Center for Advanced Knowledge Enablement (CAKE)

A CISE-funded Center

Florida International University, Naphtali Rishe, Director, 305.348.2025, rishe@fiu.edu

Florida Atlantic University, Borko Furht, 561.297.3180, borko@cse.fau.edu

Dubna International University (International Site)

Center website: http://cake.fiu.edu

Autonomous Lighter Than Air (ALTA): Sub-centimeter Aerial Photography

Lower than a plane, higher than any street view, Autonomous Lighter Than Air (ALTA) platforms provide a brand new vantage point. ALTA is a smart balloon that transmits images and environmental information from distant locations. ALTA models may be tethered to continuously monitor specific patches of land or sail on air currents along predicted wind-driven paths.

Autonomous Lighter Than Air (ALTA) Platform and its view of Washington DC.

ALTA's remote-controlled, smart balloons transmit to the Internet geo-referenced super-resolution, oblique imagery, as well as air quality data. The US patented, Federal Aviation Administration (FAA) compliant technology operates from otherwise unreachable places, higher than ground cameras and lower than most planes. Its low cost makes frequent updates cost effective. The ALTA balloons' (R series) ability to operate at low altitudes makes possible aerial photography with unprecedented detail. ALTA's image clarity is not hampered by propulsion vibration or cloud cover. Because of these advantages, ALTA photography often yields higher fidelity images than are produced by other geo-referenced imagery.

Atmospheric graphical information can be attached to the images to depict air quality and other atmospheric data. The ALTA imaging technology http://altadrifter.com is being combined with TerraFly at http://TerraFly.com, see Figure , "TerraFly Maps Enable Monitoring of Airborne Cameras," on page 31. As such, it will be provided as a service via http://TerALTA.com. The ALTA team services public safety, news agencies, agriculture, construction, real estate, travel and tourism. A demonstration prototype is at http://cake.fiu.edu/ALTA.

Economic Impact: ALTA opens a new multi-billion market for aerial photography as it produces images that are orders of magnitude higher resolution than the current state of the art. Image collection is currently accomplished from ground-based cameras, aircraft or satellites. Because ALTA is higher than ground-based cameras, it sees more. ALTA is lower than an aircraft and therefore sees better. Additionally, compared to other aerial platforms, ALTA has low cost components. For the capital outlay of one manned aircraft, 300 ALTAs could be purchased; balloons are also dramatically cheaper than drones. It costs $100 to build one ALTA. It costs many thousands dollars to build other aerial platforms. Operations costs also are a fraction of what other aerial collections cost. ALTA missions eliminate cost of pilots, aircraft, and airports.

In addition, ALTA can be deployed in minutes and have information and images returned instantaneously. The ALTA technology is thus poised to produce much higher-quality imagery at much lower cost than current technologies, thus opening up new markets and bringing new capabilities to existing markets such as public safety, real estate, construction, environmental monitoring, disaster mitigation, and disaster recovery. By economically providing virtual "see for yourself" access, ALTA imagery will capabilities to civic and real estate land data and imagery, which are already multi billion dollar markets. In addition, it is expected that retail markets of socially networked users will approach hundreds of billions of dollars. It can also be used in public safety, homeland security, government and military arenas wherein ALTA can replace UAVs and drones with lower cost and longer, more stable flights, the potential markets are estimated to be in excess of $500 million.

For more information, contact Naphtali Rishe, 305.672.6471, rishe@fiu.edu.

Automated Asset Management in Data Centers

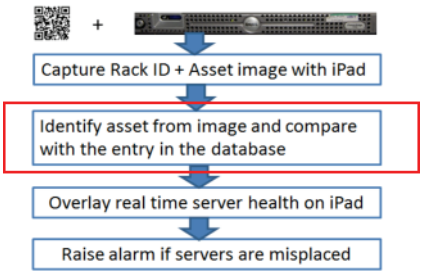

Visual identification of computing assets for efficient data center operations

CAKE researchers at Florida Atlantic University developed an innovative solution for visual asset identification using visual features of an image. Visual features of asset images are computed using complex mathematical methods. These visual features are used to identify and match asset images. A database with visual features of asset images was built for every distinct asset that is typically present in large data centers.

A data center is a facility that hosts computer systems, servers, power supplies, storage systems, and other related computing equipment, referred to as assets. The size and number of these data centers are continuously increasing to accommodate the need and demand for web based applications and services. Assets are mounted in racks and a typical rack can accommodate up to 42 assets depending on the asset size. Large data centers have thousands of racks and keeping track of these large numbers of assets manually makes it very tedious and highly prone to errors.

Human errors continue to be the greatest cause of unplanned downtime in data centers. Downtime of assets in data centers lead to slow or unavailable information services on the Internet. Solutions that minimize human input in asset management will lead to higher productivity and reduced downtimes.

Portable devices such as tablets and mobile phones are ideal devices to perform asset management operations in data centers. Information technology (IT) personnel can effortlessly carry these devices in data centers to conduct management operations. Such devices have become computationally powerful and are equipped with cameras and other sensors. Cameras on these devices provide a unique opportunity to simplify asset monitoring in a data center. Cameras on mobile devices can be used to visually recognize the assets in a rack and provide real-time information about the operating health of the assets. With a camera-based solution, IT personnel have to just point the camera at a rack and select the device to monitor. Any mismatch between the asset identified in the rack and the asset that was expected is immediately flagged. Additionally, the health of the asset is instantaneously displayed on the mobile device without having to login to the asset.

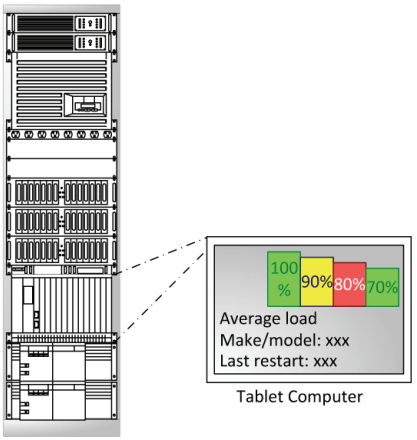

Assets in data center racks are identified using the camera on a tablet and real-time asset information is displayed on tablet.

Assets needing identification are captured using a camera on a mobile device. The device then extracts and transmits the visual features to the server for matching and asset information retrieval. This breakthrough, an optimized version of visual feature extraction and comparison methods, was developed to improve matching accuracy and reduce computational complexity of feature extraction as well as matching. This innovation introduced methods to prioritize and reduce the number of visual features used to identify and match asset images. This reduction in complexity enables efficient asset management solution on mobile devices.

This work represents an improvement over previous state-of-the-art technology because it introduces simplified asset management tools based on visual features of assets. This innovative asset management system allows IT personnel to assess the state of computing assets by just pointing a mobile device camera at the asset.

Economic Impact: The advantage of this process is that it enables immediate identification of problematic assets using real-time operational data from the assets without having to explicitly and manually logging into the asset management system. This leads to reductions in data centers' operational costs by using relatively inexpensive portable devices, such as mobile phones and tablets to minimize human errors, while improving productivity and reducing downtime. According to Emerson's Networks white paper, the average cost of a single data center downtime event was approximately $550,500 or $5,600 per minute; or one third of the total cost when indirect and opportunity cost was taking into consideration. This illustrates the importance of this project and the potential impact on this industry sector. This new approach to asset management using visual asset identification methods and mobile devices has the potential to significantly reduce the time spent on identifying problems in data centers. This should lead to improved uptime of servers and computing assets and thus increase the profitability of information service providers. Improved uptime of computing assets has direct impacts on the revenue generated in the Internet economy. The economic impact could total multibillions of dollars. This technology is being deployed in commercial products. Doing so requires minimal research and

development. In addition to favorable economics, implementation of the technology will relieve pressure on data center management thereby reducing chances of consumer dissatisfaction over delayed or unavailable services and minimize any economic impact on the service providers.

For more information, contact Hari Kalva, 561.297.0511, hari.kalva@fau.edu.

System for Reducing Hospital Readmissions

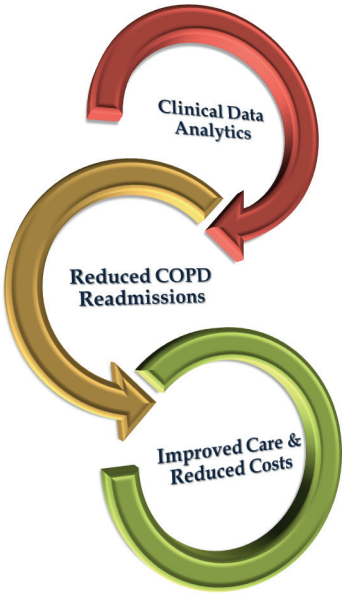

Reducing COPD readmission costs through analytics

In the United States., hospitalization is one of the largest factors impacting healthcare expenditures. Hospital readmissions are one of the indicators of the quality of care. Hospital readmission rates have direct monetary relationships with patients' financial burden and the government/private medical care organizations such as Medicare, Medicaid, and insurance companies.

There is enough evidence to conclude that almost 70% of readmissions are related to four diseases. These are Congestive Heart Failure (CHF), Chronic Obstructive Pulmonary Disease (COPD), Diabetes (DM), and Syncope (fainting or passing out). Deploying technologies to avoid such readmission should not only improve the quality of patient care but should also help offload a great deal of mounting financial burden.

Researchers at CAKE, in partnership with Soren Technology, developed a system for reducing hospital readmissions. The system integrates several telemedicine, patient care coordination, and clinical decision support systems to identify patients at high risk for re-admission. This is all based on data mining and a statistical analysis engine. The current system focuses on the readmission issues related to COPD.

Predictive modeling of readmission is a complex effort. This work represents an improvement over previous state of the art because it enables comprehensive autonomous statistical analyses based on the mining of patient data using a unique process/algorithm. The clinical decision support system developed at CAKE is designed to predict hospital readmission risk for COPD using electronic health records (EHR) information. The COPD clinical decision support system is based on predictive analytics using structured and unstructured patient data to develop a readmission risk profile for a patient being discharged after an initial COPD related hospital admission. The data in most hospital systems are unstructured, such as in physician's notes, patient discharge summaries, and X-ray radiology reports. In this new approach, information is extracted using a breakthrough clinical natural language processing system. This system automatically extracts useful information. It identifies in patient records the use of clinical named entities such as diseases and disorders, signs and symptoms, anatomical sites/ and procedures, and drugs. This information is combined with structured data from EHRs that include relevant laboratory test results and appropriate non-clinical patient demographic data. The system then uses statistical models to predict the patient's readmission risk. This CAKE system enhances prediction reliability by integrating unstructured

and structured data. Since readmission rates have remained nearly constant for the past few years, such efforts are not just timely, but they are also critical in helping improve quality of service while reducing associated healthcare costs.

> **Economic Impact**: A study reported in 2009 that 19.6% of Medicare fee-for-service beneficiaries who had been discharged from a hospital were readmitted to the hospital within 30 days, 34.0% within 90 days, and more than half (56.1%) within one year of discharge. MedPAC also reported that readmissions, within 30 days accounted for $15 billion of Medicare spending. Medicare is the payer for about half of these readmissions. Current data shows that COPD accounts for about 22% of readmissions, hence taking this into account together with the 76% of preventable readmissions yields a 16.7% of potential readmissions reduction or potentially over $4 billion of annual savings. This new decision support system is well positioned to help quantify these savings. Hospitals and clinics can integrate the developed system with their current medical information systems to leverage clinical data and provide meaningful clinical decision support. This should result in improved care and reduced hospital readmissions. The system can be further expanded to include other diseases such as Congestive Heart Failure (CHF) that would further reduce the overall hospital readmissions and thereby provides much higher savings. While this technology will have significant economic impacts through cost savings associated with reduction in COPD hospital readmissions, the greater potential economic impact could be realized as the solution is expanded to include other diseases and chronic medical conditions.

For more information contact Ankur Agarwal, 561.297.3496, ankur@cse.fau.edu or Ravi Behara, 561.297.2778, rbehara@fau.edu.

TerraFly Maps Enable Monitoring of Airborne Cameras

TerraFly is a technology and tool for the visualization and querying of geospatial data. It provides users with the experience of virtual "flight" over maps comprised of aerial and satellite imagery overlaid with geo-referenced data. The data drilling and querying component of the system allows the users to easily explore geospatial data, create geospatial queries, and get instant answers supported by high-performance multidimensional search mechanisms. TerraFly's server farm ingests, geo-locates, cleanses, mosaics, and cross-references 40TB of basemap data and user-specific data streams. The interface allows rapid deployment of interactive Web applications. It is accessible from anywhere via any standard Web browser, with no client software to install.

Although video surveillance recording is the state of the practice, the video collected is normally used only after the fact - it cannot easily be accessed in real time, does not have accurate geolocation capabilities, and cannot be easily integrated with other forms of critical information. This state-of-the-practice lack of situational awareness will be overcome by the Context Aware Rich Media Extensible Middleware technology (CARMEL) TerraFly system. This system integrates cutting-edge CARMEL technology from IBM Research Haifa (http://www.haifa.ibm.com) with the TerraFly Geospatial System at CAKE. The CARMEL-TerraFly system provides geographically anchored streaming services, delivered via IBM's state-of-the-art technology and accessible via TerraFly's intuitive spatio-temporal interface. This integrated system offers innovative situational awareness technology, while helping expand the Center's international influence and connections. By combining IBM Haifa's Geographic Information Systems (GIS) and streaming technology research, CARMEL is a geographically anchored, video-on-demand streaming infrastructure that provides: 1) scalable, end-to-end low-delay and resilient streaming technologies; 2) on-demand bandwidth

adaptation (transcoding); 3) highly accurate geographical searches; 4) real-time, geo-located notification; and, 5) high performance, service-oriented, architecture-enabled technologies.

The novel CARMEL-TerraFly technology is transforming public safety assurance systems. It is also making possible more timely responses to situations by providing geographically anchored streaming services. These can be combined with and accessed via the intuitive TerraFly user interface. Users are able to select a geographic area of interest, retrieve multimedia data from sensors in the area, and view streaming video of moving objects in real time (e.g., vehicles, people, animals, etc.). Users are also able to set temporal and geographic constraints to view the path traversed by a specific moving object or group of objects.

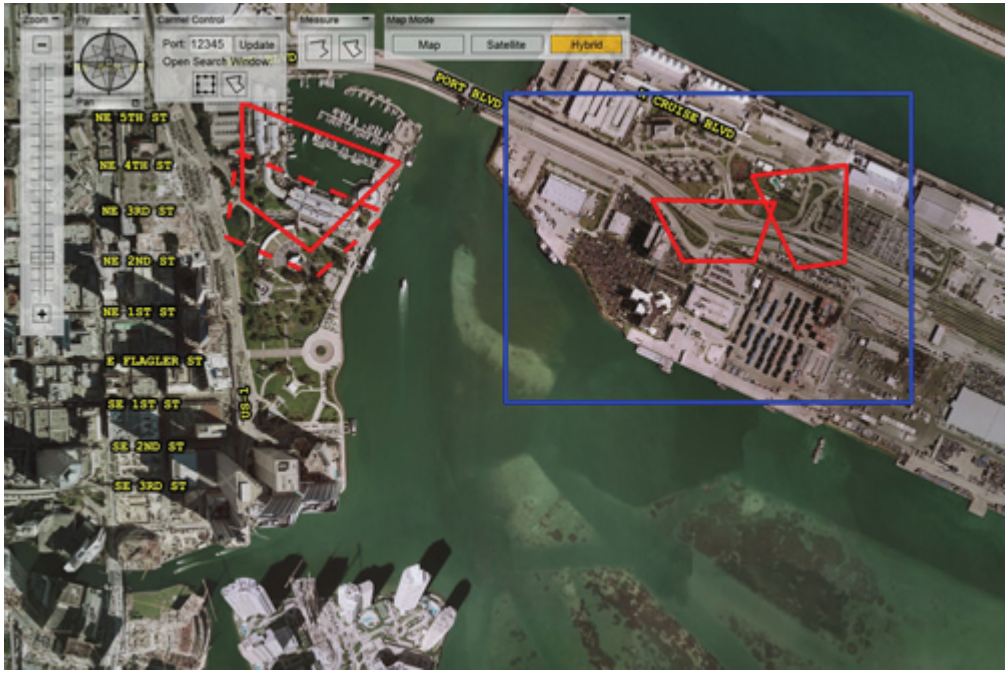

CARMEL-TerraFly user interface. The map shows the Port of Miami with moving traces of areas videotaped by airborne cameras. Solid trapezoids are ground projections synchronized with playback and dotted trapezoids are real-time projections of camera views. The blue rectangle allows selection of video fragments at times and locations of interest.

There are numerous potential applications for this advanced technology, particularly for command and control operations such as homeland security, law enforcement, and disaster response. For example, using the CARMEL-TerraFly system, law enforcement could be alerted to a situation such as a hit-and-run accident. Officers would be able to quickly pin-point the geographic location, view streaming media of the current location to quickly assess the situation, and, through the use of additional sensors, track the offender's vehicle.

Economic Impact: The potential economic impact of CARMEL-TerraFly is substantial because it can be a cost-effective public safety tool. It reduces law enforcement costs, increases effective-

ness of situational evaluations and responses, and contributes generally to the economic improvements of municipalities and regions. Litigation costs should also be decreased as more timely and accurate evidence becomes available for use in and out of the courtroom. In addition, the system could improve the effectiveness of situational evaluations and subsequent responses by providing tools for better resource allocation, thus improving the safety of responders and the public, ultimately saving lives and property. Finally, use of this system could ultimately reduce crime, which, in turn, would lower the cost of doing business and contribute to local and national economic improvement. CARMEL-TerraFly is the subject of an NSF SBIR Phase II project awarded to NOA, Inc., DBA TerraFly. IBM and NOA, Inc. have entered into an agreement concerning the project and are currently in the process of strategizing how best to propose the CARMEL-TerraFly System to government agencies in addition to the current CARMEL clientele. To support this, IBM and NOA Inc. have already produced a showcase video (please see: http://cake.fiu.edu/Carmel-TerraFly-video/). Our estimate is that this project will be producing $5 million in annual revenues by 2017 and will have beneficial consequences to consumers in the range of $50 million per year.

For more information, contact Naphtali Rishe, 305.348.2025, rishe@fiu.edu.

Business Continuity Information Network: Faster Community Driven Disaster Recovery

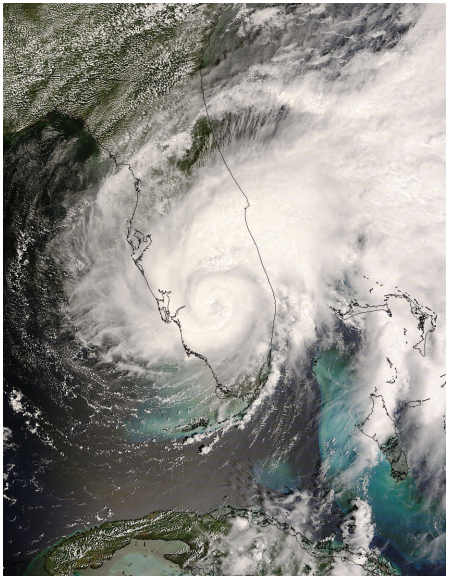

Technology used during the storm Fay, which impacted Florida.

In coastal areas throughout the United States, information sharing is critical for community resilience and protection of economic interests. Studies indicate that following hurricanes, approximately 40% of companies fail within 36 months when they were closed for 3 or more days. Years of meteorological data have demonstrated that South Florida is particularly prone to extensive damage from hurricanes. There are a myriad of toolkits, checklists, and other business continuity tools available that address how to prepare businesses for disaster. None of these stand-alone tools provide a means for business users to connect with local governments to monitor ongoing situations before, during, and after natural disasters.

The Business Community Information Network (BCIN), at CAKE, provides a platform for public and private sector communities to work in a coordinated fashion, providing the right information to the right person at the right time in the right format. Florida International University, its public and private sector partners, including Office Depot, Wal-Mart, IBM, the Greater Miami Chamber of Commerce, and county and city government agencies, have developed BCIN; a unique information sharing web-based software that provides a means for

at-risk local businesses to receive and share timely and vital preparedness, response, and recovery information. This information helps protect critical infrastructure and provide high demand recovery resources.

CAKE researchers have captured processes, workflow, and continuity "best practices" in an intuitive user interface that displays, queries, and reports on over 26 different situational categories such as ports, roads, utilities, fuel, and other critical infrastructure and recovery resources. The BCIN is available year-round as a service. This business-to-business community network provides participating companies with a new powerful tool to track their key employees and supply chain status and locate needed recovery goods and services. The system helps government agencies assess damage and prioritize recovery needs.

> **Economic Impact:** Based on training exercises, surveys, and other feedback, our participants feel they will significantly benefit by utilizing the system and its capabilities. Information sharing is critical for community resilience and overall economic well being in coastal areas throughout the United States. Since May 2009, the system has been operational in four South Florida counties: Miami-Dade, Broward, Palm Beach, and Monroe. The system was tested in response to storms Fay (see photo), Gustav, and Ike and used in numerous state and local hurricane and terrorist disaster training situations. Hundreds of individuals from government agencies, NGOs, and businesses have been acquainted with and trained on the system. Based on data from the Insurance Information Institute, if 5% of the companies in South Florida could gain the capability to speed up their hurricane recovery by one week, then $220 million of non-property economic losses could be avoided.

For more information, contact Shu-Ching Chen, 305.348.3480, chens@cs.fiu.edu.

Distributed Cloud Computing: 3-D Visualization Services for Climate Data on Demand

This breakthrough results from very successful collaborations involving two I/UCRCs, the Center for Advanced Knowledge Enablement (CAKE) at Florida International University (FIU) and Florida Atlantic University (FAU) and the Center for Hybrid Multicore Productivity Research (CHMPR) at the University of Maryland, Baltimore County (UMBC). See the entry called "Distributed Cloud Computing: 3-D Visualization Services for Climate Data on Demand" on page 97.

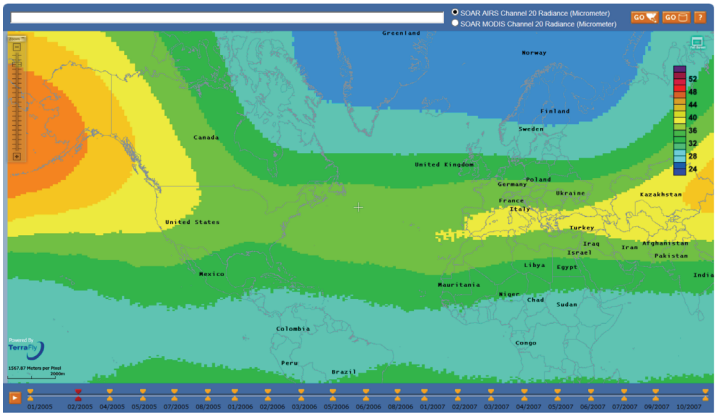

Satellite imagery enables precise measurement of global temperatures. This image presents the 8-year global average surface temperature (as Brightness Temperature colors higher on the scale are representative of warmer temperatures); by comparing successive average surface temperatures, global temperature changes can be detected.

Measuring the surface temperature of the entire earth on a daily basis is a difficult challenge because 75% of the planet is covered with oceans and ice. Continuously determining, for several days to weeks, the vertical thermal (i.e., temperature) field around a hurricane surrounded by dynamically rotating clouds is needed for more accurate landfall predictions. Thus, for applications ranging from climate change to hurricanes, satellites measure the earth's emitted infrared radiation twice daily with sufficiently high spatial and spectral (related to the spectrum) resolution to provide an estimate of vertical profiles of regional or global surface brightness temperature (BT). However, in order to assess global warming, these temperatures need to be measured to within an accuracy of 0.10 °C per year since models indicate CO_2 warming of ~20-30 °C over 100 years. Moreover, to resolve the structure around hurricanes, infrared data at resolutions of 1-5 km are needed. Not until 2002, when the Aqua (Latin for water) satellite was launched, has there been a single satellite with instruments that can meet both the accuracy and the precision required.

This breakthrough work makes it possible to deliver a decade of 3-D animated visualizations of spectral infrared (IR) satellite radiance data from instruments on Aqua. These animations use 3-D to show the vertical structure of a decade of global and regional temperature trends occurring at the surface and lower troposphere. In addition, the algorithms developed by CHMPR have been providing CAKE with 3-D temperature profiles that specify the thermal structure around hurricanes in order to improve their landfall prediction.

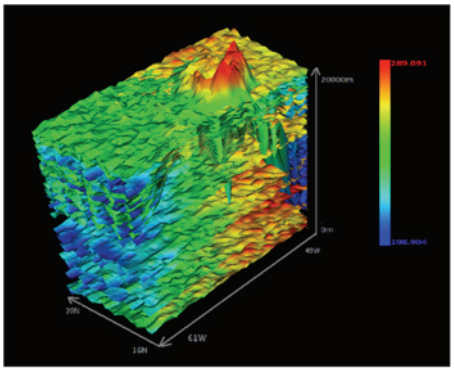

Atmospheric temperature layers up to 20,000 meters (65,619 feet). The vertical axis shows the height above sea level. The coldest (blue) and hottest (red) points in the eye of the hurricane are shown. The horizontal axes show the location of the hurricane (latitude and longitude).

CAKE and CHMPR have implemented a distributed cloud computing web-based service, called SOAR. This service incorporates visualization as a public service available on a multi-core IBM-based server cluster. This system provides researchers and students with the ability to select regional and chronological periods and automatically transform IR orbital satellite data into spherical grid arrays of 3-D temperature profiles for viewing the continuous changing thermal structure of the atmosphere. The FIU site at CAKE added value to the satellite data visualization by providing spatiotemporal (i.e., space-time) visualization and animation of the data (http://cake.fiu.edu/SOAR) using the FIU TerraFly Geospatial Data Management Service (http://terrafly.com). The FAU site at CAKE developed tools for 3-D visualization of the vertical temperature profiles; when coupled with CHMPR's data-gridding techniques, this partnership has created the first integrated, scientifically-validated, multi-year infrared brightness temperature record.

Economic Impact: Fundamental Decadal Data Records are highly desired products recommended by the National Academy of Science/National Research Council. The SOAR distributed cloud computing web-based service enhances NASA's ACCESS program by providing fundamental brightness temperature records. This can go a long way towards improving scientific and public understanding of the nature of global and regional climate change. As a result, everyone can be better positioned to design policies and actions for mitigating negative climate impacts on the economy, which could include billions of dollars of property value lost to sea-level rise and billions of dollars of insurable losses due to increases in extreme weather-related disasters.

For more information, contact Milton Halem, 410.455.2862, halem@umbc.edu, or Naphtali Rishe, http://cake.fiu.edu/Rishe, or Borko Furht, 561.297.3180, borko@cse.fau.edu.