these ideas create a framework for building object detectors in images that are simple, powerful, and extremely fast. This approach is now a standard technique in computer vision.

The first application of the Viola-Jones ideas was to face detection, where it achieved state-of-the-art accuracy with speeds an order of magnitude faster than previous approaches. Many devices and applications that use face detection today (such as smart phone cameras and photo organization software) are built on the approach Viola-Jones perfected at MERL.

Further extensions combining motion features with image features succeeded in detecting pedestrians both more accurately and faster than prior approaches.

Three R&D 100 Awards

MERL-developed computer vision algorithms are at the heart of three Mitsubishi Electric products that won R&D 100 awards in recent years.

In 2006, Mitsubishi Electric introduced Heli-Tele—the world's first disaster-response aerial camera system. At its core was an algorithm from MERL that registers live aerial video images with a conventional map so that corresponding locations could be identified. With this system, emergency managers can give instructions to the workers in the field using the correct local street addresses.

In 2014, Mitsubishi Electric introduced the MELFA-3D robot vision system, which enables a robot to identify and pick up parts from a bin. Unlike previously commercial systems, MELFA-3D works well even with a jumble of complex parts in a bin and even if the parts are shiny, both conditions being beyond the capabilities of prior systems.

In 2018, Mitsubishi Electric introduced a river monitoring system that automatically determines the water level. It uses a fast algorithm developed by MERL that enables efficient real-time monitoring to give early warning of potential flooding.

3D Reconstruction

Having accurate 3D models is important in many situations, such as planning, navigating, visualizing, and 3D printing. However, many real-world objects do not have 3D models, or no longer match preexisting models. Therefore, it is important to be able to rapidly construct accurate models.

MERL researchers developed an easy-to-use 3D reconstruction system consisting of a low-cost 3D sensor and a tablet. Our system reconstructs a 3D model in real time using the tablet for computation while an operator scans a scene or object. Large scenes can be modeled as the tablet is moved around. Planes, which are the dominant structure of indoor (and urban outdoor) scenes, are used to register multiple frames obtained from the 3D sensor. The use of planes enables more accurate registration than conventional point-based systems, and generates more compact 3D models.

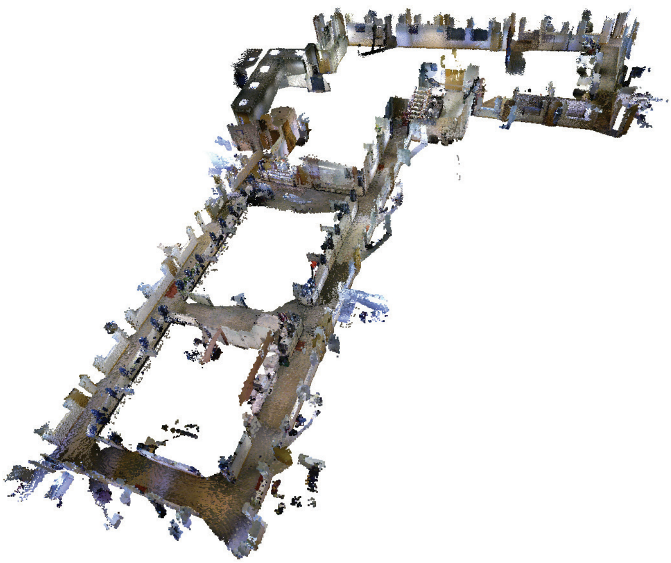

Video Anomaly Detection

Automatically noticing when something unusual is going on in a scene is an important ability that has many useful applications. This is the goal of video anomaly detection.

The key idea is to build models of normal activity given video of a particular scene that does not contain anomalous activity. Once a normal model is learned for a scene, it can be used to estimate how unusual each part of new video from the same scene is. Any part of test video that has large distance to the model of normal video will be detected as anomalous. An important property of the model is that it can easily be extended as new normal video becomes available.

MERL's video anomaly detection algorithms have been applied to detecting faults in the machinery of a factory assembly line as well as unusual activity in outdoor scenes with people and cars.

Point Cloud Processing

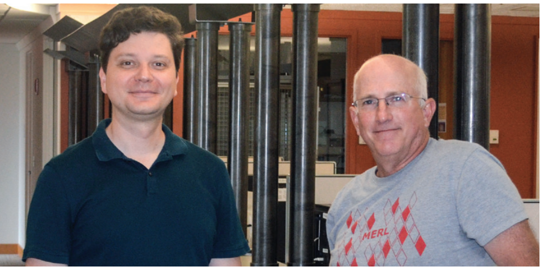

Point clouds are commonly generated by laser radars or stereo sensors in applications such as automotive and robotics. They contain valuable 3D information about objects in the scene for identification (classification), tracking, and motion prediction. MERL has worked with point clouds for many years for the purpose of 3D reconstruction of a static scene such as an elevator motor room during equipment replacement, or of a robotic workspace for scene understanding, programming, and collision avoidance.

We have developed graph signal processing algorithms that generalize Fourier transforms to graphs constructed from point clouds to enable the points to be high-pass filtered. This filtering predominantly removes points lying on flat planar areas of the data while preserving those on edges and corners. The resulting filtered point cloud is more compact and provides enhanced interpretability of the data without losing important information.

More recently we have used 4D (space and time) graph neural networks to perform spatio-temporal convolution on point clouds that enables the classification of the points into a number of object classes such as car, truck, or pedestrian, as well as the prediction of the future motion of these objects. This is essential technology for Advanced Driver Assistance Systems (ADAS) as well as for autonomous vehicles.

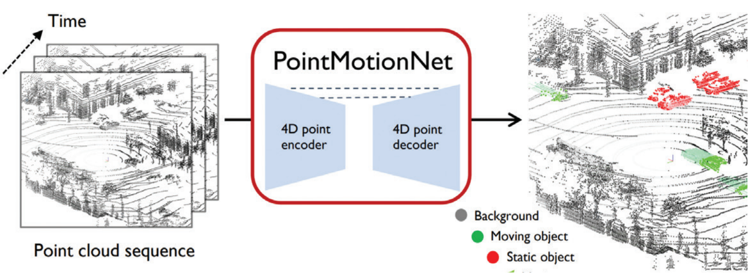

Robot Manipulation

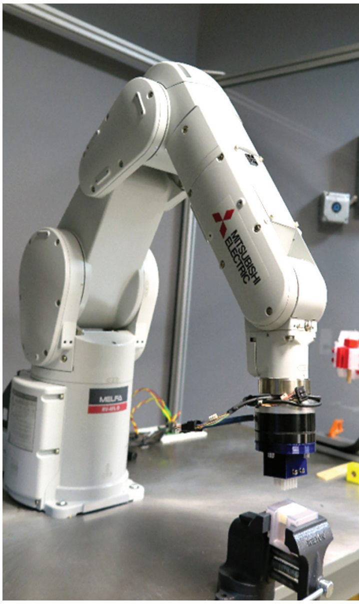

Although the cost of robot hardware has been decreasing steadily, the need for long, laborious, and costly robot programming has made robot applications economically feasible only for high-volume, low-mix industries. MERL's research on robotic manipulation aims to decrease the cost of robot deployment and facilitate their use in low-volume, high-mix industries and non-manufacturing environments such as households.

For many years, MERL has worked on grasping technology for picking parts lying in bins in random position and orientation, based on visual input. Now we are researching the use of force and tactile sensors to support dexterous, contact-rich manipulation tasks such as robotic assembly without jigs and fixtures.

A key technology reducing the need for costly robot programming is robot learning. We have pursued machine learning (ML) methods to solve difficult tasks, such as solving marble mazes via self-experimentation. We have also experimented with learning from demonstration, where the robot learns how to insert parts during assembly from a single human demonstration. ML, simulation, and self-exploration are used by the robot to learn how to respond to force feedback in a compliant manner during operations featuring tight tolerances and imprecise position estimation.

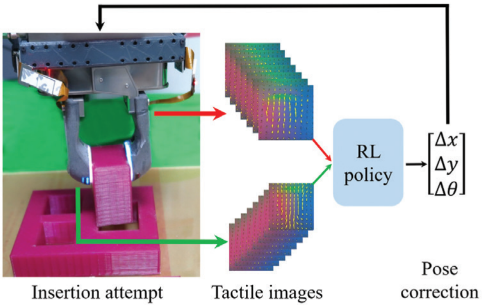

Robot Bin Picking

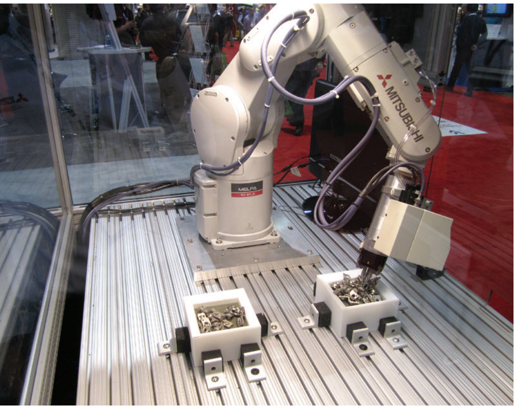

Bin picking is the task of picking up a particular part from a bin full of randomly oriented parts. People do this largely by feel. Robots rely on seeing the exact position and orientation of a part and then grasping it with their very limited dexterity and touch sensing. This “seeing” is a traditional challenging problem in computer vision, and is very hard, particularly when the parts are shiny, partially obscure each other, and the task must be done quickly.

Over the past 15 years, MERL has solved major parts of this problem. Our multi-flash camera captures depth information. We designed fast algorithms for object detection and pose estimation using 2D and 3D features, which are robust enough to handle a wide range of real industrial situations. One of our algorithms was commercialized as a part of Mitsubishi Electric’s MELFA-3D robot Vision system in 2014.

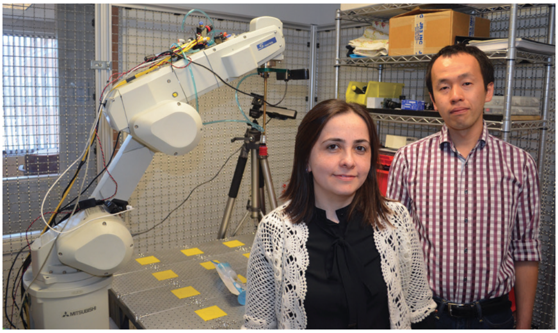

Robot Ball Maze Game

Model-free Reinforcement Learning (RL) is a popular way for a computer to learn to perform a complex task. However, RL typically requires a large number of trials which is costly for robots. MERL's research has focused on reducing the number of trials needed for robot RL by exploiting Sim2Real, Imitation Learning, and Model-Based RL concepts.

Sim2Real involves learning in a simulation domain, followed by transfer to, and refinement in, a real-world setting. Imitation Learning uses expert demonstrations to pre-train control policies, which can then be further trained with RL.

As a difficult experimental task, we chose solving a marble maze game: multiple marbles have to be gathered in the center of a maze by steering them through the maze by tilting the puzzle. A digital simulation of the marble maze game was developed using computer graphics and a physics engine. Solving the marble maze game was then learned in simulation using RL. The learned control policy was then transferred to the controller for a robot arm holding the marble maze. MERL researchers demonstrated that the transferred control policy could be refined for the real-world setting in 60x less time compared to learning the task directly in the real-world setting. Using imitation learning caused a further 10~20x speed up.

Tactile Feedback Control

Tactile sensing is essential to enable factory automation to escape the limitations of object manipulation without any awareness of an object once a robot arm has grasped it. Based on tactile information, MERL has developed algorithms that solve a range of robotic manipulation problems. These include: object identification based purely on tactile feedback, autonomous peg insertion & item packing, grasped object slip detection, grasp pose prediction, and minimal pressure grasping. These algorithms serve as components for high-complexity manipulation tasks such as food handling, warehouse order fulfillment, and autonomous assembly.

Researchers at MERL have also developed a novel high-resolution tactile sensor based on surface deformation. An elastic 3D-printed membrane is instrumented with markers which are tracked by an endoscopic vision system to determine qualitative touch conditions such as object contact, relative torque, and relative shear.

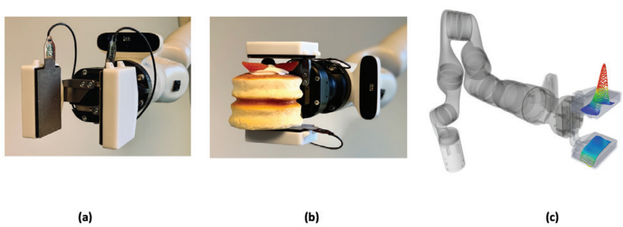

Illustration of tactile sensing to determine critical grasping. (a) Touch sensing array. (b) Critical grasping of soft object using gripper equipped with touch sensing arrays. (c) High resolution tactile response map is calibrated for true pressure measurements (kPa)

MERL Headquarters

The heart of MERL is its research staff, but much more is needed to run an effective operation. MERL's headquarters' staff is equally important, making the impact of the researchers' efforts possible.

"MERL's philosophy of collaboration is impressively open and broad, which I believe is a significant contributor to its success."

— Stark C. Draper, Ph.D.
Professor, ECE Department
University of Toronto
At MERL 2006-2007

"My three years at MERL were some of the most productive and happy times in my career. I joined a fantastic team consisting of leading experts in computational sensing and we tackled some of the most stimulating problems in our area."

— Ulugbek Kamilov, Ph.D.
Assistant Professor
Washington University in St. Louis
At MERL 2015-2018

MERL Headquarters

Management Team

MERL is led by a closely collaborating management team of six. The President and Executive Vice President have more than 30 years of U.S. research management experience and an equal amount of experience within Mitsubishi Electric's CR&D organization, respectively. The ability to meld these two points of view in a mutually reinforcing way has been an essential basis for the success of MERL over the years. The President and Executive Vice President work in close collaboration with two research Directors, a Deputy Director, and a Director of HR & Administration to guide MERL's operations. MERL uses a relatively flat organizational structure to enhance collaboration opportunities among MERL researchers. Research directions and projects are generated bottom up, based on ideas from researchers.

Patents

Filing patents allows MERL to capture our results as competitive advantages for Mitsubishi Electric, while publishing widely and participating in the global research community. In order to file high-quality patents quickly, MERL has an in-house patent department with a Patent Attorney, two Patent Agents, and a Patent Paralegal. They work with researchers to write patent applications and support the granting of these patents in countries around the world. More than 1,500 patents have been granted to MERL.

Liaisons

For our research results to impact Mitsubishi Electric's products, MERL collaborates with the 2,000 researchers in Mitsubishi Electric's CR&D labs in Japan. This collaboration is facilitated by MERL's EVP and a liaison staff of three, all four of whom are on loan from these labs. They play key roles connecting MERL researchers to researchers in Japan and aiding communication.

HR, Administration, Accounting

As a stand-alone company, MERL has the ability (and the need) to provide its own human resources, administration, and accounting. MERL gains significant advantages from being part of the 5,000-employee-strong group of Mitsubishi Electric subsidiaries in the U.S., but is not burdened by the kinds of red tape and delays that would be inevitable if we were not, in fact, a separate company.

Central Services

MERL's work is computer-intensive. A team of four supports MERL's central infrastructure of servers and networking, as well as the individual machines our researchers use every day. All told, MERL has hundreds of computers, and hardware and software taken together, no two are the same.

University Relations and Internship

Actively participating in the world research community and publishing our work have always been key features of MERL's culture. We maintain close relations with key universities and collaborate with the leading academic research groups in our areas of interest. MERL often hosts seminars presented by external researchers and collaborators. Being located at the heart of Kendall Square allows our members unfettered access to the many talks and seminars taking place at MIT, Harvard, and other nearby universities.

A central aspect of our overall academic relationship is our graduate student internship program. The internship program provides interns with experience that helps them enhance and accelerate their professional careers, while contributing to initiatives at MERL and helping us to identify good researchers to hire. Although we host students throughout the year, the main influx is during the summer when our research staff doubles as students from all over the world arrive at our Cambridge office.

Benefits of a MERL Internship

- Experience: At MERL, you work closely with top researchers and participate in a lab-wide R&D culture with a unique mix of curiosity-driven research and market-oriented prototyping.
- Publication: MERL is an open research lab with a strong tradition of publication in high-impact, peer-reviewed venues. Internships typically aim at producing publication-worthy results and many interns are co-authors on papers each year.
- Compensation: MERL offers competitive salaries based on relevant education, skills, and work experience.
- Perks: MERL provides relocation assistance including travel costs, subsidies for commuting costs, and entertainment events where interns can get to know Boston and each other.
- Networking: Interns are encouraged to network with MERL's research staff, fellow interns, and faculty at local universities. Weekly socials and seminars provide many opportunities.
- Opportunity: Many MERL interns have gone on to distinguished careers at MERL. MERL research hosts have often provided letters of reference supporting their ex-interns' candidacies for jobs, fellowships, and tenure.

"The work environment is very intellectually stimulating and I got a chance to speak with experts from various different fields."

— Sambarta Dasgupta
Intern from Iowa State University

"My internship at MERL was both challenging and fun. The people are great and very open to sharing their passion and knowledge about their research."

— Walter Weiss
Intern from Queen's University

"I had a great experience at MERL. My projects ranged from fundamental academic problems to high impact applications. The research diversity and expertise at MERL were impressive and I benefited a lot from my collaborations. Finally, Boston is definitely the place to be in the summer!"

— Sercan Ö. Arik
Intern from Stanford University

Selected MERL Publications

Below are MERL publications that have garnered particularly significant recognition over time. Each has been cited more than 500 times according to Google Scholar and 21 have been cited more than 1,000 times. Using citation counts as part of the selection criteria skews the list toward our earlier work, but highlights the significant impact of that work.

"Working at MERL was not only enjoyable, it taught me many things I have been benefiting from over the years."

— Andreas F. Molisch, Ph.D.
Professor, Department of Electrical and Computer Engineering
University of Southern California
At MERL 2002-2008

"The culture at MERL is great, respecting the effort of everybody and advocating a balanced and fruitful life. I enjoyed the freedom in fundamental research and collaboration inside and outside MERL."

— Ziming Zhang, Ph.D.
Assistant Professor, Electrical and Computer Engineering
Worcester Polytechnic Institute
At MERL 2016-2019

Selected MERL Publications

Ahmed, E., Jones, M., and Marks, T.K., 2015. An improved deep learning architecture for person re-identification. *Proceedings of the IEEE conference on computer vision and pattern recognition.* pp. 3908-3916. **(1,218 citations)**

Avidan, S. and Shamir, A., 2007, August. Seam carving for content-aware image resizing. *ACM Transactions on graphics (TOG)*. 26(3), p. 10. **(2,145 citations)**

Avidan, S., 2007. Ensemble tracking. *IEEE transactions on pattern analysis and machine intelligence.* 29(2), pp. 261-271. **(1,943 citations)**

Boufounos, P.T. and Baraniuk, R.G., 2008, March. 1-bit compressive sensing. *In 2008 42nd Annual Conference on Information Sciences and Systems, IEEE.* pp. 16-21. **(708 citations)**

Brand, M., 2002. Charting a manifold. *Advances in neural information processing systems.* pp. 961-968. **(586 citations)**

Brand, M., 2002. Incremental singular value decomposition of uncertain data with missing values. *Computer Vision—ECCV 2002*. Springer Berlin Heidelberg, pp. 707-720. **(606 citations)**

Davenport, M.A., Boufounos, P.T., Wakin, M.B., and Baraniuk, R.G., 2010. Signal processing with compressive measurements. *IEEE Journal of Selected Topics in Signal Processing,* 4(2), pp. 445-460. **(699 citations; 2015 IEEE Signal Processing Society Best Paper Award)**

Dietz, P. and Leigh, D., 2001, November. DiamondTouch: a multi-user touch technology. *Proceedings of the 14th annual ACM symposium on User interface software and technology, ACM.* pp. 219-226. **(1,854 citations)**

Efros, A.A. and Freeman, W.T., 2001, August. Image quilting for texture synthesis and transfer. *Proceedings of the 28th annual conference on Computer graphics and interactive techniques, ACM.* pp. 341-346. **(3,056 citations)**

Freeman, W.T. and Roth, M., 1995, June. Orientation histograms for hand gesture recognition. *In International workshop on automatic face and gesture recognition*, Vol. 12, pp. 296-301. **(997 citations)**

Freeman, W.T. and Weissman, C.D., 1997. Hand gesture machine control system. *U.S. Patent 5,594,469.* **(1,132 citations)**

Freeman, W.T., Pasztor, E.C., and Carmichael, O.T., 2000. Learning low-level vision. *International journal of computer vision*, 40(1), pp. 25-47. **(2,117 citations)**

Freeman, W.T., Jones, T.R., and Pasztor, E.C., 2002. Example-based super-resolution. *Computer Graphics and Applications, IEEE, 22(2)*, pp. 56-65. **(3,144 citations)**

Frisken, S.F., Perry, R.N., Rockwood, A.P., and Jones, T.R., 2000, July. Adaptively sampled distance fields: a general representation of shape for computer graphics. *Proceedings of the 27th annual conference on Computer graphics and interactive techniques,* pp. 249-254. ACM Press/Addison-Wesley Publishing Co. **(803 citations)**

Gezici, S., Tian, Z., Giannakis, G.B., Kobayashi, H., Molisch, A.F., Poor, H.V., and Sahinoglu, Z., 2005. Localization via ultra-wideband radios: a look at positioning aspects for future sensor networks. *Signal Processing Magazine, IEEE*, 22(4), pp. 70-84. **(2,456 citations)**

Hershey, J., Chen, Z., Le Roux, J., and Watanabe, S., 2016, March. Deep clustering: Discriminative embeddings for segmentation and separation. *IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. **(780 citations)**

Liu, M.Y. and Tuzel, O., 2016, December. Coupled generative adversarial networks. *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pp. 469-477. **(1,179 citations)**

Marks, J., Andalman, B., Beardsley, P.A., Freeman, W., Gibson, S., Hodgins, J., Kang, T., Mirtich, B., Pfister, H., Ruml, W., and Ryall, K., 1997, August. Design galleries: A general approach to setting parameters for computer graphics and animation. *Proceedings of the 24th annual conference on Computer graphics and interactive techniques,* pp. 389-400. **(750 citations)**

Matusik, W. and Pfister, H., 2004, August. 3D TV: a scalable system for real-time acquisition, transmission, and autostereoscopic display of dynamic scenes. *ACM Transactions on Graphics (TOG)* (Vol. 23, No. 3, pp. 814-824). ACM. **(702 citations)**

Molisch, A.F., Foerster, J.R., and Pendergrass, M., 2003. Channel models for ultrawideband personal area networks. *Wireless Communications, IEEE, 10(6)*, pp. 14-21. **(1,073 citations)**

Molisch, A.F., 2007. *Wireless communications*. John Wiley & Sons. **(3,531 citations)**

Pfister, H., Hardenbergh, J., Knittel, J., Lauer, H., and Seiler, L., 1999, July. The VolumePro real-time ray-casting system. *Proceedings of the 26th annual conference on Computer graphics and interactive techniques,* pp. 251-260. ACM Press/Addison-Wesley Publishing Co. **(598 citations)**

Pfister, H., Zwicker, M., Van Baar, J., and Gross, M., 2000, July. Surfels: Surface elements as rendering primitives. *Proceedings of the 27th annual conference on Computer graphics and interactive techniques,* pp. 335-342. **(1,250 citations)**

Porikli, F., 2005, June. Integral histogram: A fast way to extract histograms in cartesian spaces. *IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'05)*. Vol. 1, pp. 829-836. **(1,052 citations)**

Porikli, F., Tuzel, O., and Meer, P., 2006, June. Covariance tracking using model update based on lie algebra. *2006 IEEE Computer Society Conference on Computer Vision and Pattern Recognition (CVPR'06) Vol. 1*, pp. 728-735. **(813 citations)**

Raskar, R., Van Baar, J., Beardsley, P., Willwacher, T., Rao, S., and Forlines, C., 2006, July. iLamps: geometrically aware and self-configuring projectors. *ACM SIGGRAPH 2006 Courses,* p. 7. ACM. **(534 citations)**

Raskar, R., Agrawal, A., and Tumblin, J., 2006. Coded exposure photography: motion deblurring using fluttered shutter. *ACM SIGGRAPH 2006 Papers*. pp. 795-804. **(709 citations)**

Rubinstein, M., Shamir, A., and Avidan, S., 2008. Improved seam carving for video retargeting. *ACM transactions on graphics (TOG)*, 27(3), pp. 1-9. **(939 citations)**

Smaragdis, P. and Brown, J.C., 2003, October. Non-negative matrix factorization for polyphonic music transcription. *Applications of Signal Processing to Audio and Acoustics, 2003 IEEE Workshop on*. pp. 177-180. IEEE. **(1,142 citations)**

Tenenbaum, J.B. and Freeman, W.T., 2000. Separating style and content with bilinear models. *Neural computation*, 12(6), pp. 1247-1283. **(908 citations)**

Tuzel, O., Porikli, F., and Meer, P., 2006, May. Region Covariance: A Fast Descriptor for Detection and Classification. *European Conference on Computer Vision (ECCV)*. **(1,612 citations)**

Tuzel, O., Porikli, F., and Meer, P., 2008. Pedestrian detection via classification on Riemannian manifolds. *IEEE transactions on pattern analysis and machine intelligence*, 30(10), pp. 1713-1727. **(1,069 citations)**

Vetro, A., Christopoulos, C., and Sun, H., 2003. Video transcoding architectures and techniques: an overview. *Signal Processing Magazine, IEEE*, 20(2), pp. 18-29. **(982 citations)**

Vetro, A., Wiegand, T., and Sullivan, G.J., 2011. Overview of the stereo and multi-view video coding extensions of the H. 264/MPEG-4 AVC standard. *Proceedings of the IEEE*, 99(4), pp. 626-642. **(692 citations)**

Veeraraghavan, A., Raskar, R., Agrawal, A., Mohan, A., and Tumblin, J., 2007. Dappled photography: Mask enhanced cameras for heterodyned light fields and coded aperture refocusing. *ACM Trans. Graph.*, 26(3), p. 69. **(764 citations)**

Viola, P. and Jones, M.J., 2004. Robust real-time face detection. *International Journal of Computer Vision*, 57(2), pp. 137-154. **(16,299 citations)**

Viola, P., Jones, M.J., and Snow, D., 2005. Detecting pedestrians using patterns of motion and appearance. *International Journal of Computer Vision*, 63(2), pp. 153-161. **(3,156 citations; early version won the 2004 Marr Prize: best paper at ICCV)**

Yedidia, J.S., Freeman, W.T., and Weiss, Y., 2000, December. Generalized belief propagation. *Neural Information Processing Systems (NIPS),* Vol. 13, pp. 689-695. **(1,338 citations)**

Yedidia, J.S., Freeman, W.T., and Weiss, Y., 2003. Understanding belief propagation and its generalizations. *Exploring artificial intelligence in the new millennium*, 8, pp. 236-239. **(1,868 citations)**

Yedidia, J.S., Freeman, W.T., and Weiss, Y., 2005. Constructing free-energy approximations and generalized belief propagation algorithms. *Information Theory, IEEE Transactions on,* 51(7), pp. 2282-2312. **(1,727 citations)**

Zwicker, M., Pfister, H., Van Baar, J., and Gross, M., 2001, August. Surface splatting. *Proceedings of the 28th annual conference on Computer graphics and interactive techniques.* pp. 371-378. **(828 citations)**

"My nine years at MERL were wonderful. I had the freedom to work on fundamental problems, to collaborate with researchers outside the lab, and to publish."

— William T. Freeman, Ph.D.
Professor, EECS Department
Massachusetts Institute of Technology
At MERL 1992-2001

"My research life at MERL was fantastic. During this period, I concentrated my full efforts on high-impact research with outstanding colleagues in a great research environment giving me numerous opportunities."

— Shinji Watanabe, Ph.D.
Associate Professor, Language Technologies Institute
Carnegie Mellon University
MERL tenure 2012-2017

MERL
201 Broadway
Cambridge, MA 02139-1955
(617) 621-7500
www.merl.com