KEN Project: Real-World Face Recognition

The silicon retina extracts object information from the unilluminated face area that cannot be gathered by the CCD camera. Left to right, compare the images taken in two-sided illumination by the CCD camera and silicon retina with the pair taken in one-sided illumination by the CCD camera and silicon retina.

THE Institute for Scientific Computing Research at the Laboratory recently developed the real-time, face-recognition system KEN. This system was developed through the integration of existing hardware and software. The name KEN is derived from the verb "to ken = to see, to recognize" and is related to the meaning of the noun in the Shakespeare quote "'tis double death to die in ken of shore."

Handwriting and speech recognition are already part of today's multimedia PCs and personal digital assistants. Object recognition, specifically face recognition, will be a natural extension of this technology. Associative query of image databases (e.g., search of medical image databases) and the interpretation of facial expressions for human-computer interfaces are likely applications. The existence of video-conferencing and the advent of interactive television widen the field of possible applications (e.g., model-based compression techniques, low-bit-rate video conferencing).

In security and surveillance, applications include ID verification at automatic teller machines (ATMs), credit card verification, preselection of crime suspects by matching sketches to a database, and identification of suspects for America's most wanted. The following scenarios are thinkable: An automatic recognition system preselects a few pictures of suspects from a crime scene out of a large database or alerts security personnel if suspicious activity occurs around an ATM or a property. Recognition systems included in valuable properties, such as cars or boats, could be used to identify the owner and prevent unauthorized access. This recognition technology is not limited to faces but can also be applied to other object classes, such as signatures and footprints.

Recognition Process

KEN demonstrates object recognition technology by recognizing human faces in a real-world environment. The demonstration system contains several modular components: The input sensor is either an off-the-shelf charge-coupled device (CCD) camera or an analog silicon retina chip (from Caltech). The silicon retina normalizes contrast locally and offers distinct advantages under strongly varying lighting conditions by extending the dynamic range beyond that of a normal CCD camera. A computer-controlled lens system with autozoom, autoiris, and autofocus and a pan-tilt unit track the possible candidates. A datacube real-time video processing board extracts features and segments the input images. The control software for mechanical components and the classification software run on a workstation.

KEN's software represents faces as labeled graphs. The vertices of these graphs are labeled by feature vectors and their links by distance vectors. To "acquaint" the system with a person during the learning phase, the person's face is extracted as a frontal view and stored as a face graph model in the system's knowledge database. The knowledge base contains all memorized faces in the form of sparse graphs. During recognition, all face models are compared to the unknown face candidates appearing

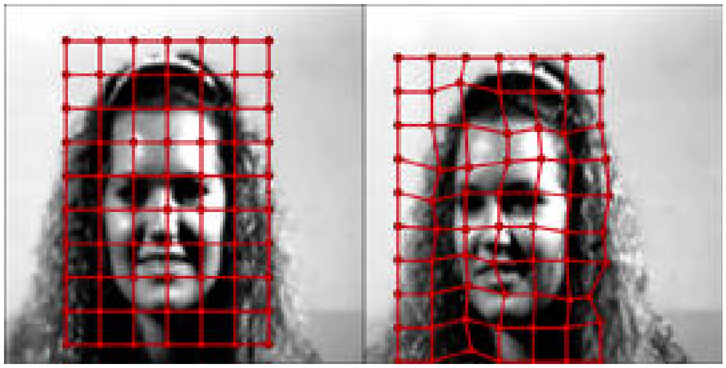

Rubber sheet effect: The labeled graph model (overlay) for the face in the left image is distorted to fit best over the face in the right image. The quality of the fit describes how closely the two face images resemble each other.

in the input. The face models are positioned over a face candidate and normalized in size and orientation by a template match. Subsequently, the graphs are locally distorted by an elastic matching process to fit the input image as closely as possible. The distortion effect is similar to the stretching of an elastic membrane (e.g., a rubber sheet). Cost values are assigned to the distorted face models according to the proximity of the final match to the input. A statistical evaluation of the resulting cost values of all models allows the system either to qualify a match or to reject poor matches and ambiguous results. KEN is currently capable of classifying a face, almost in real time, from a database of 100 faces with a positive identification rate of up to 90% and less than a 1% false positive recognition rate.

Environment

The hardware for KEN is off the shelf with the exception of the silicon retina prototype. The software is called AVision. It is written as a collection of object libraries in C and C++ and runs on major workstations (Sun, SGI, PC) and supercomputers, depending on the availability of C++ compilers. An upcoming parallel implementation in SISAL will help to speed up the development of learning algorithms that require the performance of supercomputers.

Future Directions

We plan an extension of this work to motion sequences of objects in a scene, such as the parts of a face in facial expressions or the body parts of a moving person. This involves constructing a choreographic database to implement a motion recognition system, because motion, such as the distinct trajectories of moving limbs, is used by humans as an additional visual clue to recognition. Motion processing appears to be an important component in a real-world application of recognition technology. Other research is aimed at organizing large databases to improve the efficiency of the elastic matching process. Clustering processes, active data acquisition, and learning algorithms for the weighting of components are under development. The associative query of pictorial databases (e.g., of medical images) is a similar problem that may also be researched.

Acknowledgments

We are grateful to Kawabena Boahen and Carver Mead, Caltech, for providing the prototype of a silicon retina. Joachim Buhmann, Bonn, Germany, contributed active vision concepts, learning algorithms, clustering schemes, and programming effort to the project.

For more information contact Martin Lades (510) 422-0893, hml@llnl.gov.

Modeling Groundwater Flow and Chemical Migration

GROUNDWATER contamination is a major environmental problem throughout the world. In the United States, for instance, numerous governmental and industrial sites require remediation. The Department of Energy is currently cleaning up several of its contaminated facilities, including LLNL.

At LLNL, chemical solvents and petroleum products were dumped onto the ground surface in the 1940s when the present site was a naval air station. Forty years of research and development activities also contributed to the contamination problem. LLNL is obligated to characterize the contamination and clean it up. To this end, various engineered remediation techniques (e.g., pump-and-treat and biofilters) are now being studied, tested, and implemented.

To understand better the efficacy of a given cleanup strategy, as well as to determine the most economical implementation for a specific site, engineers frequently employ mathematical modeling tools to aid in their analysis and design of remediation procedures. Unfortunately, many models are based on unrealistic assumptions about the subsurface media and flow behavior. For example, many models in use today ignore the fact that the subsurface is heterogeneous in composition and in spatial distribution. These simple homogeneous models are incapable of capturing important physical phenomena (such as nonuniform diffusion) that have a significant impact on contaminant migration.

Another flaw of many models is that they employ outdated and inefficient numerical methods, which preclude running realistic simulations on even the largest conventional vector supercomputers.

To address these deficiencies, we are developing ParFlow, a simulator for modeling fluid flow through heterogeneous porous media. To enable detailed simulations of large sites, this code uses the latest numerical methods and high-performance computing technologies.

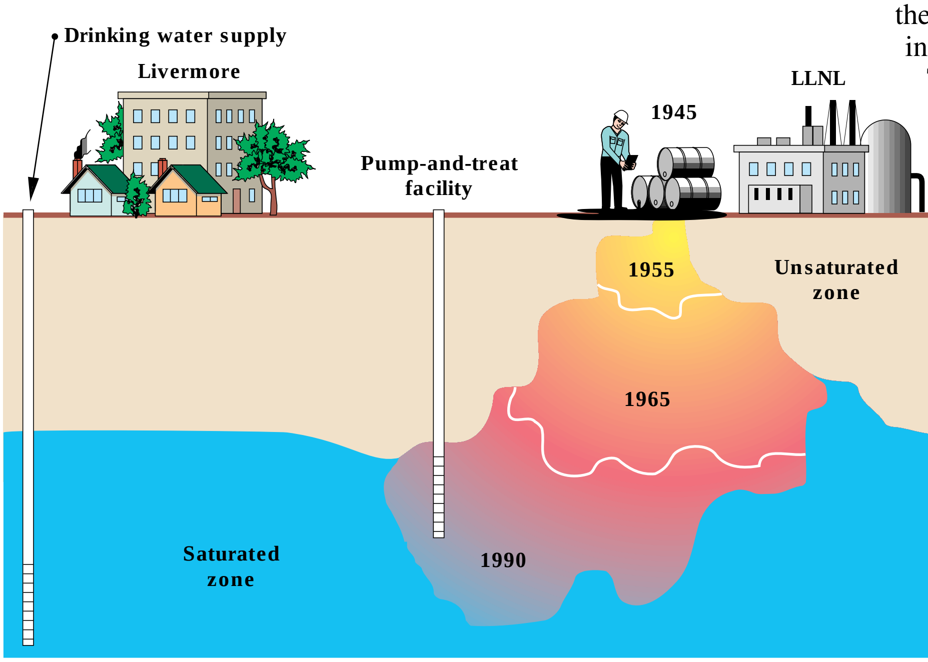

Contaminants have migrated through the unsaturated zone into the more mobile groundwaters. LLNL is designing and executing remediation procedures.

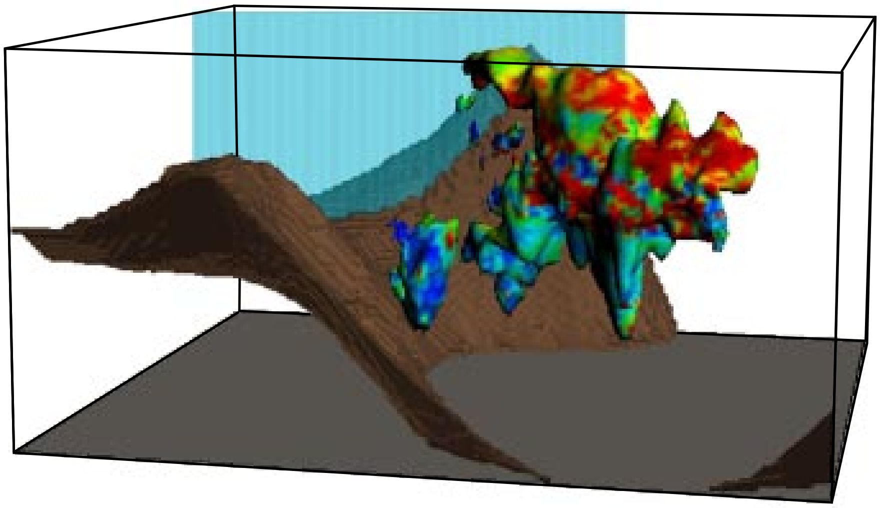

A snapshot in time of contaminant migration through the LLNL subsurface (hypothetical, multilayered, heterogeneous realization). A clay layer is shown in brown. The blue plane represents a fault zone. (The vertical scale is exaggerated 10 times.)

Modeling Flow and Contaminant Migration on MPPs

ParFlow is a portable and parallel code for modeling multiphase fluid flow and multicomponent chemical transport through heterogeneous porous media. This scalable code presently runs on the CRAY T3D, IBM SP-1, and nCUBE/2 massively parallel computers, as well as on a two-processor SGI Onyx and a cluster of Sun Sparcstations. It was used recently to simulate groundwater flow at LLNL (using a hypothetical subsurface model). The size of the physical site (several square kilometers) and the need to resolve the subsurface heterogeneities (to within a few meters) resulted in a grid with more than one million spatial zones.

To obtain the necessary hydraulic conductivity values, we implemented an algorithm that generates a statistically accurate subsurface realization from the given field data. We then discretized the modeling equations and solved the resulting linear system with preconditioned conjugate gradients. The velocity field was then passed to a particle-in-cell code to simulate contaminant migration.

It is important to note that the subsurface heterogeneities give rise to preferential flow channels. These channels, which are not reproducible by homogeneous codes, lead to nonuniform contaminant migration and have a major impact on the efficacy of engineered remediation procedures.

Graphical User Interface

To facilitate code development and maintenance, we are developing XParFlow, an X-Windows (Motif) interface to ParFlow. This program displays the current simulator configuration as a connected graph. By selecting a node in the graph, the user is able to customize the simulator. For example, the user can choose from several linear solvers and preconditioners.

To facilitate problem specification, we are currently using the GMS package from Brigham Young University and the Department of Defense. This new graphical user interface was specifically designed to be a front-end to groundwater modeling codes such as ParFlow.

Multidisciplinary Collaboration

The ParFlow project is an interdisciplinary effort involving scientists from various programs within the Laboratory. LLNL and International Technology Corporation are partnering to commercialize some of this work under the auspices of a cooperative research and development agreement (CRADA).

For further information contact Steven Ashby (510) 423-2462, sfashby@llnl.gov.

Gas and Oil National Information Infrastructure

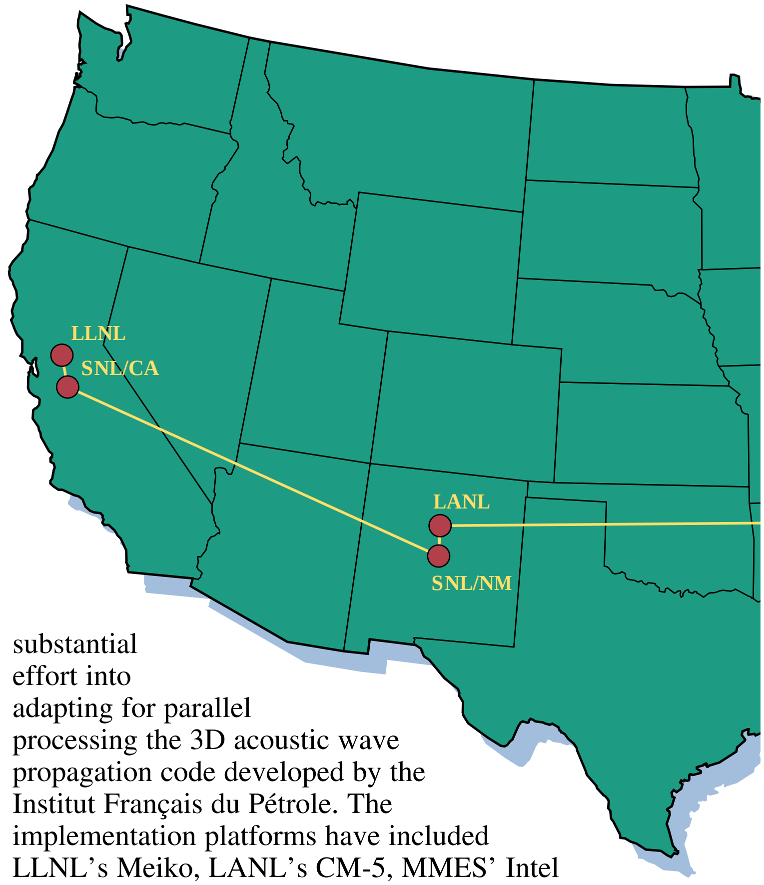

THE U.S. Department of Energy pilot National Information Infrastructure (NII) programs are designed to create computing and communications testbeds with U.S. industries. One such program is the Gas and Oil National Information Infrastructure (GO-NII), which unites four of DOE's national laboratories: Lawrence Livermore National Laboratory (LLNL), Los Alamos National Laboratory (LANL), Martin Marietta Energy Systems (MMES) at Oak Ridge, and Sandia National Laboratories (SNL). GO-NII focuses on networking, high-speed input and output, mass storage, collaborative and productivity tools, scientific data management, data visualization, data security, and distributed computing. Industry can evaluate the feasibility of using new advanced technologies in the GO-NII program while reducing the financial investment and risk of supporting its own research program. Technology and information will be transferred from the national laboratories to all segments of the gas and oil industry.

GO-NII Demonstration

At the Supercomputing '94 conference in Washington, D.C., GO-NII demonstrated the use of an advanced high-speed, long-haul communication system; a new architecture in high-performance storage systems; calculating, searching, browsing, and cataloging of three-dimensional (3D) seismic data and metadata; and productivity tools to enhance remote collaboration among scientists. Although this demonstration focused on the gas and oil industry, the technology can be applied to other U.S. industries.

Synthetic Seismic Dataset

A major undertaking of the GO-NII effort is a project with the Society of Exploration Geophysicists and the European Association of Exploration Geophysicists to design salt and overthrust 3D models and then simulate realistic 3D surveys based on those models. The four GO-NII laboratories are working jointly to generate this 3D synthetic seismic dataset for the verification and validation of seismic processing tools used in the gas and oil industry. These national laboratories have put substantial effort into adapting for parallel processing the 3D acoustic wave propagation code developed by the Institut Français du Pétrole. The implementation platforms have included LLNL's Meiko, LANL's CM-5, MMES' Intel 860 and Paragon, and SNL's Intel Paragon. The laboratories will be applying massively parallel processing at the equivalent of approximately 7000 hours at a billion calculations per second (Gflops) toward this calculation between October 1994 and September 1995.

The salt model and overthrust model synthetic seismic datasets calculated at the four laboratories will be stored at LLNL's Facility for Advanced Scalable Computing Technology (FAST). FAST will provide information and expertise in emerging storage technology, as applied to the use and distribution of these massive datasets. FAST will provide a base on which to build distributed visualization, remote access, and collaborative tools to enable more efficient analysis and processing of seismic information.

At the end of the project, the resulting model, tools, and data will be made publicly available to companies within the gas and oil industry for further commercialization and research.

ARIES Testbed Partnership

The Amoco ARIES testbed, as described in *Network World* and other trade journals, is a partnership between Amoco and about 15 industrial communications and computing vendors to explore the use of asynchronous transfer mode networks in the oil and gas business sector. The four GO-NII laboratories offer their capabilities in network technology, distributed computing, and system integration to partner with the ARIES testbed in

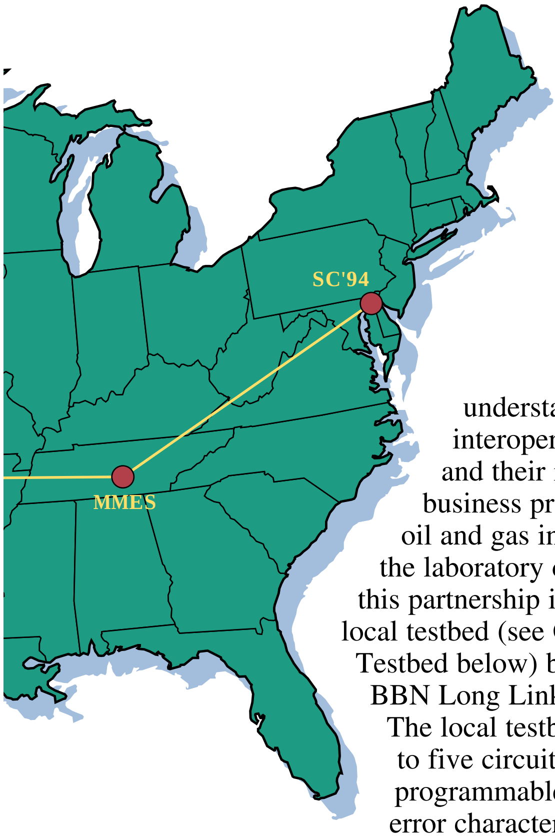

At the IEEE Supercomputing '94 Conference, DOE national laboratory researchers demonstrated long-haul, high-speed communications, high-performance storage, and collaborative tools, all of which can be used in gas and oil industry applications.

understanding interoperability issues and their impact on business processes in the oil and gas industry. Part of the laboratory contribution to this partnership is to establish a local testbed (see GO-NII Testbed below) based on the BBN Long Links Emulator. The local testbed provides up to five circuits of programmable delay and error characteristics that will be used to emulate nationwide networks within a more controlled testing environment, eliminating the need for long-haul communication lines in the early stages of testing and development.

GO-NII Testbed

The GO-NII testbed will provide an environment through which the gas and oil industry can evaluate, benchmark, and access new computing infrastructure technologies of interest to them. The GO-NII testbed will simulate long-haul, high-performance networks with access to supercomputers, workstation clusters, and high-performance mass storage in a transparent distributed computing environment. These include the National Storage Laboratory and High-Performance Storage System environment, massively parallel systems, emerging technologies in parallel input/output, cluster environments, and dynamic software to manage the environment and its resources. The GO-NII testbed will provide integrated access tools to effect the desired seamless environment.

Petroleum Technology Transfer Council

A GO-NII focus area is to work with independent oil and gas producers in making use of advanced computing technologies to increase domestic oil production. Independent of GO-NII, the Petroleum Technology Transfer Council (PTTC) was formed in January 1993 as an *ad hoc* council to address the technology needs of the U.S. oil- and gas-producing community and to identify the best mechanisms for improving the transfer and communication of technology to domestic producers. Ten regional lead organizations (RLOs) representing the oil-producing areas of the United States have been identified, and regional resource centers will be established at these sites. A producer advisory group in each region will help to define the priority technical problems encountered in the region and to determine the needs of the producers in accessing available technologies to resolve them.

The four national laboratories participating in GO-NII are working with the PTTC RLOs to provide expertise, advice, and technology transfer possibilities and thus to help in forming the electronic information system linking the ten RLOs and the approximately 8000 independent oil and gas producers. Potential areas in which the laboratories and RLOs will team up are computing infrastructure for electronic commerce, full spectrum network access, network security, distributed visualization, data integrity and conversion, system integration, electronic storage, compression, and database technologies.

ACTS Testbed

Amoco, the American Petroleum Institute, Naval Research Laboratory, DOE laboratories, and NASA are planning a joint demonstration of satellite communications capabilities as applied to the gas and oil industry. NASA's Advanced Communications Technology Satellite is operating multiple T1 (1.5-megabit-per-second) circuits, with the capability of OC-3 (155-Mb/s) and OC-12 (622-Mb/s) speeds. This capability could have a dramatic impact on the ability of the gas and oil industry to retrieve seismic and reservoir data from remote sites, such as the Gulf of Mexico, Alaska, and foreign locations. A low-speed demonstration of this capability is scheduled for December 1994.

For further information contact Carol Hunter (510) 422-1657, chunter@llnl.gov.

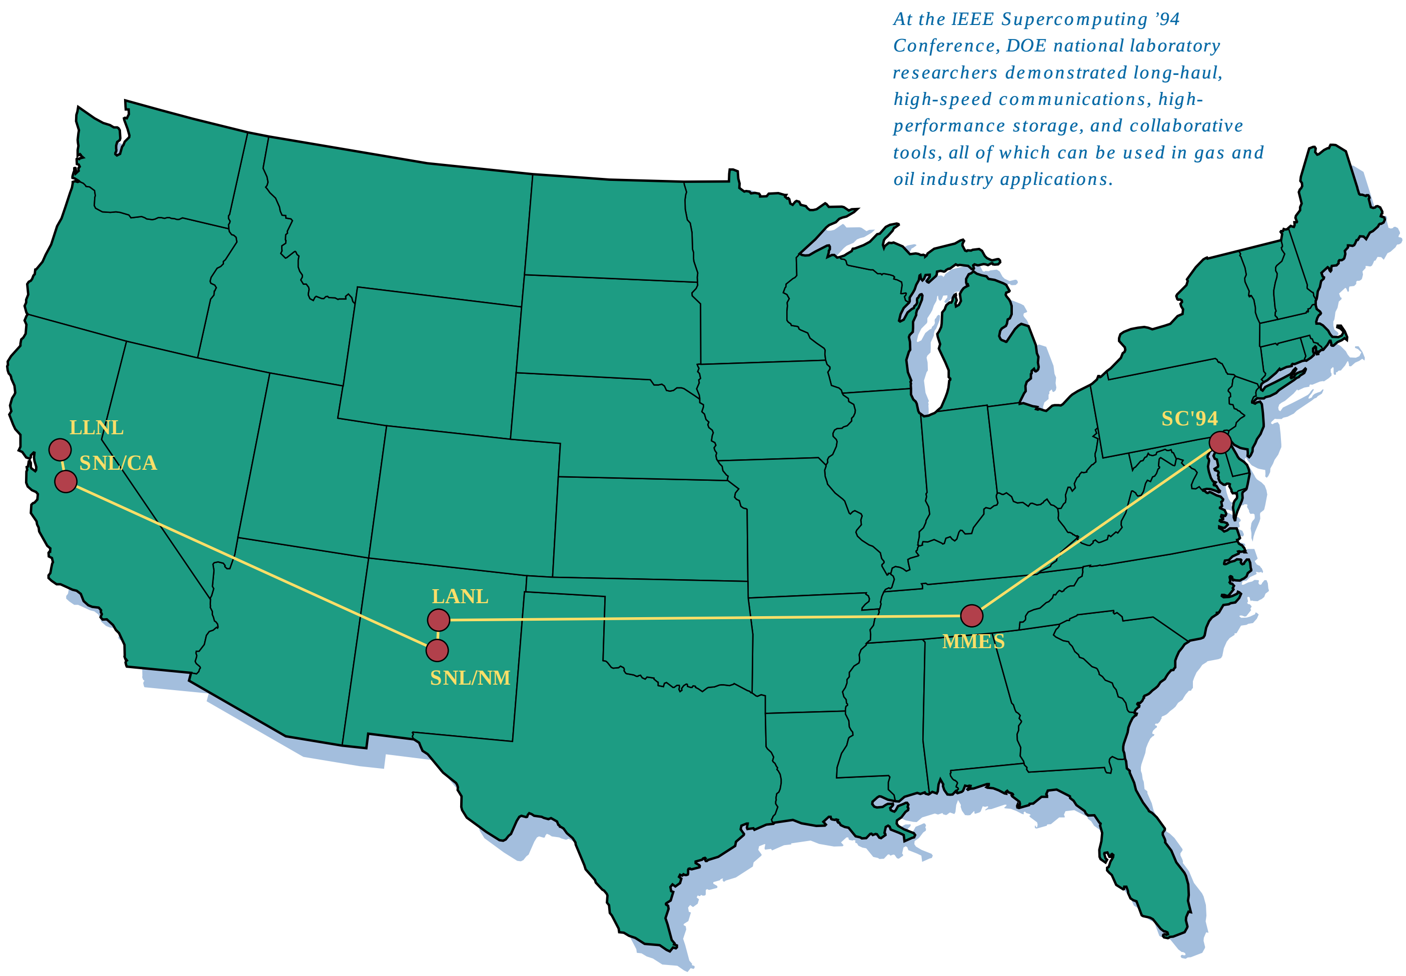

Figure for Gas and Oil National Information Infrastructure Highlight