Mitsubishi Electric Research Laboratories

25 YEARS

1991 — 2016

INNOVATION | IMAGINATION | INSPIRATION

MITSUBISHI ELECTRIC

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

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

Production:
Anthony Vetro & Richard C. Waters

Photos:
Fred Igo

Design:
Frances Scanlon, Scanlon Design

Introduction

Celebrating 25 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.

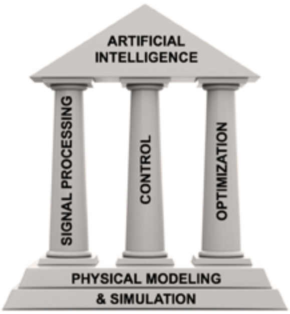

With 60 researchers, MERL is small enough to be flexible and agile, while gaining leverage from our large global parent Mitsubishi Electric. We turn our technical achievements into business successes by partnering with the tens of thousands of researchers and engineers in Mitsubishi Electric's operations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

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 too was 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

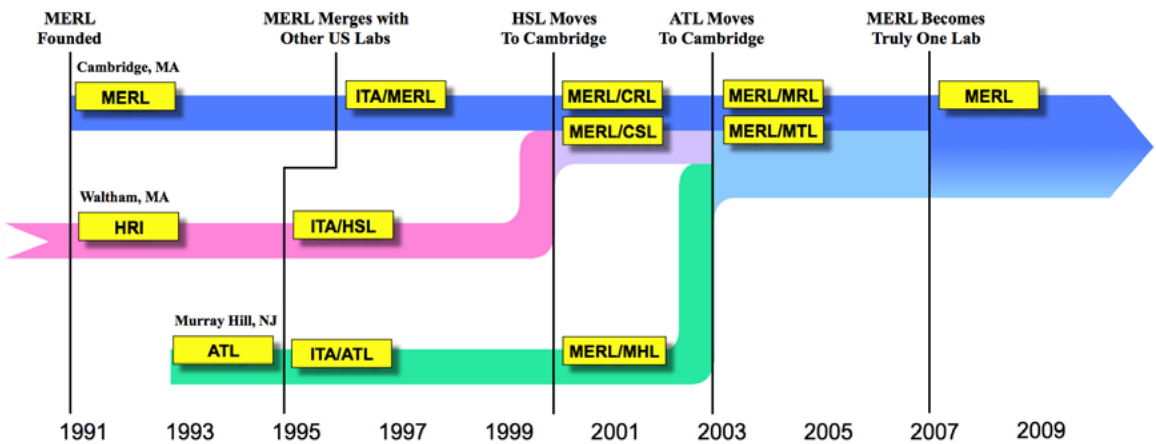

focused on long-range research and one on research & advanced development.

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 it has produced. The following pages summarize the different areas of research MERL is engaged in, interspersed with short descriptions of major results in those areas.

Research Areas and Results

The following pages summarize areas of research at MERL, interspersed with short descriptions of significant research results in those areas.

“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

Optical Communications

In the early 2000s, MERL's optical communications work focused on forward error correction (FEC) using low-density parity-check (LDPC) codes, resulting in several algorithms that achieved world's best coding performance in optical transmission equipment with extremely low bit-error rates.

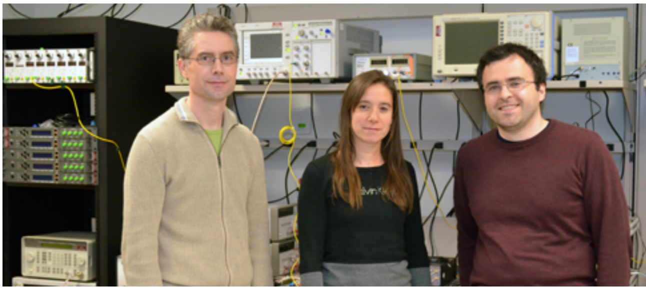

Since 2009, we have expanded our research into technologies for digital coherent optical devices and optical systems. On the devices side, we have developed novel designs with advanced multi-parameter optimization methods and have verified experimental performance using test facilities at MERL.

On the systems side, we have developed new signal processing algorithms to compensate for linear and non-linear impairments introduced by optical fiber and by non-ideal components. We have also developed novel high-dimensional modulation formats and showed the first experimental demonstration of a 24-dimensional format, which achieved significant performance improvement over traditional methods. We continue to contribute to the development of new FEC methods and optimized codes.

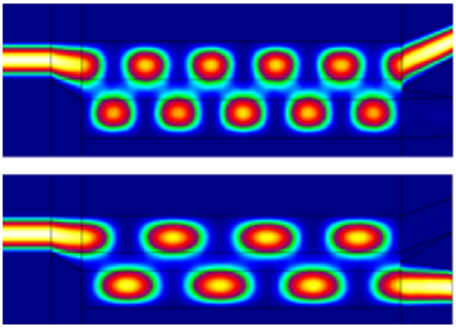

1 Tbps Optical Receiver

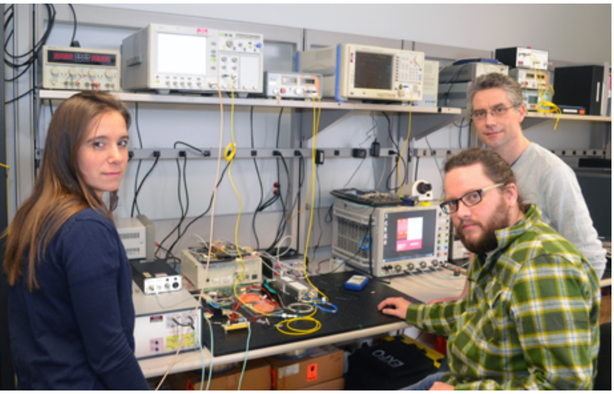

Optical fibers carry much of the world's data traffic, and the capacity requirements continue to grow by more than 20% annually to support connectivity among people and machines. Currently, commercial optical network data rates are limited to 100-200 Gbps per receiver.

MERL researchers developed a new multi-subcarrier transceiver technology that experimentally achieved a capacity of 1 Tbps using a single optical receiver. Spectral efficiency of 9.2 b/s/Hz was achieved, which was a world record for single receiver 1 Tbps transmission. A key feature of MERL's approach is the use of a pilot signal to enable high-quality compensation for signal impairment over a fiber. MERL's technology is compatible with currently installed optical fiber systems, allowing for a smooth upgrade path.

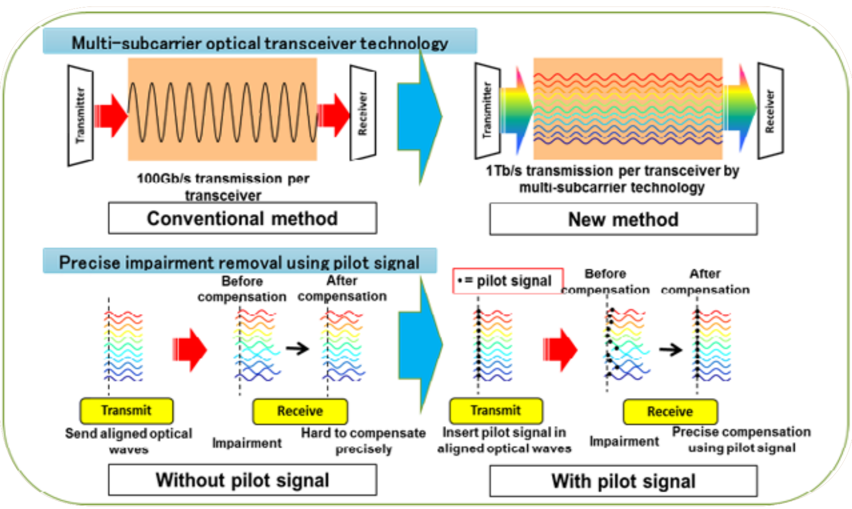

Wireless Communication

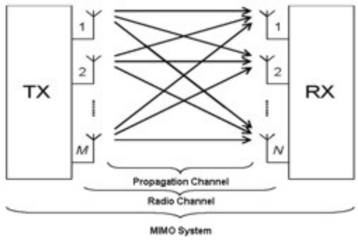

MERL's initial wireless communication research was triggered by the rapid growth wireless cellphone usage. In the late 1990s, we initiated signal-processing work for 3GPP standardization. Research topics included efficient channel state estimation [Molisch 2003] and feedback for pre-coding, multi-user detection and separation, channel equalization, Orthogonal Frequency Division Multiplexing (OFDM), and antenna selection. 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 standard in 2009.

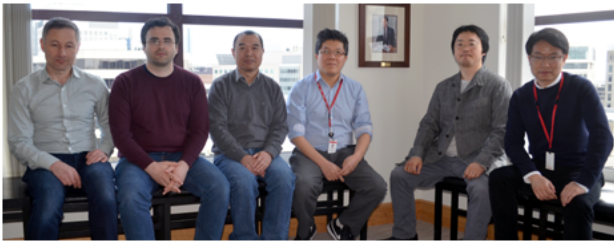

We expanded our activities in the early 2000s to include research on Ultra-Wide Band (UWB) systems, and new application areas such as short-range, high-definition video transmission, and indoor localization. We also addressed system-level issues such as ad hoc mesh networking, wireless channel access, and multi-hop routing, which led to work in routing methods for smart meter and sensor networks.

In 2008, we shifted our focus to ultra-reliable wireless for machine-to-machine (M2M) communication and new technologies for the Internet of Things (IoT). Our current research focuses on advanced multi-antenna systems for millimeter wave systems, error-correction coding for small-message M2M communication, and network-protocol design.

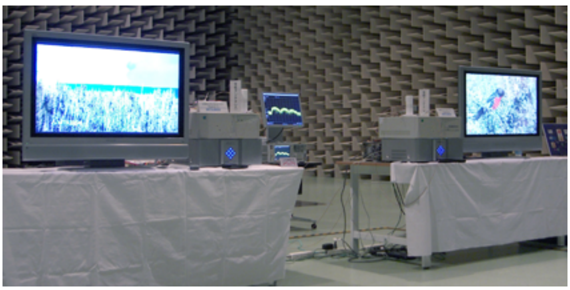

UWB Impulse Radio

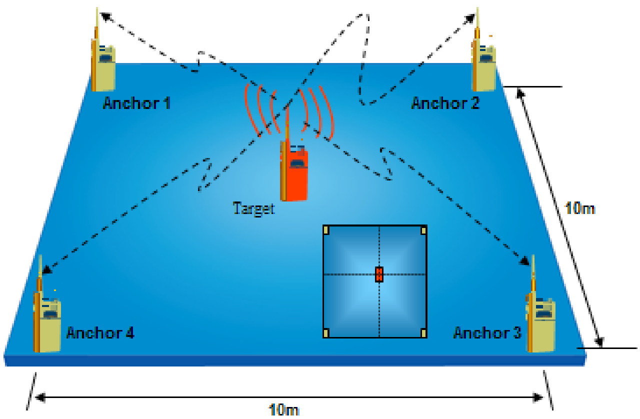

Ultra wideband (UWB) systems can simultaneously provide robust communication and accurate self-localization capabilities, making them ideal for wireless sensor networks and location-aware applications such as surveillance, infrastructure monitoring, and healthcare.

In 2007, MERL demonstrated the world's first prototype using UWB impulse radio technology to provide not only short-range wireless communication but also high-precision ranging [Gezici 2005, Zhang 2009]. 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.

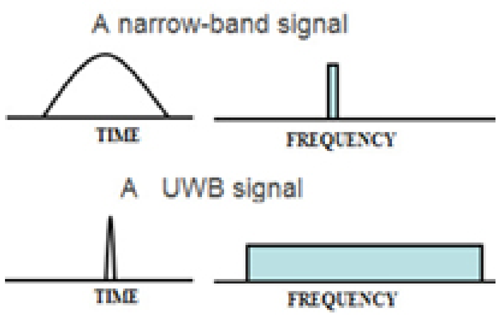

Mesh Networking

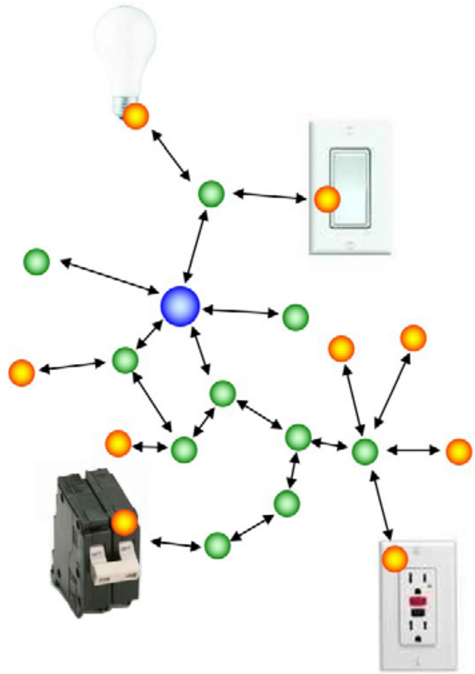

MERL's mesh networking research focuses on lightweight low-complexity network formation and routing algorithms so that simple low-cost devices can efficiently form robust mesh networks. Such networks are the backbone of what is now called the Internet of Things (IoT).

Mitsubishi Electric was a founding member of the Zigbee® Alliance in 2003. MERL researchers were key contributors to, and editors of, the first version of the protocol. Based on the ZigBee® standard, 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 [Pekhteryev 2005]. More recently, MERL has demonstrated the use of mesh networking technologies for a range of applications including smart metering of electric power consumption.

Our current research in this area focuses on reducing the overhead associated with network management, improving reliability by pre-computing backup routes, and distributing workload among nodes based on their capabilities and resources such as power, memory, and processing speed.

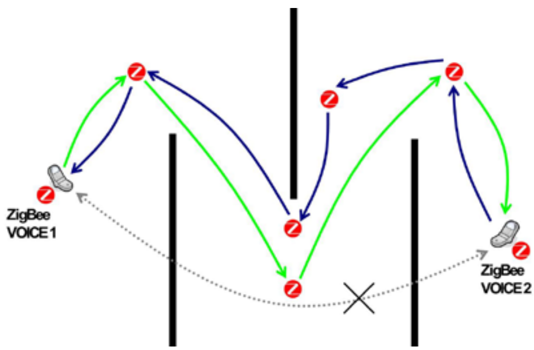

Antenna Selection

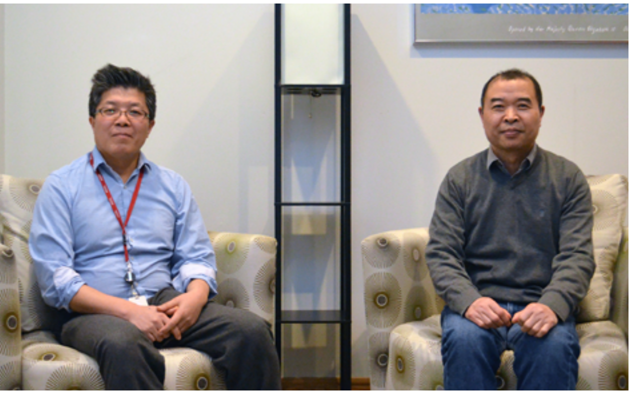

The simplest form of wireless communication is between a single transmitting antenna and a single receiving antenna, for example, between a base station and a mobile device. Unfortunately, at certain specific relative positions, multi-path interference and other problems can significantly degrade communication. Adding multiple antennas, preferably at both ends, increases the potential data rate and adds robustness, because the relative position of one transmit/receive antenna pair will not be the same as another pair. However, intelligent selection of which transmit/receive pairs to use at a given moment is needed to realize the full potential of a multi-antenna system.

MERL developed technology that enables the base station to dynamically determine the best antenna configuration for user equipment in terms of time and frequency transmission settings. This is achieved with the use of training signals that can be analyzed to determine the most effective transmission pattern. MERL's scheme was adopted as part of both the IEEE 802.11n and 4G LTE cellphone standards.

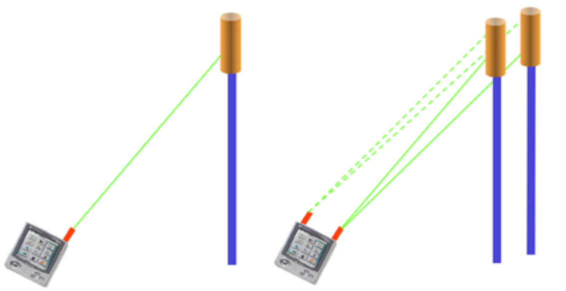

Ultra-Reliable Wireless

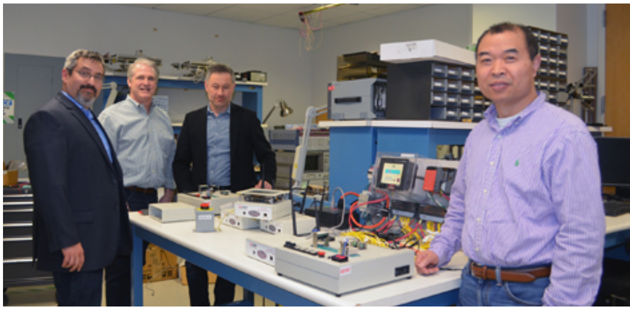

From 2008 to 2011, MERL researchers developed a proprietary wireless communication system, termed Ultra-Reliable Wireless (URW). This system achieves a message-loss rate of less than 10^{-8}, which is comparable to what can be achieved with a wired system. Through the use of a deterministic channel-access protocol, URW also ensures a delay of less than 10 milliseconds [Guo 2010].

MERL continues to work on novel wireless transmission techniques optimized for machine-to-machine (M2M) communication. By adapting ideas from multi-antenna cellular concepts to inexpensive,

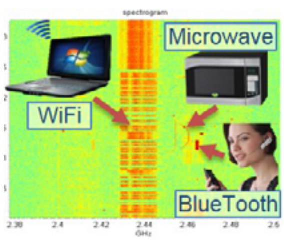

low-power M2M devices, we have achieved performance approaching theoretical limits. Additionally, MERL has developed machine learning-based methods to detect and classify interference sources using commercially available WiFi receivers, enabling the deployment of interference avoidance methods on WiFi-based M2M devices.

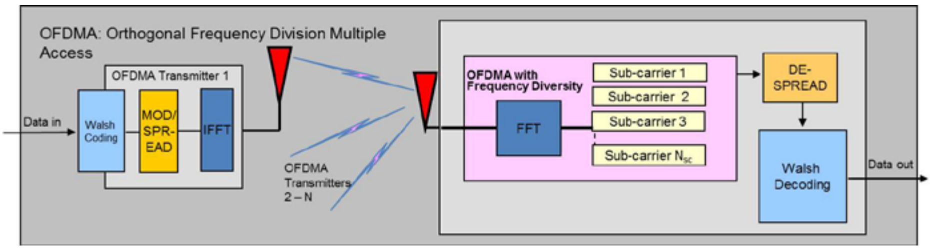

Indoor Localization

There are many potential applications for accurate positioning inside buildings or underground facilities where GPS signals are not available, such as walking navigation, warehouse logistics, and asset management. To address these needs, MERL developed techniques that augment existing WiFi equipment with acoustic signal processing to achieve position determination with less than one-meter error. The system locates a target based on the time it takes for audio signals to travel from the target to surrounding wireless access points equipped with positioning capability.

Using audio time-of-flight relaxes the synchronization requirements needed for radio time-of-flight measurements. This, combined with the use of standard WiFi hardware, makes the system significantly less costly than other solutions.

Power and RF Amplifiers

Since 2011 MERL has been investigating Galium Nitride (GaN) devices to take advantage of their high power and high frequency/switching-speed capabilities. We began our exploration with 3D multi-channel architectures. More recently, we investigated GaN power devices using diamond as the back barrier. Preliminary simulation results confirm the large potential of incorporating such layers to increase the breakdown voltage and improve thermal performance.

We developed several power amplifiers, achieving 10 W 1 GHz bandwidth class-J operation at 2 GHz, and broadband Doherty operation of 400 MHz bandwidth at 3.5 GHz.

We are currently researching complete digital transmission circuits, focusing on novel transmitter architectures and power-encoding algorithms. MERL was the first to propose a multi-dimensional out-phasing power-coding scheme. Our approach achieved world record coding efficiency.

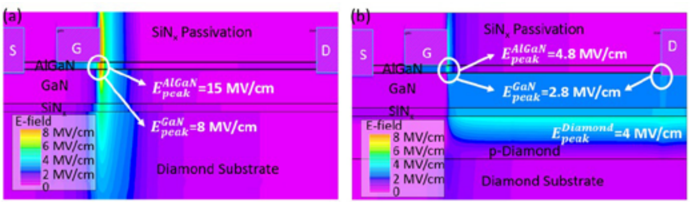

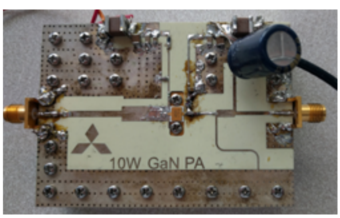

Wireless Power Transfer

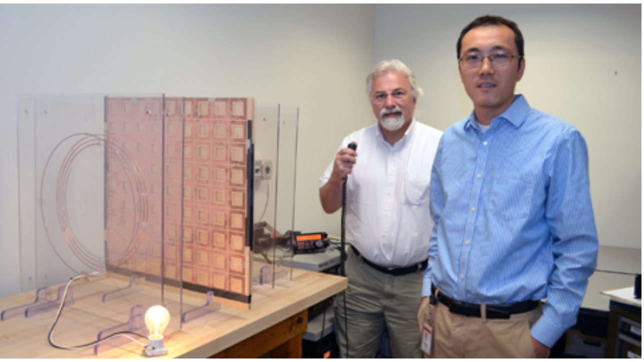

Wireless power transfer (WPT) is a promising technology for many applications, from consumer electronics to industrial automation. MERL's research extended the range and improved the flexibility of WPT systems.

In 2009 MERL was the first to use meta-materials for enhanced WPT [Wang 2011]. These artificially engineered materials are composed of structures designed to achieve unique properties that cannot be obtained with natural materials. We showed with simulation and experimental results that a thin slab of metamaterial can boost the near-field coupling between two resonant coils, resulting in a tripling of the power transfer efficiency between them.

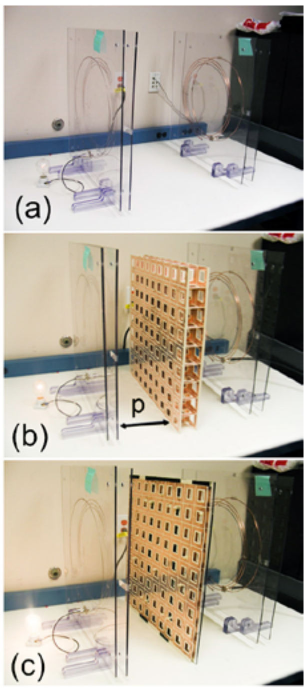

We also introduced the concept of using an array of coupled resonators between a WPT transmitter and receiver. This greatly extends the range of the WPT link and allows the receiver to move over significant distances while continuously getting power.

Computational Sensing

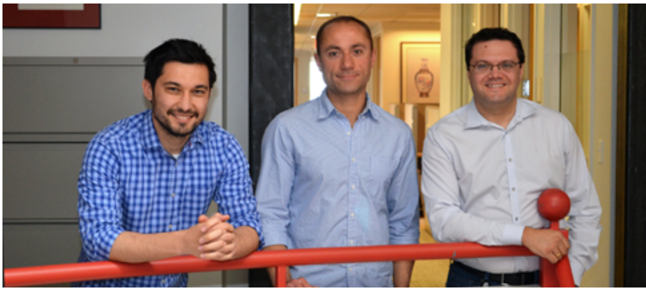

Sensing and signal acquisition are indispensable technologies in our increasingly digital world. MERL exploits widely available computational power to overhaul the signal acquisition paradigm and significantly enhance sensing capabilities [Davenport 2010].

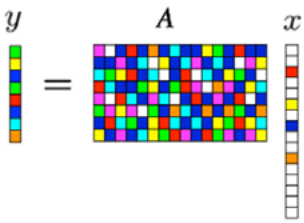

Sensing systems aim to capture physical signals and reconstruct them digitally or extract useful information from them. We aim to fundamentally understand how signals behave and propagate in the environment and the sensing systems. Using this knowledge, we develop reconstruction algorithms to recover signals with significantly improved fidelity and robustness.

Co-developing sensing systems with novel computational methods enables us to realize the full potential of the sensing hardware. Our research results in high resolution sensing of the environment at significantly lower cost.

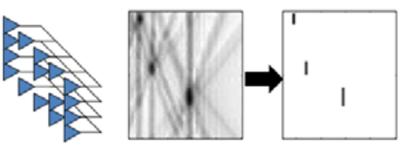

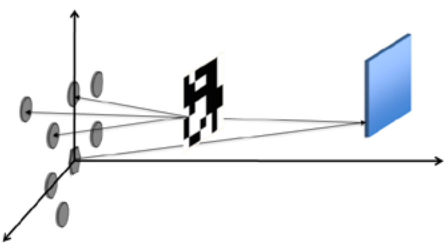

We have successfully applied our methods to advance the state of the art in a number of applications, including high-accuracy low-cost active depth sensing, fusion of signals from multiple sensors with different modalities, accurate sound localization with microphone arrays, and improved resolution in radar imaging systems.

Compressive Sensing 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.

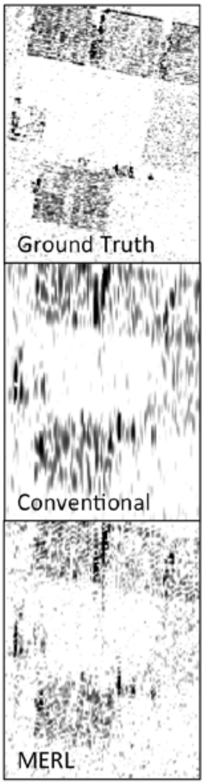

One of our first designs demonstrated substantially increased azimuth resolution using a high-rate non-uniform pulsing scheme, allowing interference between pulses. The same design can be configured to double the observation area instead. Rather than alter the pulsing rate, follow-up designs utilized 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 [Liu 2013] was recognized with the 2014 Symposium Prize Paper Award by the IEEE Geoscience & Remote Sensing Society.

Information Security

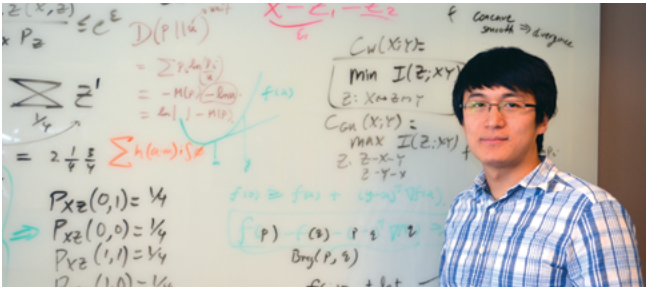

Particularly in the emerging world of the Internet of Things (IoT), privacy and security concerns are becoming ever more important. Information security research at MERL covers a range of topics addressing these concerns.

We have developed protocols using homomorphic encryption and secret-sharing techniques that enable privacy-preserving data analysis in the presence of untrusted parties — gaining new insights into the fundamental limits of secure computation.

We conduct research on statistical inference under privacy constraints, providing a computationally inexpensive method for data analysis while assuring strong statistical guarantees of security via anonymization of data.

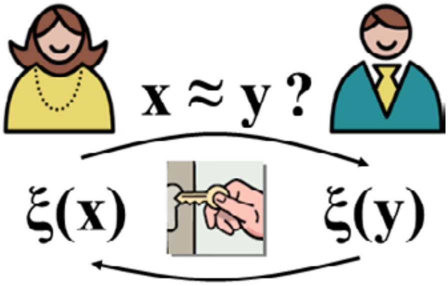

MERL researchers have achieved world-leading performance in terms of accuracy, robustness, and security for fingerprint authentication systems. Our theoretical developments have become the basis for international standards.

Predictive Modeling

MERL's research on predictive modeling of systems and processes from collected sensor data is based on statistical machine learning algorithms. The unifying purpose of these models is to enable fast and accurate prediction of outputs and effects resulting from the application of controllable decision variables as well as uncontrollable operating conditions, for the purpose of finding the most optimal settings for the decision variables. They can be used for the optimal control of equipment and also to predict when equipment will fail or need maintenance. They can also be used to predict how people will react in given situations based on their past behavior.

MERL employs advanced statistical machine learning algorithms for classification, regression, and anomaly detection to learn predictive models entirely from data, eliminating the expense and effort of their manual construction and calibration. We have focused on predictive modeling algorithms that can operate under demanding industrial conditions: large data volumes, streaming measurements, multiple data sources, and limited memory and computational resources. We have developed algorithms for abrupt change detection of streaming data and exemplar learning with limited memory that are best in their class.

Decision Analysis and Optimization

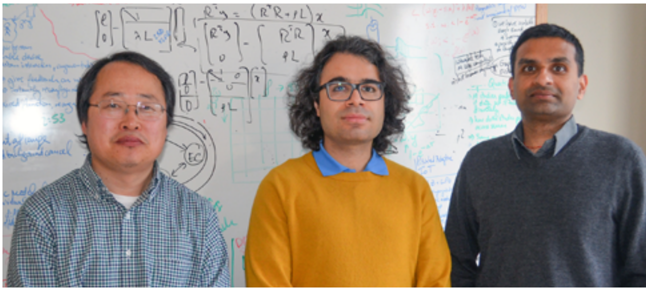

The field of decision analysis and optimization is concerned with making optimal decisions that improve a performance metric, such as cost, comfort, price, or efficiency. MERL has developed advanced optimization algorithms to attack hard industrial problems such as voltage & power optimization in electrical power networks, temperature set-point scheduling for air conditioners, group elevator scheduling, and various transportation problems including vehicle routing, speed profile optimization, and battery management.

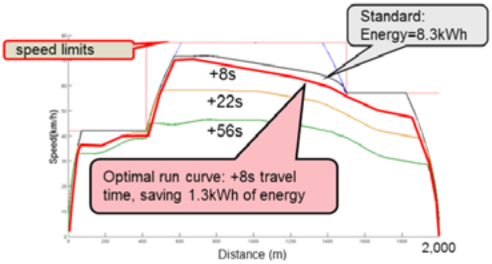

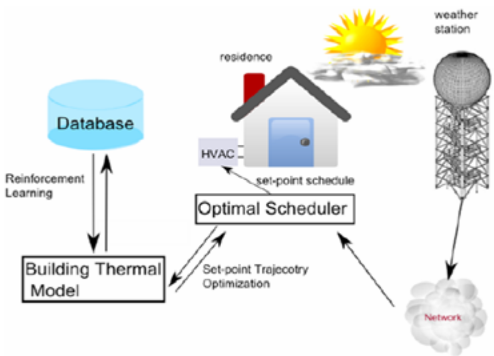

For optimization problems with continuous decision variables, such as feeder voltage optimization for electrified trains, we have been able to solve very large problems (many thousands of variables) in a fraction of a second, even for non-convex situations. In addition, we have developed several algorithms for sequential decision making based on approximate dynamic programming that can solve problems with very high-dimensional state spaces (thousands of dimensions) without suffering from the "curse of dimensionality," enabling efficient operation on embedded devices with limited computing power.

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.

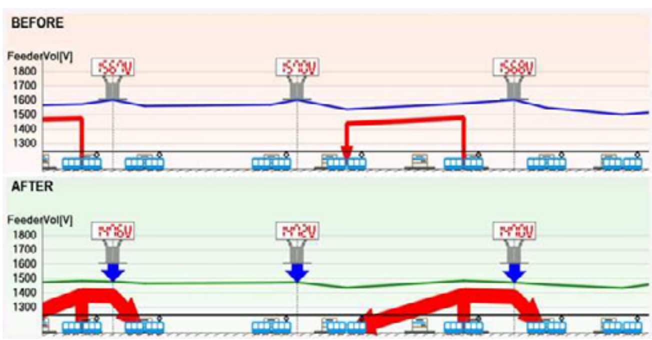

Digital Video

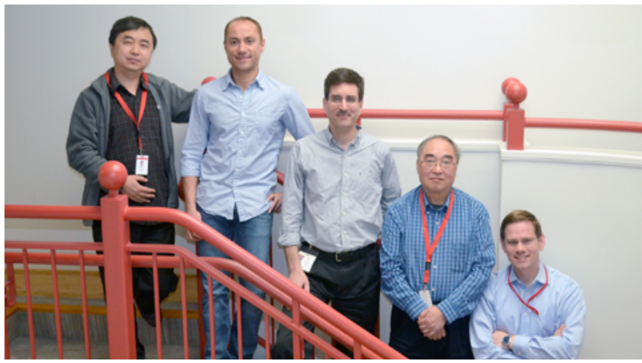

The transition from analog to digital television in the U.S. triggered the beginning of digital video work at MERL. In the mid-1990s, we contributed to the design of the first HDTV receiver chip set and began participating in the development of MPEG video-coding standards. During this period, we also worked on encoding algorithms to optimize video quality, and fast video-signal-processing techniques that can remove compression artifacts for consumer electronic devices.

Our work began to shift in the early 2000s, addressing a wider range of problems in video delivery and analysis. We developed leading algorithms and architectures for video transcoding, and became a driving force in both academic and standardization communities on advanced solutions for multimedia adaptation. 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.

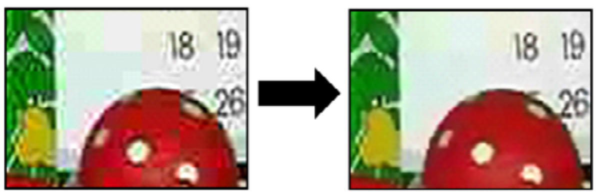

The current thrust of our work is centered on advanced compression technologies, including compression of rich and immersive media formats, visual features, and remotely sensed imagery, as well as advanced visual analysis techniques, including the analysis of human activities and motion flows.

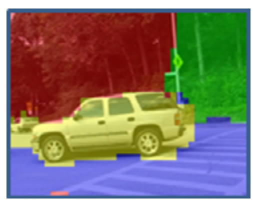

Video Transcoding

Video transcoding is the process of converting a video in one compressed format into another. MERL developed a highly efficient method that intelligently reused compression information from the input stream, rather than merely decoding and then recoding the video [Vetro 2003]. A theoretical analysis of the errors accumulated due to changes in quantization and spatial resolution was described in MERL's award-winning papers [Vetro 2002, Yin 2002].

In 2004, Mitsubishi Electric began shipping a PC-based product for transcoding. The system simultaneously transcoded multiple MPEG-2 streams from surveillance cameras into compact, low-resolution MPEG-4 streams, delivering them to remotely-connected devices. The transcoding was done entirely in software using MERL's technology.

Video transcoding has since matured into a billion dollar industry due to the rise in media streaming, coupled with the explosion of different media formats and device capabilities.

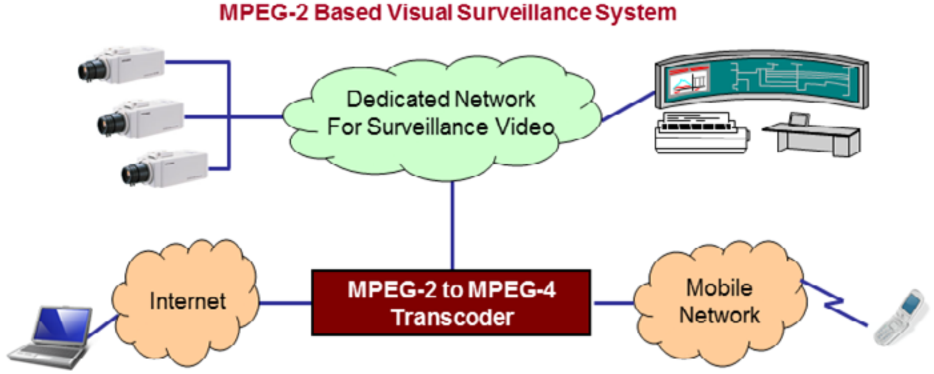

Video Coding Standards

MERL has actively contributed to the development of video-coding standards for more than 20 years. For example, MERL made contributions to the base specification of the state-of-the-art 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 held leadership roles and provided key technology to all the major 3D and multi-view video-coding standards developed and published in the past decade. Most notably, this includes a Multiview Video Coding amendment to the H.264 Advanced Video Coding (AVC) standard, which has been commercially deployed for 3D Blu-ray Discs and specified as a 3D broadcast format by the Digital Video Broadcasting (DVB) and Advanced Television Systems Committee (ATSC) organizations [Vetro 2011].

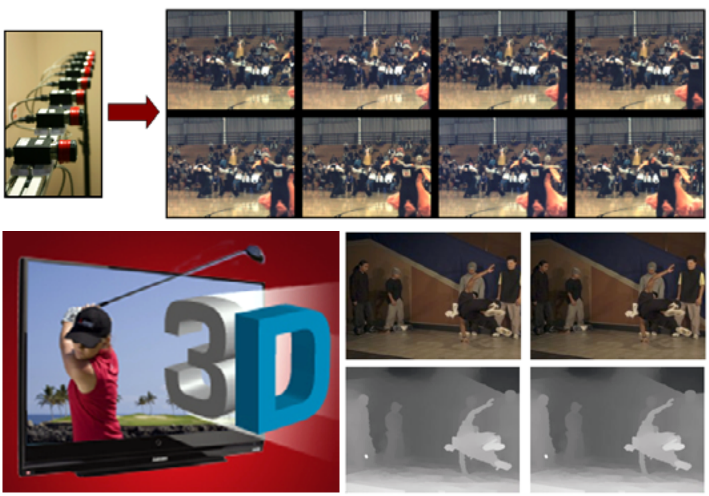

More recently, MERL led and contributed to the development of the first coding standard to incorporate depth information for view-rendering purposes, and pioneered new techniques to leverage such information as part of the coding process.

Video Highlights Playback

In 2005, Mitsubishi Electric began selling a DVD recorder with 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 [Xu 2001]. 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.

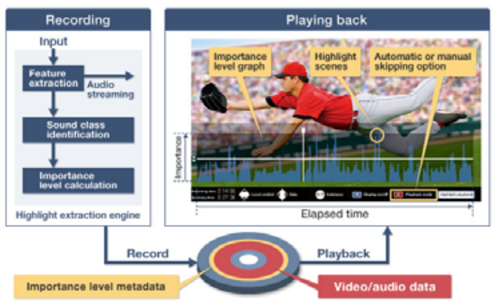

3D TV System

In 2004, MERL developed the first TV system supporting real-time 3D acquisition and 3D display of dynamic scenes [Matusik 2004]. The system used an array of video cameras to simultaneously capture multiple views of a scene. Clusters of network-connected PCs were used to encode, transmit, decode, and process for viewing the captured video. An array of video projectors was used to project multiple video streams onto a lenticular screen. Viewers looking at the screen could see stereo video pairs from multiple viewpoints without the need for special glasses. To manage the high demands on computation and bandwidth, a distributed, scalable architecture was designed. To facilitate setting up the system, automatic digital alignment of the various cameras and projectors was developed.

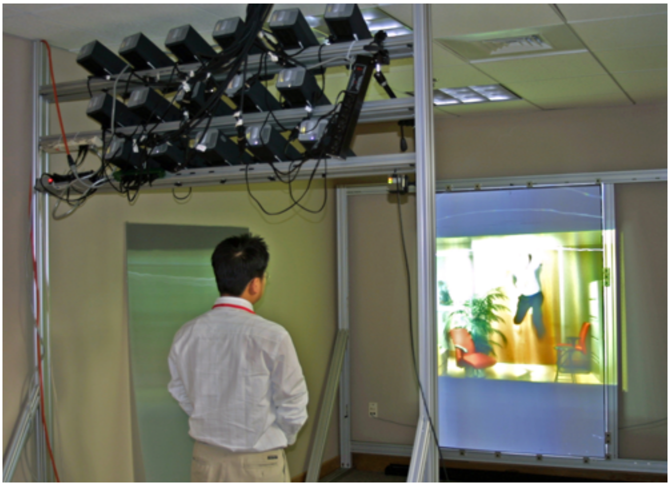

HDTV Receivers

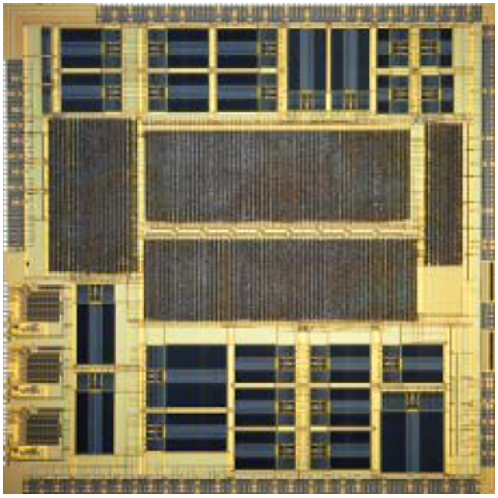

In 1998, Mitsubishi Electric and Lucent Technologies' Bell Labs jointly completed the development of the first integrated chipset meeting the North American standard for receiving and displaying High-Definition Television (HDTV) signals. The chipset supported demodulation, demultiplexing, audio and video 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.

For a 2009 display-processing chip, MERL developed advanced methods of noise reduction, enhancing the picture quality of LCD TVs.

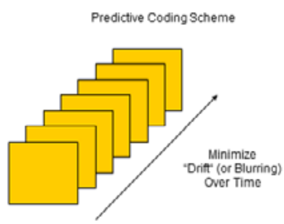

Computer Graphics

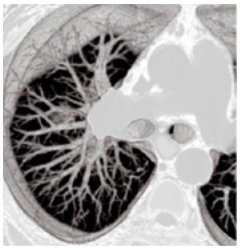

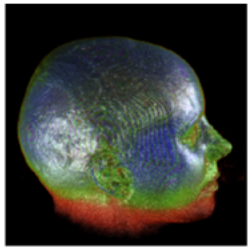

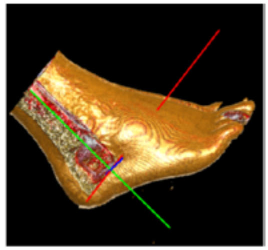

Although Computer Graphics is presently not a large area of research at MERL, it was a major area, particularly from the late 1990s into the early 2000s.

The dominant representation for 3D graphics models is collections 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 two entirely new representations: Surfels – a 3D cloud of oriented surface elements, and Asymmetrically-Sampled Distance Fields (ADFs) – a parametric representation of the distance between each point in 3D space to the nearest surface.

In addition, MERL developed a number of innovative applications combining cameras and projectors, and using graphic-rendering techniques to pre-distort images sent to the projectors, in order to create innovative effects.

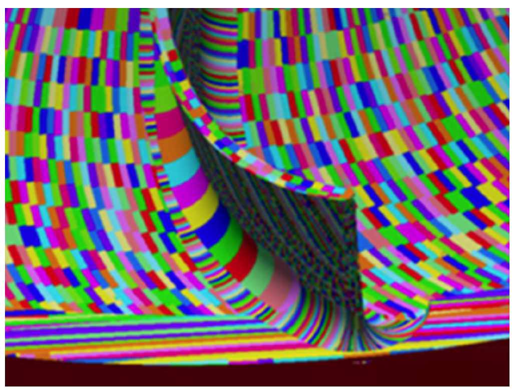

Real-Time Volume Graphics

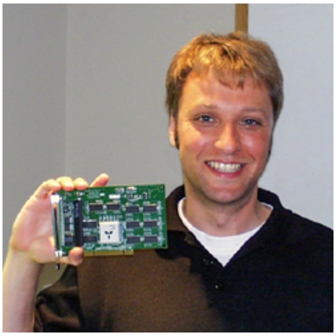

First demonstrated at the SIGGRAPH conference in 1999, VolumePro [Pfister 1999] was the first PC graphics card for real-time visualization of 3D voxel data from scientific and medical instruments such as computed tomography (CT) and magnetic resonance imaging (MRI) machines. At the time, it was the fastest 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. The technology was subsequently acquired by TeraRecon and is still in use by customers in medicine, biology, and engineering.

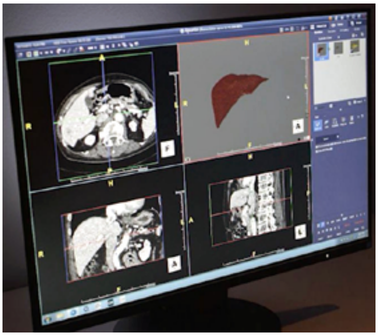

Surfels

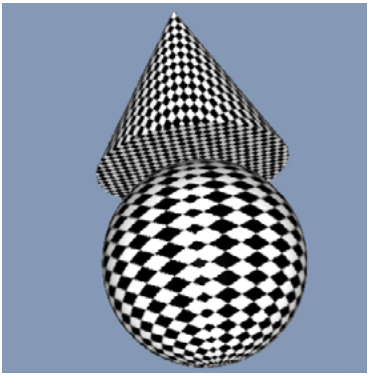

“Surfels” (short for Surface Elements) [Pfister 2000] are small patches with 3D positions and lighting information such as color, texture, and orientation. The key objective of Surfels is to act as a *lingua franca* that can unify the rendering of polygonal graphics models, volume data, and scanned point clouds into a single rendering framework. The concept of Surfels spawned a point-based graphics research sub field, with contributions in the areas of dynamically changing models, graphic data compression, and GPU accelerated rendering.

From a modeling perspective, Surfels provide a discretization of the geometry reducing the object representation to the essentials needed for rendering. In a preprocessing step, the surfaces of complex geometric models are sampled along three orthographic views. At the same time, computation-intensive calculations such as texture mapping are performed. By moving rasterization and texturing from the core-rendering pipeline to this preprocessing step, rendering cost is dramatically reduced. This is particularly beneficial for models with rich, highly complex shapes or highly detailed surface coloring.

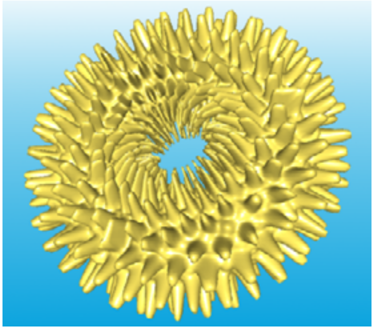

Adaptively Sampled Distance Fields

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 [Frisken 2000]. ADFs use a detail-directed adaptive partitioning of space and can represent very high levels of detail, using much less memory than polygonal approaches.

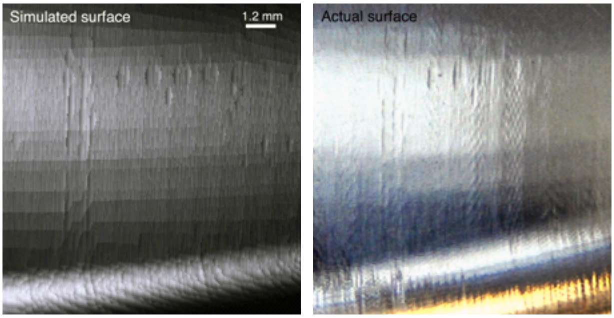

In 3D, ADFs have been applied to Computer Numerical Control (CNC) milling to create extremely accurate simulations of complex parts that can reveal minute details of a milled surface to help diagnose manufacturing problems.

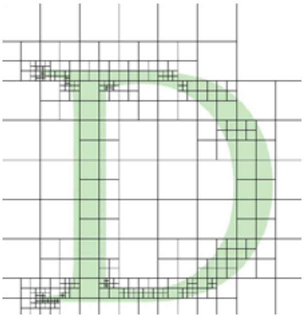

In 2D, ADFs have been used to compactly represent fonts that can be scaled to arbitrarily large sizes and rendered clearly even at very small sizes. They require less memory than other approaches and can produce high-quality rendering across all sorts of displays, from computer screens to hand-held devices.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

Projectors

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” [Raskar 2006].

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.

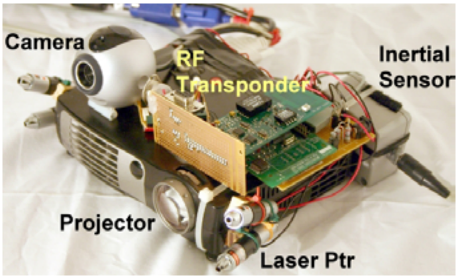

Human-Machine Interfaces

HMI research per se is not a major focus of MERL's research today, but it is an important component of many projects, especially in speech interfaces and Computer Graphics. In addition, several MERL projects in past years had a strong influence on the HMI research community and ultimately on the innovations appearing in current and future human-machine interfaces.

MERL's DiamondTouch table is still the world's only 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.

MERL also combined innovative sensors and hardware to explore new concepts of interaction among people, computers and data. Simple IR sensors were integrated with cameras to develop ambient sensing systems for building security and safety [Ivanov 2007].

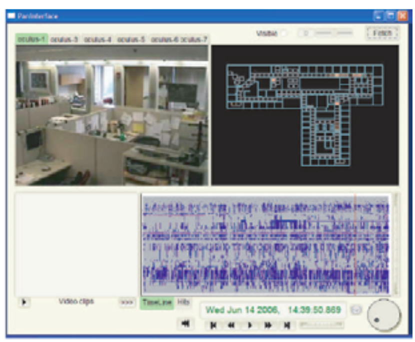

DiamondTouch

The DiamondTouch table [Dietz 2001] is a front-projected interactive display that allows multiple users to sit face to face and work together on the same screen. The key difference between DiamondTouch and other multi-touch interfaces is that multiple users can touch the table at the same time and DiamondTouch can keep track of who is doing what. 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.

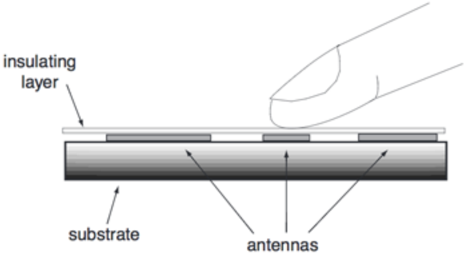

The surface of the table contains two perpendicular transmitter antenna arrays. When users touch the table, they complete circuits between these arrays and receivers in the chairs they are sitting on via capacitive coupling. Because each user has a different receiver, DiamondTouch can determine which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

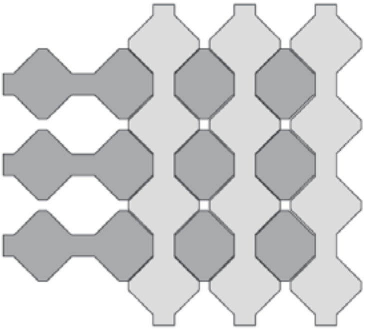

Diamond Park

Demonstrated at COMDEX in 1995, Diamond Park [Waters 1997] 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 important, it was the first to support arbitrary modification and extension of the environment during operation, which are essential features of the massive online gaming environments of today.

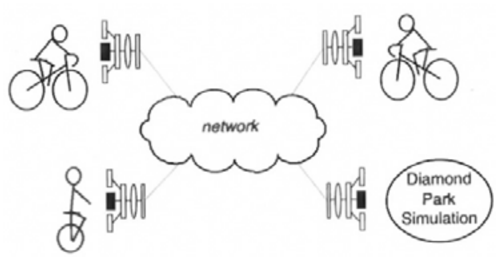

Collaborative Agents

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 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 [Rich 1998]. Collagen 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 also been influential on subsequent dialogue-processing research.

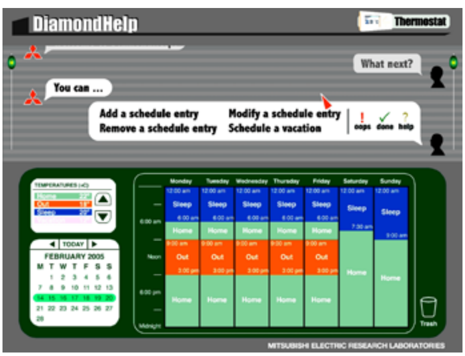

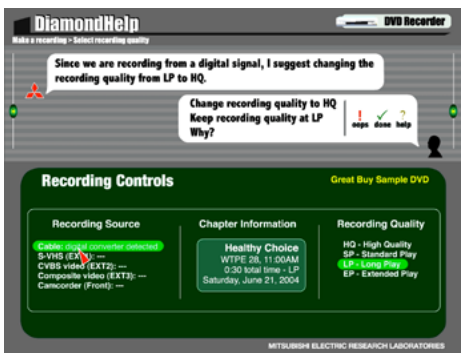

Speech and Audio

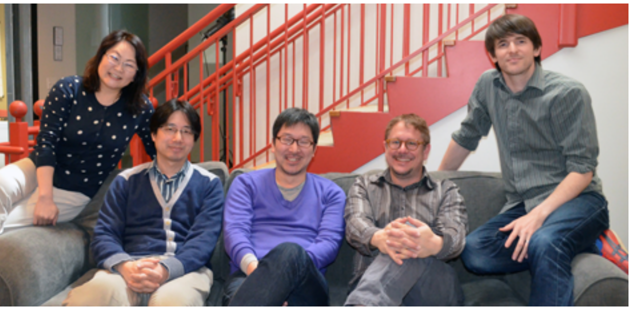

The speech and audio team at MERL is pursuing a range of challenging speech and audio machine-perception problems involving acoustic signals, human language, and everything in between. On the signal acquisition side, we focus on source separation, both in single- and multi-channel settings. For speech recognition, we develop novel deep-learning methods for acoustic and language modeling. On the language side, our efforts include natural-language understanding, topic modeling, user-action prediction, and interactive multi-modal human-machine interfaces.

A particular difficulty in this field is bridging the gap between two differently structured domains: language and acoustic signals. Our approach has been to develop novel machine-learning approaches 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.

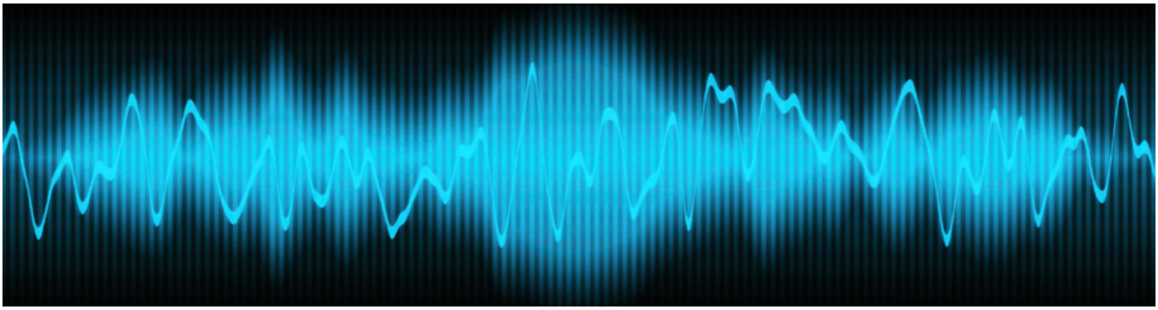

SpokenQuery

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

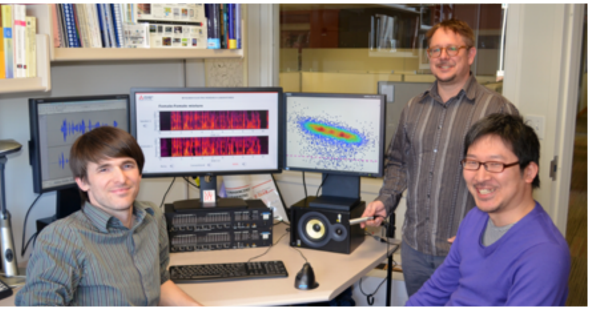

The human auditory system enables us to do an extraordinarily good job conversing above the chatter of a lively cocktail party. Endowing a computer with this ability has been the holy grail of speech processing for over 50 years.

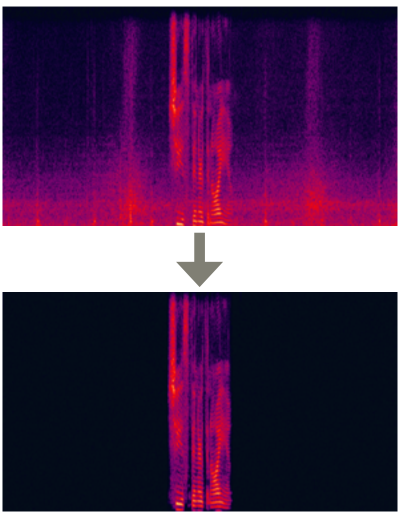

For 15 years, MERL has pioneered advances in source separation. Combining originality with solid theoretical foundations, MERL's methods have dramatically influenced the field, and solved problems that were previously out of reach. Early highlights include seminal work on non-negative matrix factorization of audio [Smaragdis 2003], audio-visual enhancement, bandwidth expansion, and speech separation.

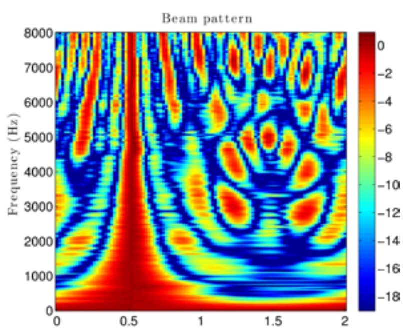

Recently, by harnessing the power of cutting-edge deep learning, our algorithms have been able to overcome far more challenging interference than ever before. Some of our current developments are based on a fundamentally new way to use domain knowledge as the basis for novel deep-learning architectures. We believe our latest technology, *deep clustering*, is poised to solve the general audio-separation problem, opening up a new era in spontaneous human-machine communication.

Optimal Control

Striving to achieve high performance under tight operating constraints, MERL develops and uses a large class of optimal and optimization-based control methods.

MERL concentrates on innovative methods to synthesize optimal control policies for a broad spectrum of dynamical systems. These methods have been applied to optimal motion planning for servomotors and automotive, and to simultaneous optimal design of multi-physical systems.

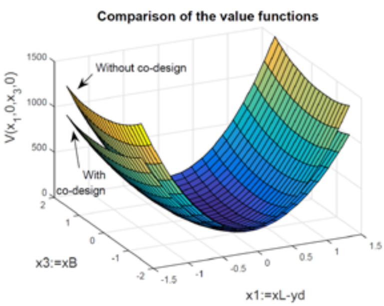

MERL is extending Model Predictive Control (MPC) for constrained multi-variable systems. Our developments include new controller designs that provide performance guarantees for systems with uncertain or switching dynamics, and optimization algorithms suitable for low-cost microprocessors. Applications include vehicle and railway control systems, communication satellite station keeping, and motion control in factory automation.

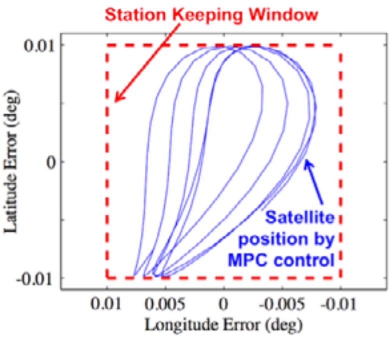

MERL's research in constrained control includes the investigation of reference and command governors, the construction of control invariant sets for uncertain and switching systems, and their use for control design with application to automotive, railways, and motion control.

High-Speed Laser Cutting

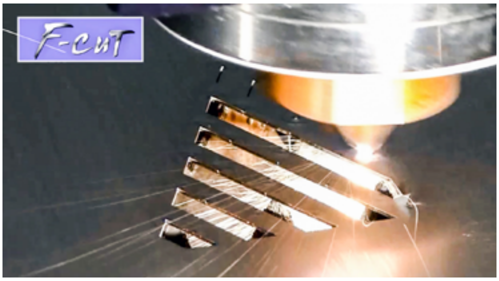

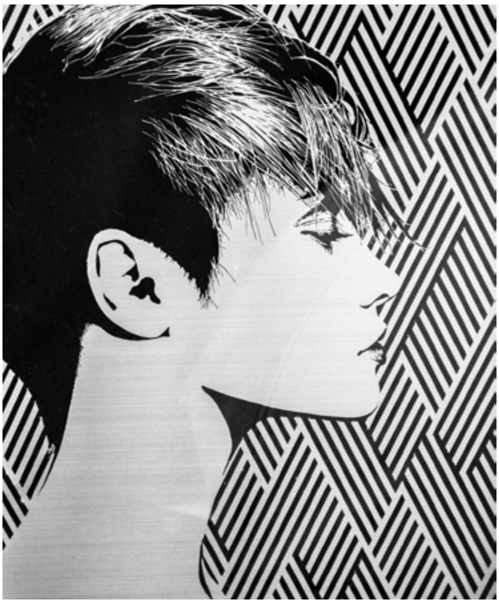

Mitsubishi Electric's Laser Cutting Machine uses a high-power laser on a moving head to rapidly and accurately make cuts in sheet metal. It can be used to make a wide range of flat metal parts. To maximize the value of the machine, one would like to make a pattern of cuts as fast as possible.

The time it takes to make a set of cuts depends on the order the cuts are made 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 compared to conventional approaches.

Nonlinear Estimation and Control

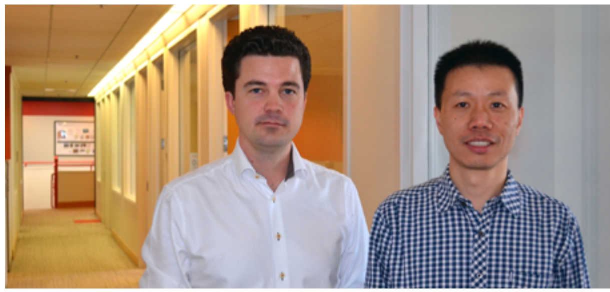

MERL does fundamental research in nonlinear estimation and control with applications to transportation, aerospace, robotics, and factory automation systems.

MERL does research on estimation of nonlinear dynamical systems with both sensing and model uncertainties. We have developed feedback-policy particle filters that demonstrate excellent accuracy and improved robustness to intermittent measurements. We have also developed particle-filtering with reinforcement-learning algorithms for state and parameter determination for transportation systems and robots.

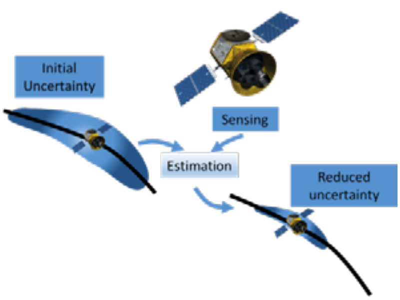

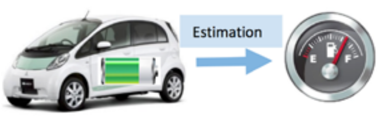

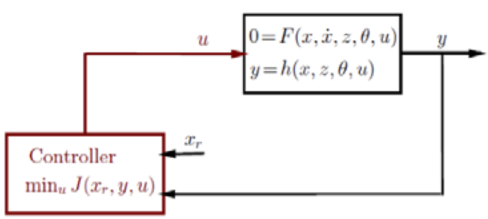

We use both normal-form observer (model-based) and extremum seeking (model-free) approaches. This research is driven by various industrial applications including asynchronous motors, energy-storage systems, elevators, transportation systems, and motion reconstruction.

Standard nonlinear control designs require high-fidelity models with scalability and performance guarantees. We leverage novel mathematical tools to relax this requirement. We are particularly interested in design tools enabling robust, adaptive, incremental learning, and optimal control of Euler-Lagrange & networked systems.

Dynamical Systems

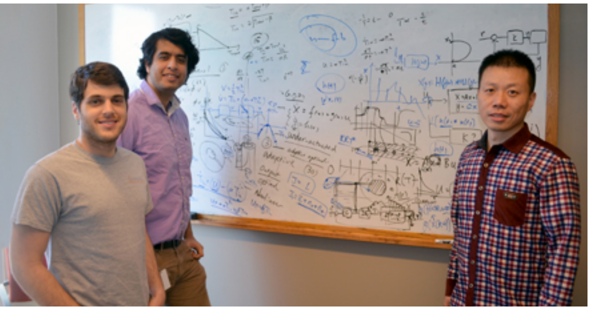

MERL conducts fundamental research and algorithm development in the area of dynamical systems, with applications in space systems, vibration mitigation, and HVAC.

Analytical and computational methods from dynamical systems theory are used to discover geometric structure in phase space of nonlinear systems. We develop algorithms that exploit this structure for applications such as space probe trajectory planning, and the design of nonlinear vibration absorbers.

We work on developing reduced-order models for thermo-fluid systems. To enable real-time estimation and control, it is crucial that essential nonlinear features be preserved in the reduced models. Using a combination of operator theory and control theory-based learning algorithms, we have developed reduced-order models for fluids that are successful in reproducing complicated dynamics of full-order models.

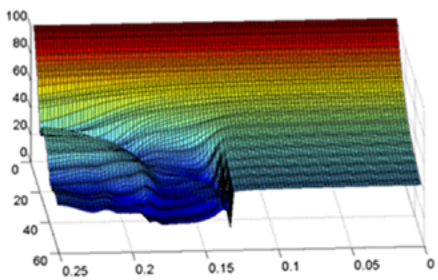

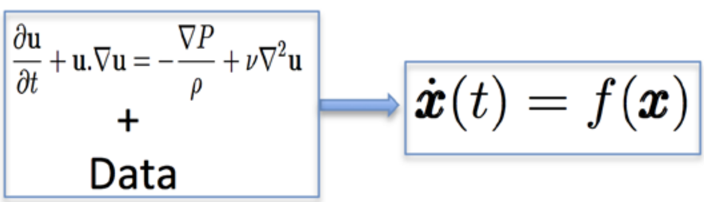

Spacecraft Trajectory Design

MERL has developed an end-to-end trajectory design framework for space missions to the moon and beyond using 3-body gravitational dynamics. Small resonant perturbations from the moon's gravity can increase the size of a spacecraft's orbit, even when it is far from the moon's primary sphere of influence. Additionally, there exist pathways in the Earth-moon system through which a spacecraft can transit from an Earth orbit to a moon orbit, without using any fuel.

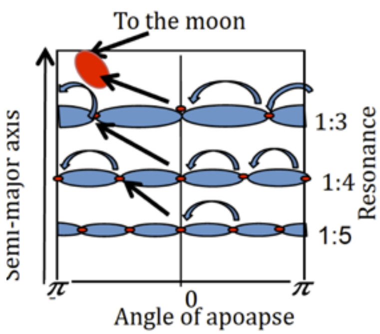

By using dynamical systems theory to compute trajectory segments that exploit these mechanisms and patching them together, we obtain an initial guess of a good trajectory from the Earth to an orbit around the moon. Using pseudospectral collocation and other nonlinear optimization tools, this guess is optimized to satisfy required mission fuel and time constraints.

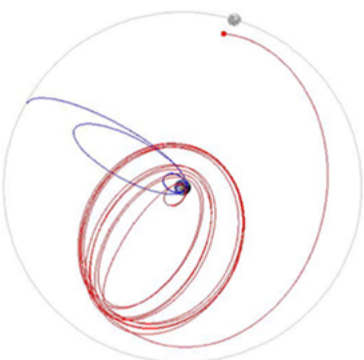

MERL has computed trajectories that can save up to 7% of fuel while reaching the moon in only five months. Since fuel weight is a large fraction of total spacecraft weight, a 7% fuel savings enables a large increase in payload.

Thermo-Fluid Systems

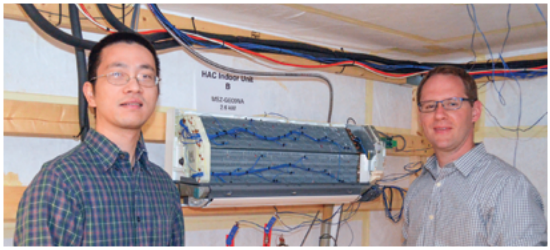

With the goal of dramatically improving the energy efficiency of buildings, MERL does research on three aspects of thermo-fluid systems.

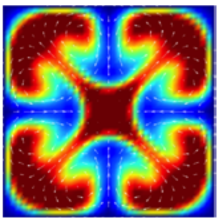

MERL is researching new design methodologies to optimize mixing and heat transfer using a combination of nonlinear dynamical systems theory and optimization. These methods can be applied to problems ranging from more compact and efficient heat-exchanger designs to next-generation building ventilation.

MERL is developing innovative system-level dynamic modeling tools for the simulation and analysis of Variable Refrigerant Flow (VRF) building-scale HVAC systems. This is difficult due to the large number of states required to accurately model both the equipment and the building, the nonlinearities associated with refrigerant phase change, and the multi-physical nature of heat transfer and mass transport in a building.

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \vec{v}) = 0$$

$$\rho \frac{D\vec{v}}{Dt} = \rho \vec{g} - \nabla p$$

$$\rho \frac{Dh}{Dt} = \frac{Dp}{dt} + \dot{Q}$$

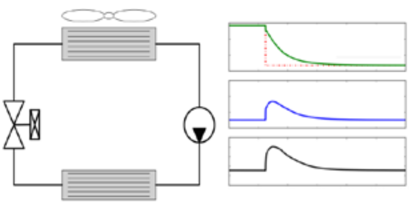

MERL has developed a new generation of control algorithms for HVAC systems based on Model Predictive Control (MPC). These algorithms optimize energy efficiency, improve temperature regulation, and enforce important constraints associated with equipment's performance.

Optimization

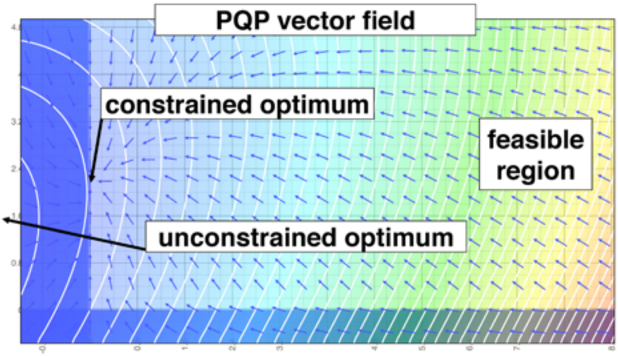

Much of MERL's research activity involves reformulating scientific and engineering problems as optimizations. In addition, we develop new algorithms for canonical optimization problems, because an advance in that area can push many other technologies forward making heretofore infeasible approaches practical.

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. Although the analysis runs to many pages, the algorithm itself is remarkably simple — a single line of code. That makes it is easy to validate and verify PQP implementations in critical systems where the software must be reliable and provably correct. 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 bi-criterion minimum-cost path, alternating direct method of multipliers (ADMM), hyperbolic programming, and numerical preconditioning.

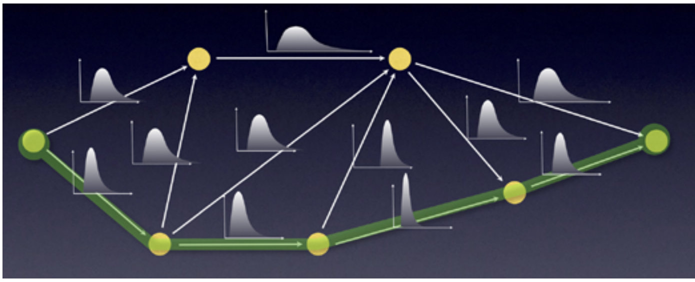

Dimensionality Reduction

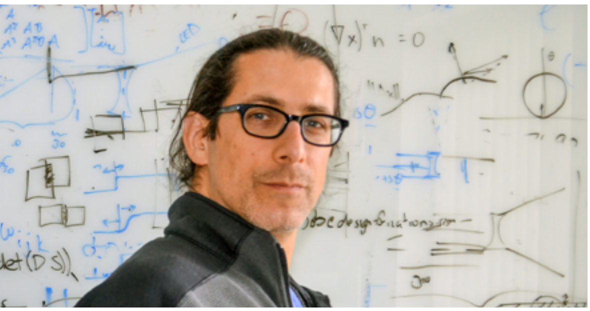

Machines often observe processes in the world that have limited degrees of freedom with sensors that have millions. These extra degrees of freedom capture irrelevant data that obscures the process being observed. (Consider the millions of pixels in a video of the one-dimensional motion of a pendulum.) It is useful to think of the data or video frames as points sampled from a low-dimensional manifold that is embedded in the high-dimensional ambient space of the sensor. Understanding the manifold can tell us a lot about the configuration space and dynamics of the process.

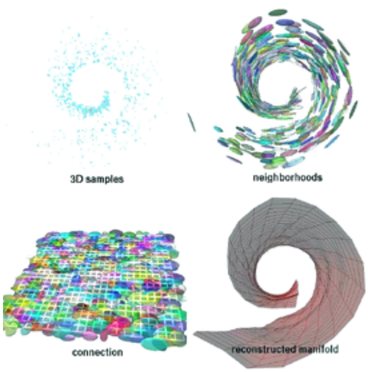

MERL's Manifold Charting [Brand 2002a] introduced a way to approximate the manifold as a smoothly blended set of subspaces. It reveals the intrinsic dimensionality of the manifold and establishes a coordinate system on the manifold so that the dynamics of the system can be modeled. Manifold Charting launched numerous applications inside and outside of MERL, including extracting pose from video and representing the semantic structure of language analogies geometrically. A successor technology, Nullspace Analysis, remains popular in the field as an extension to singular value decomposition that "unrolls" embedded manifolds.

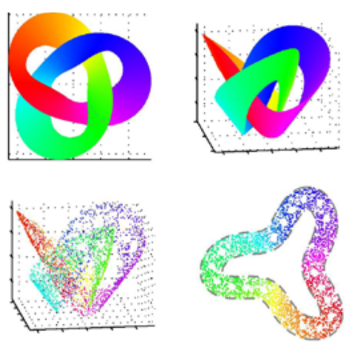

Incremental SVD

Singular Value Decomposition (SVD) reveals the most interesting variations in a data set and is the core engine of thousands of algorithms in signal processing, data analysis, and machine learning. However, traditional SVD requires that all the data be present at once and has cubic-time complexity, limiting its usefulness for really big "big data."

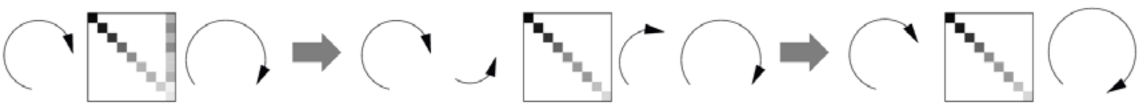

MERL's incremental SVD [Brand 2002b] is an online SVD-updating algorithm that averaged over many elements, runs in linear time. This makes it possible to compute the SVD of terabyte-scale data streams. SVD-based machine-perception algorithms, such as determining non-rigid 3D structures from single-camera observations, suddenly become practical in real time.

MERL implemented a real-time Movie Recommender that used real-time SVD updating to learn and exploit correlations in customer's tastes. Years before the Netflix Challenge, MERL's

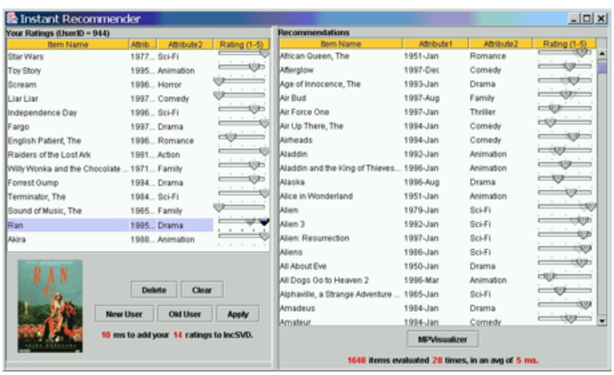

Movie Recommender made the most accurate recommendations on the largest movie-watching database then available. Today, the software is widely used in fields ranging from e-commerce and data mining to seismology and weather prediction.

Seam Carving

When a picture does not have the right aspect ratio to fit in the space available, there are two basic options: crop it or distort it. Seam Carving [Avidan 2007] introduced a new option: recompose 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 to go viral on the Internet. MERL 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. Ten years on, these algorithms remain popular in photo-editing software.

Computational Mathematics

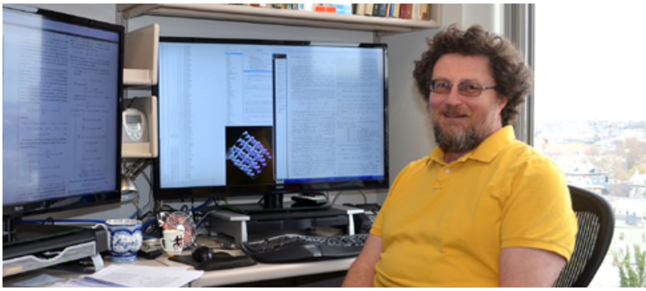

MERL exploits its expertise in computational mathematics to design algorithms that go beyond traditional paradigms to improve computational performance by orders of magnitude.

Model Predictive Control (MPC) uses a model of the physical properties and constraints of a system to control it better, but requires significant computation. Nonlinear models are the most general and at the same time most computationally challenging. MERL has developed real-time control algorithms that operate 10-100 times faster than traditional methods.

Graph-based approaches extend the traditional notation of spatial signals to signals determined on graphs, for example, data from sensors in an electrical network. MERL signal reconstruction and denoising algorithms enable 10-15 dB noise reduction with 5-20 times less computation than previous methods.

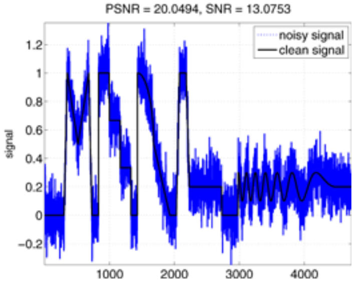

Clustering is one of the most important tasks in data mining and machine learning. Using our expertise in eigenvalues and iterative methods, MERL develops novel algorithms for spectral clustering, including high-accuracy/low-cost clustering of data from multiple sensors and modalities.

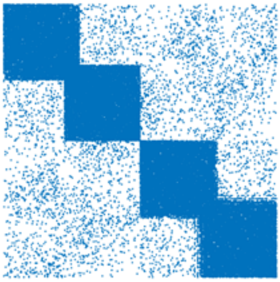

Machine Learning

MERL has a long history of research activities in machine learning. Our early work was on classical machine-learning algorithms and graphic-interference algorithms. The cascade Adaboost algorithm extended at MERL enabled the face detection that is now available in almost all consumer cameras and smart phones. MERL's generalized belief-propagation algorithm is critical for solving various inference problems in big data analysis. More recently, MERL developed a boosting algorithm for data lying on analytical manifolds. MERL has also contributed to several other machine-learning fields including conditional graphical models and dictionary learning.

Lately, MERL has focused on deep learning and developed novel algorithms for a wide range of applications including semantic scene labeling, activity recognition, object detection, image generation, predictive interfaces, and speech processing. We are also working on advanced Artificial Intelligence systems for industrial automation, smart transportation, and living environments.

Super-Resolution

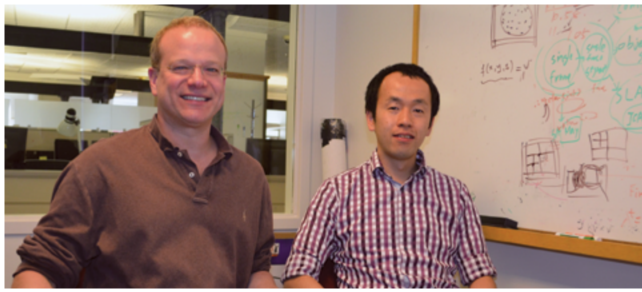

Super-resolution, which is the task of generating a high-resolution image from a low-resolution one, is a classic image-processing problem. Freeman's seminal paper on this topic [Freeman 2002] was an early demonstration of the potential of machine-learning techniques to address such problems. The paper inspired much follow-up work, including later research at MERL. Mitsubishi Electric sells a surveillance camera with super-resolution technology from MERL.

Most recently, MERL has put forward a new deep neural network architecture, called global-local network, which learns both holistic and local constraints for super-resolution of faces. This method provides state-of-the-art super-resolution accuracy, both visually and numerically, for magnification rates up to 8x.

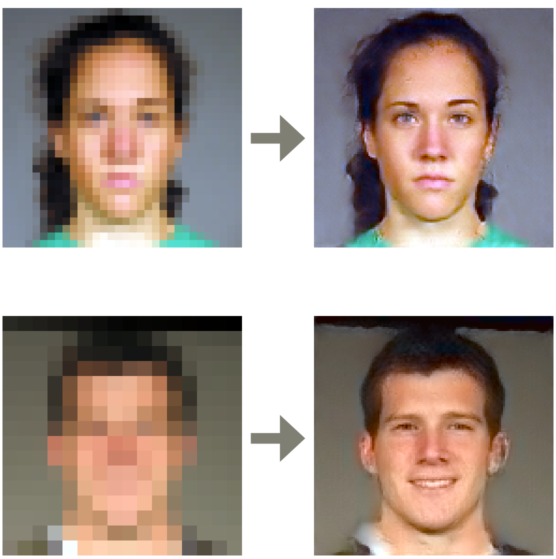

Generalized Belief Propagation

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.

MERL researchers pioneered the development of Generalized Belief Propagation algorithms [Yedidia 2000, 2003, 2005] 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.

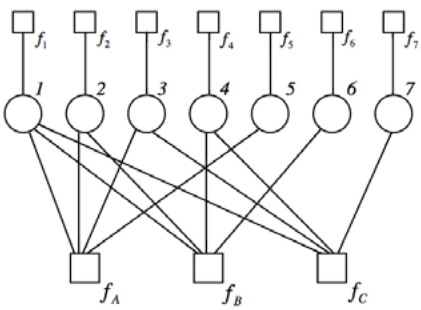

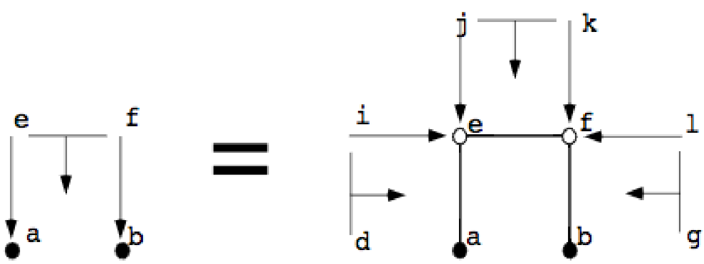

Computer Vision

Over its history, MERL has performed research in almost every area of computer vision. Our early work stemmed from computer graphics research and applied camera modeling to more realistic scene rendering. 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, precedence-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. The introduction of inexpensive 3D sensors has spawned many new MERL projects such as providing object position and pose for robotic grasping, and the creation of 3D models of the world for the purpose of navigation and planning.

For many years, we have applied machine learning to computer-vision problems [Freeman 2000] and developed new fundamental methods for detection and classification of people and objects. In recent years, we have developed new deep-learning algorithms for super-resolution, segmentation, noise reduction, and many variants of detection and classification.

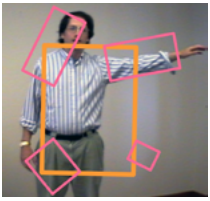

Object Detection

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.

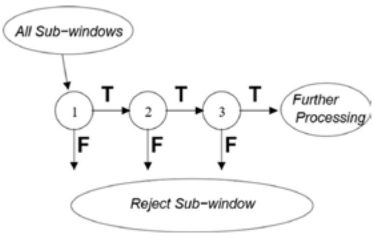

The Viola-Jones approach includes the concepts of Haar-like features, integral images, AdaBoost learning, and classifier cascades. These ideas combine to create a framework for building object detectors in images that is simple, powerful, and extremely fast, and is now a standard technique in computer 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 [Viola 2004].

At MERL these techniques were extended to face recognition. Further extensions combining motion features with image features succeeded in detecting pedestrians both more accurately and faster than prior approaches [Viola 2005, Venkataraman 2010].

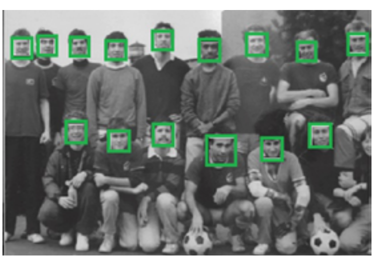

Computational Photography

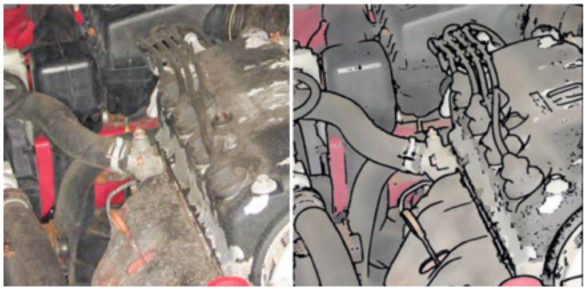

A traditional digital camera differs from a film camera only in the fact that the output is digital. Computational photography takes a step further — instead of just capturing a directly viewable image, it captures more complex data that may not be directly viewable, but with proper post-processing can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

MERL has been a leader in computational photography. In early work we created the "multi-flash camera," which takes multiple images using lights in different positions to make the extraction of depth edges easy. Our "flutter shutter" camera takes 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 takes images through programmable masks that make it possible to refocus an image during post-processing [Veeraraghavan 2007].

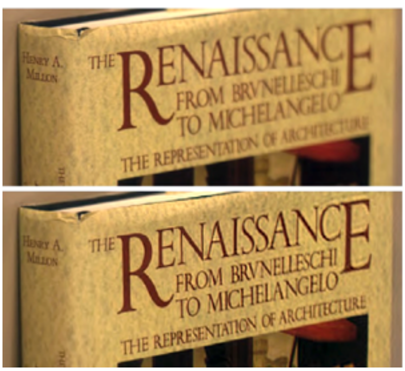

Today's cutting-edge commercial cameras make extensive use of computational photography principles to make photography more forgiving and precise.

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 10 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, which won an R&D 100 award in 2014.

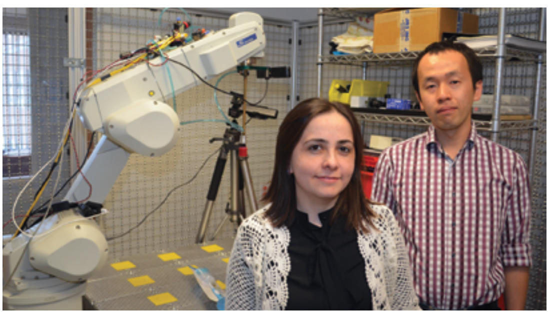

Heli-Tele

Unlike in the United States, street addresses in Japan are typically not sequential. GPS can help when one is in a car, but an alternative solution is needed for emergency management from the air. 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 accomplished registration of the live aerial video image 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.

This project was recognized with an R&D 100 award in 2006. In 2009, the Heli-Tele system was upgraded to high-definition cameras, but ground stations were only capable of handling lower-resolutions, so automatic tracking was added to crop the region of interest from the large high-definition image for transmission.

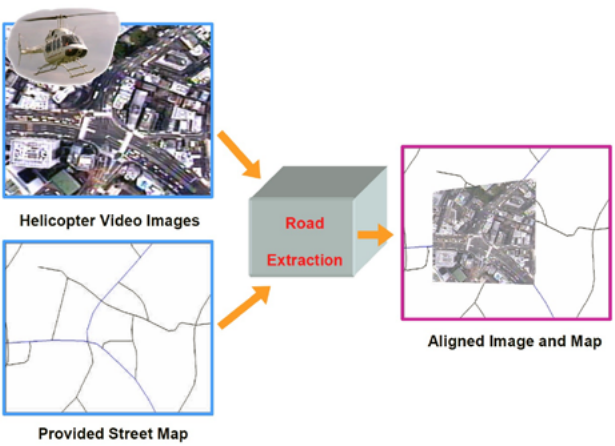

3D Reconstruction

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.

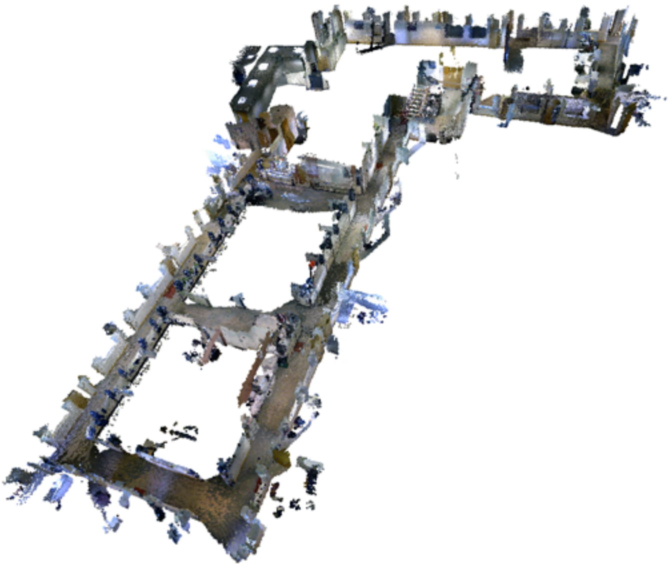

Covariance Features

In order track or locate something in an image, it is important to create a compact description of what it looks like. In 2005, MERL researchers developed “Covariance features,” a new region descriptor that combines powerful representation with efficient computation [Tuzel 2006]. It surpassed state-of-the-art accuracy when it was introduced and retained its leading status for many years afterword.

Over the years, this descriptor has been utilized for many computer-vision tasks such as image matching, object detection, tracking, semantic segmentation, and texture analysis. In follow-up research, it was used as the main building block of a new class of machine-learning algorithms for smooth manifold structured data, which was one of the seminal advances in this research area [Tuzel 2007].

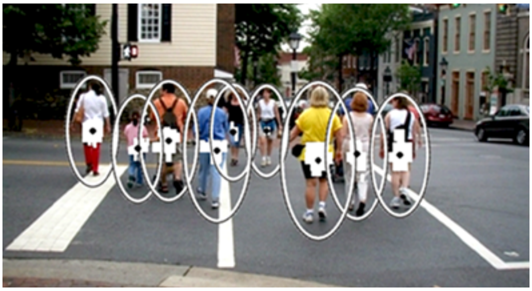

Autonomous Driving

The change from ordinary driving to autonomous driving promises to have as great an impact on society as the change from horses to cars. The requirements of true autonomy have spawned MERL research in several different areas.

The perceptual requirements include localization of the vehicle in the world, fusion of information from the many sensors on a vehicle, and understanding the 3D relationship of objects and motions in the scene.

The motion-planning requirements must support hundreds of different scenarios and make predictions and provisions for various possible actions on the part of the other agents. The control requirements mean that the plan must be dynamically feasible and that the vehicle faithfully executes the plan.

MERL researchers have been applying and extending techniques in the domains of computer vision, deep learning, compressive sensing, particle filtering, and model predictive control to address these challenges.

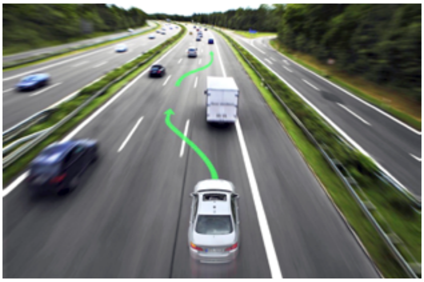

MERL Headquarters

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.

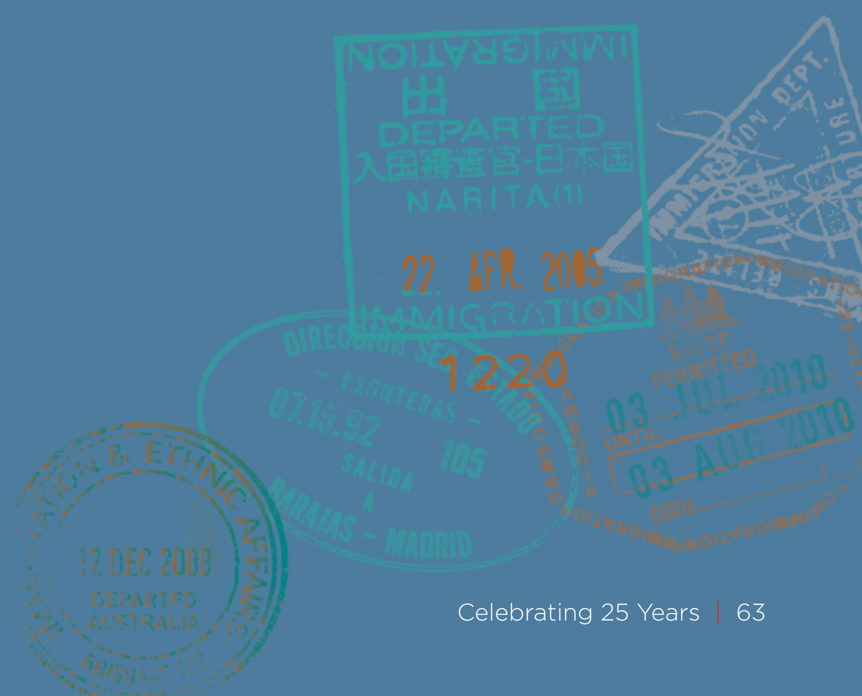

"MERL's philosophy of collaboration is impressively open and broad, which I believe is a significant contributor to its success."

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

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 US 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 Manager 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 two Patent Attorneys, a Patent Agent, 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 1,100 patents have been granted to MERL.

Liaisons

In order 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 Executive Vice President and a liaison staff of three, all four of whom are on loan from these labs. They play key roles connection MERL researchers to appropriate researchers in Japan and aiding communication between them and MERL.

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 US, 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

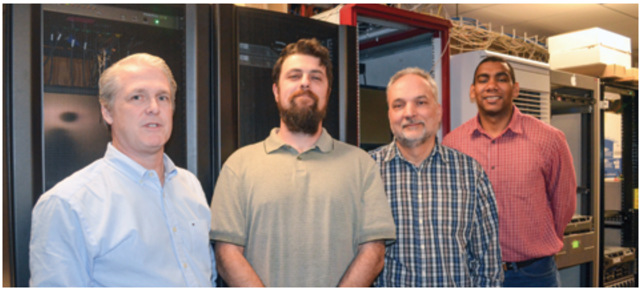

MERL's work is computer intensive. A team of four supports MERL's central infrastructure of servers and networking, as well as the individual machines our researchers use every day. All told, MERL has hundreds of computers, and hardware and software taken together, no two are the same.

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 simulating 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

The following pages list the MERL publications that have garnered the most recognition over time. Each has either won an award, or has been cited more than 100 times according to Google Schalar. Twelve of these publications have been cited more than 1,000 times.

“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, EE Department
University of Southern California
At MERL 2002-2008

Selected MERL Publications

[Avidan 2007] Avidan, S. and Shamir, A., 2007, August. Seam carving for content-aware image resizing. *ACM Transactions on graphics (TOG)* (Vol. 26, No. 3, p. 10). ACM. **(1,260 citations)**

[Brand 2002a] Brand, M., 2002. Charting a manifold. *Advances in neural information processing systems* (pp. 961-968). **(454 citations)**

[Brand 2002b] Brand, M., 2002. Incremental singular value decomposition of uncertain data with missing values. *Computer Vision—ECCV 2002* (pp. 707-720). Springer Berlin Heidelberg. **(375 citations)**

[Davenport 2010] Davenport, M.A., Boufounos, P.T., Wakin, M.B. and Baraniuk, R.G., 2010. Signal processing with compressive measurements. *Selected Topics in Signal Processing, IEEE Journal of, 4*(2), pp. 445-460. **(374 citations; 2015 IEEE Signal Processing Society Best Paper Award)**

[Dietz 2001] 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* (pp. 219-226). ACM. **(1,471 citations)**

[Efros 2001] 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* (pp. 341-346). ACM. **(1,923 citations)**

[Freeman 2000] 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. **(1,456 citations)**

[Freeman 2002] 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. **(1,826 citations)**

[Frisken 2000] 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. **(626 citations)**

[Gezici 2005] 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. **(1,595 citations)**

[Guo 2010] Guo, J.L. and Zhang, J.Y., 2010. Safety message transmission in vehicular communication networks. *17th ITS World Congress*. **(ITS 2010 Outstanding Paper Award)**

[Ivanov 2007] Ivanov, Y.A., Wren, C.R., Sorokin, A. and Kaur, I., 2007. Visualizing the history of living spaces. *Visualization and Computer Graphics, IEEE Transactions on, 13*(6), pp. 1153-1160. **(VIS 2007 Best Paper Award)**

[Liu 2013] Liu, D. and Boufounos, P.T., 2013, July. Synthetic aperture imaging using a randomly steered spotlight. *Geoscience and Remote Sensing Symposium (IGARSS), 2013 IEEE International* (pp. 919-922). IEEE. **(IGARSS 2014 Symposium Prize Paper Award)**

[Matusik 2004] 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. **(529 citations)**

[Molisch 2003] 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. **(862 citations)**

[Molisch 2007] Molisch, A.F., 2007. *Wireless communications*. John Wiley & Sons. **(1,713 citations)**

[Otsuka 2007] Otsuka, I., Suginohara, H., Kusunoki, Y. and Divakaran, A., 2007. Detection of music segment boundaries using audio-visual features for a personal video recorder. *Consumer Electronics, IEEE Transactions on, 53*(1), pp. 150-154. **(ICCE 2007 Best Paper Award)**

[Pekhteryev 2005] G Pekhteryev, G., Sahinoglu, Z., Orlik, P. and Bhatti, G., 2005, May. Image transmission over IEEE 802.15. 4 and ZigBee networks. *Circuits and Systems, 2005. ISCAS 2005. IEEE International Symposium on* (pp. 3539-3542). IEEE. **(142 citations)**

[Pfister 1999] 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. **(507 citations)**

[Pfister 2000] 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). **(961 citations)**

[Porikli 2011] Porikli, F. and Ozkan, H., 2011, August. Data driven frequency mapping for computationally scalable object detection. *Advanced Video and Signal-Based Surveillance (AVSS), 2011 8th IEEE International Conference on* (pp. 30-35). IEEE. **(AVSS 2011 Best Paper Award)**

[Raskar 2006] 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. **(417 citations)**

[Rich 1998] Rich, C. and Sidner, C.L., 1998. COLLAGEN: A collaboration manager for software interface agents. *User Modeling and User-Adapted Interaction*, 8(3-4), pp. 315-350. **(339 citations)**

[Shen 2004] Shen, C., Vernier, F.D., Forlines, C. and Ringel, M., 2004, April. DiamondSpin: an extensible toolkit for around-the-table interaction. *Proceedings of the SIGCHI conference on Human factors in computing systems* (pp. 167-174). ACM. **(401 citations)**

[Shibata 2010] Shibata, H., Kato, M., Kori, M. and Yerazunis, W., 2010. An automatic training data collection method for confidential email detection. *The Forum on Data Engineering and Information Management (DEIM)*. **(DEIM 2010 Best Paper Award)**

[Smaragdis 2003] 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. **(723 citations)**

[Thejaswi 2009] Thejaswi C., P.S., Zhang, J., Pun, M.O. and Poor, H.V., 2009, April. Distributed opportunistic scheduling with two-level channel probing. *INFOCOM 2009, IEEE* (pp. 1683-1691). IEEE. **(INFOCOM 2009 Best Paper Award)**

[Tuzel 2006] Tuzel, O.; Porikli, F.; Meer, P., "Region Covariance: A Fast Descriptor for Detection and Classification", *European Conference on Computer Vision (ECCV)*, May 2006. **(927 citations)**

[Tuzel 2007] Tuzel, O., Porikli, F. and Meer, P., 2007, June. Human detection via classification on riemannian manifolds. *Computer Vision and Pattern Recognition (CVPR), IEEE Conference on* (pp. 1-8). IEEE. **(460 citations, CVPR 2007 Best Paper Runner Up)**

[Vetro 2002] Vetro, A., Hata, T., Kuwahara, N., Kalva, H. and Sekiguchi, S.I., 2002. Complexity-quality analysis of transcoding architectures for reduced spatial resolution. *Consumer Electronics, IEEE Transactions on, 48*(3), pp. 515-521. **(2nd place 2002 Chester Sall Award)**

[Vetro 2003] Vetro, A., Christopoulos, C. and Sun, H., 2003. Video transcoding architectures and techniques: an overview. *Signal Processing Magazine, IEEE*, 20(2), pp. 18-29. **(768 citations)**

[Vetro 2011] Vetro, A., Wiegand, T. and Sullivan, G.J., 2011. Overview of the stereo and multiview video coding extensions of the H. 264/MPEG-4 AVC standard. *Proceedings of the IEEE*, 99(4), pp. 626-642. **(373 citations)**

[Veeraraghavan 2007] 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. **(441 citations)**

[Venkataraman 2010] Venkataraman, V. and Porikli, F., 2010, June. RelCom: Relational combinatorics features for rapid object detection. *Computer Vision and Pattern Recognition (CVPR) workshop on Object Tracking and Classification Beyond and in the visible Spectrum (OTCBVS), 2010 IEEE Computer Society Conference on* (pp. 23-30). IEEE. **(OTCBVS 2010 Best Paper Award)**

[Viola 2004] Viola, P. and Jones, M.J., 2004. Robust real-time face detection. *International Journal of Computer Vision,* 57(2), pp. 137-154. **(9,865 citations)**

[Viola 2005] 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. **(2,106 citations; early version won the 2004 Marr Prize: best paper at ICCV)**

[Wang 2011] Wang, B., Teo, K.H., Nishino, T., Yerazunis, W., Barnwell, J. and Zhang, J., 2011. Experiments on wireless power transfer with metamaterials. *Applied Physics Letters*, 98(25), p. 254101. **(109 citations)**

[Waters 1997] Waters, R.C., Anderson, D.B., Barrus, J.W., Brogan, D.C., Casey, M.A., McKeown, S.G., Nitta, T., Sterns, I.B. and Yerazunis, W.S., 1997. Diamond park and spline: Social virtual reality with 3d animation, spoken interaction, and runtime extendability. *Presence: Teleoperators and Virtual Environments*, 6(4), pp. 461-481. **(191 citations)**

[Xin 2008] Xin, J., Vetro, A. and Sekiguchi, S.I., 2008, January. A study of MPEG-2 to H. 264/AVC transcoding with half-horizontal resolution. *Consumer Electronics, (ICCE 2008) Digest of Technical Papers. International Conference on* (pp. 1-2). IEEE. **(ICCE 2008 Best Paper Award)**

[Xu 2001] Xu, P., Xie, L., Chang, S.F., Divakaran, A., Vetro, A. and Sun, H., 2001, August. Algorithms and system for segmentation and structure analysis in soccer video. *IEEE International Conference on Multimedia and Expo (ICME)* (Vol. 1, pp. 928-931). **(282 citations)**

[Yedidia 2000] Yedidia, J.S., Freeman, W.T. and Weiss, Y., 2000, December. Generalized belief propagation. *Neural Information Processing Systems (NIPS)* (Vol. 13, pp. 689-695). **(1,012 citations)**

[Yedidia 2003] 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. **(1,309 citations)**

[Yedidia 2005] 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. **(1,190 citations)**

[Yin 2002] Yin, P., Vetro, A., Liu, B. and Sun, H., 2002. Drift compensation for reduced spatial resolution transcoding. *Circuits and Systems for Video Technology (CSVT), IEEE Transactions on*, 12(11), pp. 1009-1020. **(115 citations, 2003 CSVT Transactions Best Paper Award)**

[Zhang 2009] Zhang, J., Orlik, P.V., Sahinoglu, Z., Molisch, A.F. and Kinney, P., 2009. UWB systems for wireless sensor networks. *Proceedings of the IEEE*, 97(2), pp. 313-331. **(217 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

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