





Center of Excellence
WIRELESS AND INFORMATION TECHNOLOGY
AT STONY BROOK UNIVERSITY

Select Research Projects



CEWIT | Center of Excellence in Wireless and Information Technology



Stony Brook University



Table of Contents

Leadership

Yacov Shamash, PhD
 Vice President of Economic
 Development and Dean of College of
 Engineering and Applied Sciences

Satya Sharma, PhD, MBA
 Executive Director

Samir R. Das, PhD
 Director, Network Technologies

Shmuel Einav, PhD
 Director, Medical Technologies

Sangjin Hong, PhD
 Director, Globalization

Arie Kaufman, PhD
 Chief Scientist

Yuanyuan Yang, PhD
 Director, Communications & Devices

Bin Zhang
 Associate Director, Computing Services

Rong Zhao, PhD
 Director, Software Systems

Security

| | |
|---|----|
| Securing Mobile Apps and Platforms..... | 6 |
| National Security Institute | 6 |
| Cloud Computing Security..... | 6 |
| Security Policy Mining..... | 7 |
| Security Policy Analysis..... | 7 |
| Fast, Architecture-neutral Static Binary Instrumentation for Security Applications..... | 7 |
| Security and Privacy in Geo-Social Networks | 8 |
| NFS4Sec: An Extensible Security Layer for Network Storage | 8 |
| Security and Privacy in Geo-Social Networks | 8 |
| NFS4Sec: An Extensible Security Layer for Network Storage | 8 |
| UID: A Strongly-Secure Usable Identity Ecosystem with Privacy..... | 9 |
| Cloudtracker: Transparent, Secure Provenance Tracking and Security Policy Enforcement in Clouds | 9 |
| Development of Security Threat Control System with Multi-Sensor Integration and Image Analysis ... | 9 |
| Internet Censorship Lab: Illuminating Internet Censorship with Principled Network Measurements | 10 |
| Detecting and Characterizing Internet Traffic Interception based on BGP Hijacking..... | 10 |

Networks

| | |
|--|----|
| Belief Evolutions in Networks of Bayesian Agents | 9 |
| Energy Efficient Reliable Data Transmissions in a Generalized Power Line Monitoring Network...9 | 9 |
| Development of Security Threat Control System with Multi-Sensor Integration and Image Analysis ... | 9 |
| Internet Censorship Lab: Illuminating Internet Censorship with Principled Network Measurements..... | 10 |
| Detecting and Characterizing Internet Traffic Interception based on BGP Hijacking..... | 10 |
| Cognitive and Efficient Spectrum Access in Autonomous Wireless Networks | 11 |
| Novel Interconnection Networks for Data Centers | 11 |
| Exploiting Control and Communications for Robust and Flexible Multi-Robot | 12 |
| Multicast in Data Center Networks..... | 12 |
| Modeling and Understanding Complex Influence in Social Networks | 12 |
| Research on Wireless Rechargeable Sensor Networks..... | 13 |
| RFID Sense-a-tags for the Internet of Things | 40 |
| Algorithmic Foundations for Hybrid Mobile Sensor Networks | 46 |

Big Data

| | |
|---|----|
| Center for Dynamic Data Analytics (CDDA): A National Science Foundation Industry/University Cooperative Research Center..... | 14 |
| Scalable Multilingual NLP through Deep Learning | 14 |
| An Ontology and Reasoning System for Access Policies to Software Services | 14 |
| The ND-Scope: Visual Analytics for High-Dimensional Data | 15 |
| Capturing Word Meanings by Language, Place and Time..... | 15 |
| Application of Parallel Computing to the Analysis of Next-generation Sequencing Data..... | 15 |
| Develop Novel Statistical Methods for Multi-loci Genetic Mapping..... | 16 |
| Efficient Resource-oblivious Parallel Algorithms | 16 |
| Cache-efficient Parallel Algorithms for fast Flexible Docking and Molecular Dynamics..... | 16 |
| FTFS: A Read/Write-optimized Fractal Free File System..... | 17 |
| Eliminating the Data Ingestion Bottleneck in Big Data Applications | 17 |
| High-performance Rule Engine for Intelligent Web Information Systems | 17 |
| Visual Correlation Analysis of Numerical and Categorical Multivariate Data..... | 18 |
| High Performance Big Data Analytics with the User in the Loop..... | 18 |
| Automatic Generation of Virtual Troubleshooting Systems Using Ontologies..... | 18 |
| Reality Deck - Immersive Gigapixel Display..... | 20 |
| NYRISE Visualization for Climate Simulation Data | 21 |
| Gigapixel Video Streaming..... | 21 |
| Fusion of GIS and Procedural Modeling for Driving Simulations in New York..... | 22 |
| Ambienizer: Turning Digital Photos into Ambient Visualizations | 28 |
| Virtual Dressing Room..... | 28 |
| An Interactive User Interface for the Smart Grid..... | 42 |
| Performance Analysis and Optimization for Logic Rule Engines | 47 |



Stony Brook University is an affirmative action/
 equal opportunity educator and employer.
 This publication can be made available in
 an alternative format upon request.
 © 2014 Stony Brook University

Table of Contents

An Efficient, Versatile, Scalable, and Portable Storage System
for Scientific Data Containers 48
Workload-aware Storage Architectures for the Optimal Performance and Energy Efficiency 49
General-Purpose Computation on Graphics Processing Units 51

Imaging and Visualization

The ND-Scope: Visual Analytics for High-Dimensional Data 15
Visual Correlation Analysis of Numerical and Categorical Multivariate Data 18
High Performance Computing for Medical Imaging 19
A Visual Framework for Health Care Analytics 19
Detecting Data Visualization Preferences Using Games 19
Visual Analytics for Open-Ended Educational Multiplayer Games 19
Energy Choices: Using an Agent-based Modeling Simulation and
Game to Teach Socio-Scientific Topics 19
Compressive Sensing Approach to Intraoperative Mass Spectrometry for Tumor Margin 20
Reality Deck - Immersive Gigapixel Display 20
NYRISE Visualization for Climate Simulation Data 21
Gigapixel Video Streaming 21
Natural Interaction with VR Environments 21
Immersive Virtual Colonoscopy 21
Interactive Immersive Visualization of Computed Microtomography Data 22
Fusion of GIS and Procedural Modeling for Driving Simulations in New York 22
Medical Volume Rendering on the Gigapixel Reality Deck 22
Lattice Simulations and Rendering 22
Plume Modeling Simulation and Visualization 23
Visual Simulation of Thermal Fluid Dynamics in a Water Reactor 23
Volumetric Segmentation of Computed Tomography Angiography 23
Automatic Spleen Segmentation for Non-Invasive Lymphoma Diagnosis 24
Mobile-based Volume Rendering Pipeline for m-Health 24
Video Streaming for Interactive Visualization on Mobile Devices 24
Volumetric Mesh Mapping 25
Conformal Mapping for Medical Imaging 25
Automatic Detection of Colon Cancer 25
Reconstruction and Registration Framework for Endoscopic Videos 26
Registration of Volumetric Prostate Scan using Curvature Flow 26
Brain Parcellation 26
Multimodal and Multivariate Visualization of the Brain 27
Meshless Point Cloud Registration by Conformal Mapping 27
Volumetric Shape Analysis 27
Saliency-aware Compression of Volumetric Data 28
Ambienizer: Turning Digital Photos into Ambient Visualizations 28
Virtual Dressing Room 28
3D Facial Recognition 29
Geometric Manifold Theory for Higher-Dimensional Data Modeling, Analysis, and Visualization ... 29
Multivariate Spline Theory, Algorithms, and Computational Techniques for
Shape Modeling and Graphics Applications 30
Wireless Sensor Network Routing 30
Human Cortical Surface Morphological Study 31
Volumetric Modeling and Shape Design in Virtual Environments 31
EmoteControl: An Interactive Multimedia Stress Management System 32
Shape Analysis with Teichmüller Shape Space 32
Conformal Wasserstein Shape Space 33
Shelterware: An Integrated Multi-Platform System to Automate the Services
of the Smithtown Animal Shelter 33
BrainRank: A Multimedia Game to Increase Fluid IQ 34
Integrating Humans and Computers for Image and Video Understanding 34
Machine Learning for the Analysis of fMRI Images 35
Learning Models for Illumination, Shadows and Shading 35
An Interactive User Interface for the Smart Grid 42



Table of Contents



Healthcare and Biomedical Applications

| | |
|--|----|
| Application of Parallel Computing to the Analysis of Next-generation Sequencing Data..... | 15 |
| Cache-efficient Parallel Algorithms for fast Flexible Docking and Molecular Dynamics..... | 16 |
| High Performance Computing for Medical Imaging..... | 19 |
| A Visual Framework for Health Care Analytics..... | 19 |
| Compressive Sensing Approach to Intraoperative Mass Spectrometry for Tumor Margin..... | 20 |
| Immersive Virtual Colonoscopy | 21 |
| Interactive Immersive Visualization of Computed Microtomography Data..... | 22 |
| Medical Volume Rendering on the Gigapixel Reality Deck..... | 22 |
| Volumetric Segmentation of Computed Tomography Angiography | 23 |
| Automatic Spleen Segmentation for Non-Invasive Lymphoma Diagnosis | 24 |
| Mobile-based Volume Rendering Pipeline for m-Health | 24 |
| Conformal Mapping for Medical Imaging..... | 25 |
| Automatic Detection of Colon Cancer | 25 |
| Reconstruction and Registration Framework for Endoscopic Videos..... | 26 |
| Registration of Volumetric Prostate Scan using Curvature Flow | 26 |
| Brain Parcellation | 26 |
| Multimodal and Multivariate Visualization of the Brain | 27 |
| Platform Comparison via Errors in Variables Models with or without Replicates | 36 |
| Biomarker Agreement and Integration Analyses across Measurement Platforms | 36 |
| Comparing the RNA-seq and Microarray Platforms with Generalized Linear EIV Model | 36 |
| Bootstrap-based RANOVA for microRNA Panel Data Analysis | 37 |
| Integrative Multi-Scale Biomedical Image Analysis..... | 37 |
| Feature-based Exploration of Extremely Large Spatio-Temporal Scientific Datasets..... | 37 |
| Unified Statistical Models for Integrative Omics | 38 |
| Home Sleep Monitoring Using Bluetooth LE Biosensors to Facilitate Sleep Health..... | 38 |
| Developing Computer Games and Virtual Reality Environments for Use in Physical Rehabilitation | 39 |
| Mobile Healthcare Device Platform..... | 39 |
| Study of How Information Technology Affects Patient Access, Quality of Care, and Administrative Costs in a Newly Formed ACO | 39 |

Internet of Things

| | |
|--|----|
| Research on Wireless Rechargeable Sensor Networks..... | 13 |
| Mobile Healthcare Device Platform..... | 39 |
| RFID Sense-a-tags for the Internet of Things | 40 |
| Design of Audio Interfaces in Adverse Acoustic Environments | 40 |
| Non-isotropic Networked Sensor Deployment for Smart Buildings..... | 40 |
| Self-powered Wireless Hybrid Langasite Sensor for Pressure/ Temperature Monitoring of Nuclear Reactors..... | 40 |
| Algorithmic Foundations for Hybrid Mobile Sensor Networks | 46 |

Smart Energy and Smart Urban Systems

| | |
|---|----|
| Non-isotropic Networked Sensor Deployment for Smart Buildings..... | 40 |
| Smart Grid Regional Demonstration – Long Island Smart Energy Corridor | 41 |
| Smart Composites for Energy Harvesting | 41 |
| Thin Film Solar Cells with Tunable Transparency | 41 |
| An Interactive User Interface for the Smart Grid..... | 42 |
| An Interactive User Interface for the Smart Grid..... | 42 |
| Enhanced Power System Operation and Control | 42 |
| Smart Grid Android Manager..... | 42 |


Table of Contents

Algorithms

| | |
|---|----|
| Security Policy Mining | 7 |
| Security Policy Analysis | 7 |
| Modeling and Understanding Complex Influence in Social Networks | 12 |
| An Ontology and Reasoning System for Access Policies to Software Services | 14 |
| Scalable Multilingual NLP through Deep Learning | 14 |
| Capturing Word meanings by Language, Place, and Time | 15 |
| Application of Parallel Computing to the Analysis of Next-generation Sequencing Data | 15 |
| Develop Novel Statistical Methods for Multi-loci Genetic Mapping | 16 |
| FTFS: A Read/Write-optimized Fractal Free File System | 17 |
| Eliminating the Data Ingestion Bottleneck in Big Data Applications | 17 |
| High-performance Rule Engine for Intelligent Web Information Systems | 17 |
| Visual Correlation Analysis of Numerical and Categorical Multivariate Data | 18 |
| Automatic Generation of Virtual Troubleshooting Systems Using Ontologies..... | 18 |
| Lattice Simulations and Rendering | 22 |
| Volumetric Mesh Mapping | 25 |
| Meshless Point Cloud Registration by Conformal Mapping | 27 |
| Volumetric Shape Analysis | 27 |
| Saliency-aware Compression of Volumetric Data | 28 |
| Design of Audio Interfaces in Adverse Acoustic Environments | 40 |
| Non-Visual Skimming: Improving the Usability of Web Access for Blind People..... | 43 |
| Divisible Load Scheduling | 43 |
| Approximation Algorithms for Geometric Optimization | 43 |
| Geometric Networks..... | 43 |
| Algorithms in Support of Flight Trajectory Analysis..... | 44 |
| Using Evolutionary Computations to Understand the Design and Evolution of Gene and Cell Regulatory Networks..... | 44 |
| Adaptive Runtime Verification and Recovery for Mission-Critical Software | 45 |
| Computational Modeling and Analysis for Complex Systems | 45 |
| Algorithmic Foundations for Hybrid Mobile Sensor Networks | 46 |
| From Clarity to Efficiency for Distributed Algorithms..... | 46 |
| Demand-driven Incremental Object Queries | 47 |
| Performance Analysis and Optimization for Logic Rule Engines | 47 |
| Building Metaknowledge Representations in Circuit Design: | |
| Symbolic Data Mining for Systematic Modeling of Analog Circuits..... | 50 |

Advanced Computing Systems

| | |
|--|----|
| Cache-efficient Parallel Algorithms for fast Flexible Docking and Molecular Dynamics..... | 16 |
| Efficient Resource-oblivious Parallel Algorithms | 16 |
| FTFS: A Read/Write-optimized Fractal Free File System..... | 17 |
| High Performance Big Data Analytics with the User in the Loop | 18 |
| Video Streaming for Interactive Visualization on Mobile Devices | 24 |
| Adaptive Runtime Verification and Recovery for Mission-Critical Software | 45 |
| Computational Modeling and Analysis for Complex Systems | 45 |
| Eco Hadoop: Cost and Energy-Aware Cloud and HPC Computing | 48 |
| Secure Provenance in High-End Computing Systems | 48 |
| An Efficient, Versatile, Scalable, and Portable Storage System for Scientific Data Containers | 48 |
| Workload-Aware Storage Architectures for Optimal Performance and Energy Efficiency..... | 49 |
| Pochoir: A Stencil Computation Compiler for Modern Multicore Machines | 49 |
| Energy-Efficient Superconductor SFQ Processor Design | 50 |
| Building Metaknowledge Representations in Circuit Design: | |
| Symbolic Data Mining for Systematic Modeling of Analog Circuits | 50 |
| Leveraging Three-Dimensional Integration Technology for Highly Heterogeneous Systems-on-Chip..... | 50 |
| General-Purpose Computation on Graphics Processing Units | 51 |



Securing Mobile Apps and Platforms

Long Lu
long@cs.stonybrook.edu

Smartphones, tablets, and other mobile devices have outgrown personal computers to become the dominant tools supporting users' day-to-day needs for communications, computations, and entertainment. While enjoying the mobility and versatility of these fairly popular devices, users are facing an emerging wave of security and privacy threats. These threats have not been seen on conventional computers or systematically studied. The current research projects in RiS3 Lab (Research in Systems and Software Security) aim at understanding and mitigating the two general causes for security threats on mobile devices: (1) vulnerable applications resulted from flawed designs or poor development practices; (2) inadequate and unsuited OS-level security mechanisms originally designed for non-mobile platforms. As the first step towards automatic detection of vulnerabilities in apps, we undertook the challenging task of building the first comprehensive static program analysis framework for Android apps. This framework addresses several open problems in analyzing mobile code, such as modeling asynchronous execution of app components and discovery of implicit code entry points. Next, based on this framework, we will design scalable and accurate vulnerability detectors desirable to app stores, enterprise IT, and security service providers. At the operating system level, our previous studies revealed some major security problems of both iOS and Android platforms. Recognizing the distinct security requirements of mobile applications, we revisited the traditional OS design assumptions that may not hold on mobile devices. Our ongoing projects are now retrofitting mobile OSes with new process and memory management methods that better safeguard user privacy even with the presence of powerful attackers. (NSF)

Cloud Computing Security

Radu Sion
sion@cs.stonybrook.edu


As computing becomes embedded in the very fabric of our society, the exponential growth and advances in cheap, high-speed communication infrastructures allow for unprecedented levels of global information exchange and interaction. As a result, new market forces emerge that propel toward a fundamental, cost-efficient paradigm shift in the way computing is deployed and delivered: computing outsourcing. Outsourcing has the potential to minimize client-side management overheads and benefit from a service provider's global expertise consolidation and bulk pricing. Companies such as Google, Amazon, and Microsoft are rushing to offer increasingly complex storage and computation outsourcing services supported by globally distributed "cloud" infrastructures. Yet significant challenges lie in the path to successful large-scale adoption. In business, healthcare and government frameworks, clients are reluctant to place sensitive data under the control of a remote, third-party provider, without practical assurances of privacy and confidentiality. Today's solutions however, do not offer such assurances, and are thus fundamentally insecure and vulnerable to illicit behavior. Existing research addresses several aspects of this problem, but advancing the state of the art to practical realms will require a fundamental leap. This project addresses these challenges by designing, implementing and analyzing practical data outsourcing protocols with strong assurances of privacy and confidentiality. It will also initiate the exploration of the cost and energy footprints of outsourcing mechanisms. This is essential as the main *raison d'être* of outsourcing paradigms lies in their assumed end-to-end cost savings and expertise consolidation. Yet, research has yet to explore and validate the magnitudes of these savings and their underlying assumptions. (NSF, ARO, Microsoft Research)

National Security Institute

Radu Sion
sion@cs.stonybrook.edu

Stony Brook University, the flagship state university of New York, is establishing the National Security Institute (NSI). The NSI vision and its core mission are bold: to secure our homeland by researching and developing technologies and insights for secure, trustworthy, and available communications and computing platforms. NSI's goal is to become a world leader in research, the education of professionals, security technology, business and policy, and raising awareness. NSI will span multiple disciplines and establish public-private partnerships to develop new holistic socio-technological solutions for securing our highly-digital societies; it will engage not only in research but also in the education of professionals in defense, national and cyber-security, assurance, healthcare, and policy. A comprehensive assurance education program will be established, to train not only Stony Brook students but also the broader corporate and academic community.

NSI will leverage the team's strengths to spawn a steady stream of successful security-centric technology startups. (CEWIT)



Security Policy Mining

Scott D. Stoller
stoller@cs.stonybrook.edu

The most widely used representation of access control policy is still the humble access control list (ACL). Expressing policies in a higher-level language, such as a role-based access control (RBAC) policy language or an attribute-based access control (ABAC) policy language, can make the policies much easier---and hence cheaper---to understand, analyze, and maintain. For example, consider the policy, "A user working on a project can read and request to work on a non-proprietary task whose required areas of expertise are among his/her areas of expertise." In ABAC, this policy can be expressed with a single rule, regardless of the number of users, projects, and tasks. In RBAC, this policy requires creation of a role for each task, creation of user-role and permission-role assignments for each of those roles, and updates to those assignments when relevant attribute values change (for example, a user gains an area of expertise). In an ACL policy, a separate entry is needed for each combination of a user and a task for which the user has permissions. The RBAC and ACL policies require significantly more management effort, and are more prone to management errors, than the ABAC policy.

Policy mining is the problem of constructing a higher-level (typically role-based or attribute-based) policy that is equivalent, or approximately equivalent (if noise is present), to a given policy represented as ACLs. Policy mining can significantly reduce the cost of migrating to a higher-level policy language in large organizations. We are developing new algorithms for mining RBAC policies and ABAC policies. Our RBAC policy mining algorithm can easily be used to optimize a variety of policy quality metrics, including metrics based on policy size, metrics based on interpretability of the roles with respect to user attribute data, and compound metrics that consider size and interpretability. In experiments with real access control policies released by HP Labs, our algorithm

achieves significantly better results than previous algorithms. Our first ABAC policy mining algorithm mines the policy from an ACL policy and data about user attributes and resource attributes. Our second ABAC policy mining algorithm mines the policy from operation logs and attribute data. It allows a controlled trade-off between policy quality (reflecting how well the resulting permissions correlate with attribute data) and the number of speculative permissions (i.e., permissions granted by the generated policy even though they do not appear in the log). (NSF, ONR)

Security Policy Analysis

Scott D. Stoller
stoller@cs.stonybrook.edu

In large organizations, access control policies are managed by multiple users (administrators). An administrative policy specifies how each user may change the access control policy. Fully understanding the implications of an administrative policy in an enterprise system can be difficult, because of the scale and complexity of the access control policy and the administrative policy, and because sequences of changes by different administrators may interact in unexpected ways. Administrative policy analysis helps by answering questions such as user-permission reachability, which asks whether specified users can together change the policy in a way that achieves a specified goal, namely, granting a specified permission to a specified user. We are developing algorithms for administrative policy analysis for role-based and attribute-based policy frameworks.

For attribute-based frameworks, the problem is particularly challenging, because administrators may change policy rules as well as facts (e.g., facts about attribute data). To provide more broadly applicable results, our algorithm performs abductive analysis, which means that it can answer user-permission reachability queries for given initial policy rules, even in the absence of specific data about initial user attributes and resource attributes. It does this by identifying minimal conditions on

the initial attribute data that enable specified users to together change the policy in a way that grants the specified permission to the specified user. (NSF, ONR)

Fast, Architecture-Neutral Static Binary Instrumentation

R. Sekar
sekar@seclab.cs.sunysb.edu

Program instrumentation techniques form the basis of many recent software security defenses, including defenses against common exploits, security policy enforcement, application monitoring and debugging. As compared to source-code instrumentation, binary instrumentation is easier to use and more broadly applicable due to the ready availability of binary code. Moreover, source-code based instrumentation may be incomplete because some of it may be eliminated by compiler optimizations, and because some low level code added by linkers (or compilers) is not instrumented.

One of the major challenges in binary instrumentation is the complexity of modern instruction sets. Accurate instrumentation requires the semantics of all instructions to be captured, since all of the analyses and transformations performed by the instrumentor are based on this semantics. Any errors in modeling instructions will likely cause instrumented programs to fail. Clearly, this is a daunting task even for a single architecture: the Intel manual describing the x86 instruction set runs to over 1500 pages describing over 1100 instructions. When this task is multiplied across different architectures such as ARM, PowerPC, SPARC, MIPS, etc, the effort involved becomes impractically large. We are therefore developing a novel approach that avoids the need for modeling instructions by leveraging knowledge embedded in retargetable code generators in today's compilers such as GCC. This approach not only simplifies the development of instrumentation, but also makes it applicable to all architectures for which a code generator is available.

Another important advance made by our approach is that of enabling a rich set of optimizations to be performed on binary instrumentations, thereby significantly improving performance over today's techniques. Moreover, our approach enables the use of today's compiler backends for generating and optimizing instrumentations, thereby achieving architecture-independent instrumentation.

Today's binary instrumentation techniques have largely been based on dynamic (i.e., runtime) binary instrumentation (DBI). DBI techniques provide two key features needed for security instrumentation: (a) it should be applied to all application code, including code contained in various system and application libraries, and (b) it should be non-bypassable. Previous static binary instrumentation (SBI) techniques have lacked these features. However, DBI techniques can incur very high overheads in several common usage scenarios, such as application startups, system-calls, and many real-world applications. We have therefore developed a new platform for secure static binary instrumentation (PSI) that overcomes these drawbacks of DBI techniques, while retaining the security, robustness and ease-of-use features. Our experimental results have demonstrated an order of magnitude improvement in performance over DBI techniques on many real-world applications. (NSF)

Security and Privacy in Geo-Social Networks

Radu Sion
sion@cs.stonybrook.edu

Location based social or geosocial networks (GSNs) have recently emerged as a natural combination of location based services with online social networks: users register their location and activities, share it with friends and achieve special status (e.g., "mayorship" badges) based on aggregate location predicates. Boasting millions of users and tens of millions of daily check-ins, such services pose significant privacy threats: user location information may be tracked and leaked to third parties. Conversely, a solution enabling location privacy may provide cheating capabilities to users wanting to claim special location status. In this paper we introduce new mechanisms that allow users to (inter)act privately in today's geosocial networks while simultaneously ensuring honest behavior. An Android implementation is provided. The Google Nexus One smartphone is shown to be able to perform tens of badge proofs per minute. Providers can support hundreds of millions of check-ins and status verifications per day. (NSF, ONR)

NFS4Sec: An Extensible Security Layer for Network Storage

Radu Sion and Erez Zadok
sion@cs.stonybrook.edu,
zadok@cs.stonybrook.edu

The Network File System (NFS) is a popular method for computers to access files across networks. The latest major version of this IETF protocol, version 4, is widely accepted and includes numerous new features to improve security, performance, and usability when used over wide-area networks. However, the NFSv4's security focus is on network-wide encryption (ensuring that user data cannot be intercepted) and user authentication (ensuring that only legitimate users can access their files); it does not address end-to-end data security (i.e., persistently stored data), does not address data integrity (malicious or benign data corruptions), and more. This project extends NFSv4 with a security layer that allows one to develop multiple, composable plugin modules to enhance the protocol's security. This layer allows for interception of protocol requests between clients and servers to perform various useful security functions: logging access to files by users and hosts, useful for regulatory compliance reports and audits; inspecting files for malware patterns and automatically quarantining them; verifying the integrity of long-lived files against malicious changes (e.g., Trojan intrusions) and benign but serious ones (e.g., storage media degradation and hardware corruptions); detecting denial-of-service attempts and ensuring quality-of-service to legitimate users through load-balancing and redirection; automatic snapshotting and logging to allow for forensic analysis and recovery from failures and intrusions. In a cloud-based era where more data lives longer and is accessed over wide-area, insecure networks, this project help elevate the level of security of every user's data files. (NSF)



uID: A Strongly-Secure Usable Identity Ecosystem with Privacy

Radu Sion
sion@cs.stonybrook.edu

uID is a secure, usable, privacy-enabling digital identity ecosystem, able to integrate, and synergize with existing governmental, commercial and open-source identity and authentication solutions.

Designing tomorrow's digital identity solution is faced with unique challenges. Identity mechanisms overwhelmingly refer to and are used by people. They need to be usable and affordable, and address individual concerns of privacy and confidentiality. At the same time, to ensure trust they need to provide accountability and be strongly secure.

Further, it is important to realize that no one platform can be a sole provider – a viable ecosystem will have standards with well specified APIs and conduits for interoperability that naturally foster a healthy market. Finally, it is essential that these mechanisms interoperate and are efficient so as to not constitute a bottleneck when deployed.

While addressing all of the above challenges, uID will focus on two key goals: privacy protection and transaction unlinkability. These properties are unfortunately conflicting and require a complex multi-layer research and development approach calling on multi-disciplinary expertise across all the layers of today's digital transactions. Simple "browser plugins" or "email-based" mechanisms alone are bound to fail by not considering the multiple cross-layer security challenges. (CEWIT)

Cloudtracker: Transparent, Secure Provenance Tracking and Security Policy Enforcement in Clouds.

Radu Sion and Don Porter
sion@cs.stonybrook.edu
porter@cs.stonybrook.edu

As companies, governments, and individual users adopt increasingly diverse computing platforms, from outsourced cloud computations to personal laptops and mobile devices, enforcing uniform security policies across these platforms becomes unwieldy.

Similarly, regulatory compliance and business auditing requires tracking the history of this data in a comprehensive, secure, and platform-independent manner. Unfortunately, technology has not kept pace with these practical concerns, and several systems and security research challenges must be addressed to make this vision a reality.

There is a natural and under-explored connection between understanding the origins of data and using that data's history to enforce security policies. To leverage this connection, this project is developing a comprehensive, general framework for automatically tracking the history of data and enforcing associated security policies in cloud computing environments. The research focuses on three key research challenges.

First, the project investigates novel applications of virtualization technologies to transparently infer data provenance by inspecting a guest operating system (OS) and applications. Second, this project is developing techniques to securely store, manage, and query provenance data at cloud scale. Finally, the project combines the first two technologies to transparently and collaboratively enforce security policies throughout the cloud and end-user systems.

The prototype system is designed to allow individual users and organizations to rapidly adopt new technology platforms, from clouds to novel end-user systems, without having to worry about the interaction of these new systems with security policies and regulatory compliance concerns. (NSF)

Development of Security Threat Control System with Multi-Sensor Integration and Image Analysis

Sangjin Hong
Sangjin.hong@cewit.stonybrook.edu

This purpose of this project is to develop a heterogeneous sensor network platform for intelligent surveillance system applications. Multiple image sensors, such as hyperspectral image sensors, are utilized in order to capture various undetectable images for discovering hidden information. Different sensors collaborate for creating consistent overall real-time information to be used to infer many abnormal surrounding situations. In the project, following researches are conducted. Multiple object association through estimation and prediction, low complexity embedded system design, large scale system modeling, multimedia database access strategy, and stochastic collaborative signal processing. (Korea MKE)

Internet Censorship Lab: Illuminating Internet Censorship with Principled Network Measurements

Phillipa Gill
phillipa@cs.stonybrook.edu

The Internet was not designed with information controls, such as censorship or surveillance, in mind. However, its importance has led many nations to repurpose Internet protocols (e.g., the Domain Name System (DNS) and Border Gateway Protocol (BGP)), and network management products (e.g., Web proxies, traffic shapers) for information control. This unintended use of networking technologies can lead to unintended international impact of censorship, and raises many ethical issues when network management products are exported to countries that use them to violate human rights.

To address these challenges, this project: (1) develops a platform which enables repeatable measurements of Internet censorship while mitigating risks to individuals performing the measurements, (2) designs techniques to analyze data from the platform to detect different types of censorship and even specific network management products used for censorship, and (3) uses the platform to quantify instances of national censorship which have unintended international impact on the Internet.

The use of technology to restrict freedom of speech and persecute dissidents around the globe is a human rights concern. This project provides high fidelity technical data to inform policy discussions. Further, the technical analysis provides insights into the global impacts of national censorship on the Internet, and how proposed improvements to existing protocols (e.g., DNSSEC, BGPSEC) can mitigate these issues. (NSF, Google)

Detecting and Characterizing Internet Traffic Interception based on BGP Hijacking

Phillipa Gill
phillipa@cs.stonybrook.edu

Recent reports have highlighted incidents of massive Internet traffic interception executed by re-routing BGP paths across the globe (affecting banks, governments, entire network service providers, etc.). The potential impact of these attacks can range from massive eavesdropping to identity spoofing or selective content modification. In addition, executing such attacks does not require access or proximity to the affected links and networks, posing increasing risks to national security. Worse yet, the impact of traffic interception on the Internet is practically unknown, with even large-scale and long-lasting events apparently going unnoticed by the victims.

As reported by Renesys Corporation in November of last year, there is evidence that traffic interception events are growing more frequent, but there are no validated methods to immediately detect them or evaluate their impact. The architectural innovation that mitigates this inherent protocol design flaw exploited by such attacks, is slow to take off, suggesting that this vulnerability will persist, leaving our critical communication infrastructure exposed.

Because of their complex dynamics, and the number of different actors involved on a global scale, devising effective methodologies for the detection and characterization of traffic interception events requires empirical and timely data (e.g., acquired while the event is still ongoing). By leveraging our experience in measuring and investigating events affecting inter-domain communication and leveraging our measurement and data processing infrastructure, this research project is working on the following goal: (i) investigate, develop, and experimentally evaluate novel methodologies to automatically detect traffic interception events and to characterize their extent, frequency, and impact; (ii) extend our measurement infrastructure to detect in near-realtime and report episodes of traffic interception based on

BGP hijacking; (iii) document such events, providing datasets to researchers as well as informing operators, emergency-response teams, law-enforcement agencies, and policy makers. In characterizing their impact, we will quantify increased latency along observed paths, the magnitude of the incident in terms of number of ASes and prefixes intercepted, and the social/political implications of interceptions that take traffic across national borders. We will augment our active measurement framework with algorithmic simulations of BGP routing policies, and qualitative analysis of the organizations involved, to better understand the both technical and political effects of hijacks. (CEWIT)

Belief Evolutions in Networks of Bayesian Agents

Petar M. Djuric
petar.djuric@stonybrook.edu

The research is on understanding the processes of belief evolution about the state of nature in time and/or space in networks of Bayesian agents and in a wide variety of settings. We study the evolution of beliefs in static networks (with fixed number of agents and fixed topology) and dynamic networks (with varying number of agents and social ties). The latter includes networks that allow new agents to join the network and its current agents to disappear and reappear or to change their social ties. Special attention is given to problems of evolution of beliefs in networks where agents communicate correlated information. Evolution of beliefs in networks whose agents make random decisions and evolution of beliefs in presence of malicious agents is also studied. (NSF)



Novel Interconnection Networks for Data Centers

Yuanyuan Yang
yuanyuan.yang@stonybrook.edu

Driven by technology advances, massive data centers consisting of tens or even hundreds of thousands servers have been built as infrastructures by large online service providers. Designing a cost-effective network topology for data centers that can deliver sufficient bandwidth and consistent latency performance to a large number of servers has been an important and challenging problem. A good network topology should meet the following requirements: (1) expandability, which means that expanding an existing network should not incur huge extra cost in either man power or device replacement; (2) plenty parallel paths between any pair of servers to guarantee a sufficient bandwidth and graceful degradation; (3) small network diameter so that a task can be assigned to any part of the network as required by cloud computing; (4) low cost of the interconnection structure. As the size of data centers becomes larger and larger, the cost of the interconnection structure becomes a key factor in practical applications. Many data center network topologies have been proposed recently, which can be basically divided into two categories: switch-centric networks and server-centric networks. We conduct extensive research in both categories. In this project, we mainly focus on server-centric networks and propose a novel server-centric network topology called BCube Connected Crossbars (BCCC) which can meet all the requirements described above. We also propose routing algorithms for different communication patterns in the BCCC. It is shown that the BCCC outperforms the current popular data center networks. (NSF)

Energy Efficient Reliable Data Transmissions in a Generalized Power Line Monitoring Network

Xin Wang
xwang@ece.sunysb.edu

Efficient power line monitoring is essential for reliable operation of the Smart Grid. A Power Line Monitoring Network (PLMN) based on wireless sensor nodes can provide the necessary infrastructure to deliver data from the extension of the power grid to one or several control centers. However, the restricted physical topology of the power lines constrains the data paths, and has a great impact on the reporting performance. In this work, we present a comprehensive design to guide efficient and flexible relay selection in PLMNs to ensure reliable and energy efficient transmissions while taking into account the restricted topology of power-lines. Specifically, our design applies probabilistic power control along with flexible transmission scheduling to combat the poor channel conditions around power line while maintaining the energy level of transmission nodes. We consider the impact of different channel conditions, non-uniform topologies for a power line corridor and the effect of reporting events. Our performance results demonstrate that our data forwarding scheme can well control the energy consumption and delay while ensuring reliability and extended lifetime. (CEWIT)

Cognitive and Efficient Spectrum Access in Autonomous Wireless Networks

Xin Wang
xwang@ece.sunysb.edu

The exponential growth of wireless traffic calls for novel and efficient spectrum access techniques and wireless network infrastructures. The recent introduction of autonomous network concept, where femtocell is an application, presents as a paradigm shift from traditional cellular network with planned deployment and

centralized management to more autonomous, uncoordinated, and intelligent rollouts of small base stations deployed by end users and overlaying wide-area cellular networks. Fundamentally different from conventional cellular networks or heterogeneous networks with nodes operating in different spectrum bands, femtocells are deployed in an ad hoc manner and share the same spectrum band as the cellular networks to increase the spectrum usage efficiency and allow a terminal to seamlessly operate in macrocells and femtocells. This network infrastructure, however, would create strong cross-tier and intra-tier interference. The objective of this project is to enable more efficient and reliable operation of autonomous femto-cell networks with agile spectrum access, autonomous interference control, as well as intelligent network self-organization and self-optimization. This project falls into four interacted thrusts: 1) Incorporate cognition into the femtocell networks to cognitively reuse the available spectrum sensed; 2) Develop distributed, dynamic and cooperative interference management schemes exploiting antenna techniques and based on sensed environmental conditions; 3) Investigate the scenarios and schemes that femtocells can be exploited to facilitate macrocell transmissions, and the potential gains in capacity, coverage and reliability; 4) Incorporate interference cancellation for data multicast, and develop techniques to support multiuser video streaming. The project also develops a testbed with open source programmable wireless platforms, for prototyping and evaluating the effectiveness of various techniques developed. The proposed research has the potential to significantly increase the capacity and resilience of existing and future wireless networks. The agility and resilience of the system will also make it instrumental to support communications and applications that are important for national security and economy. (NSF)

Exploiting Control and Communications for Robust and Flexible Multi-Robot Coordination

Xin Wang

xwang@ece.sunysb.edu

The objective of this project is to investigate a set of control and communication techniques for establishing and maintaining communication connections among multiple collaborating mobile robots, in response to varying communication conditions in practice, to ensure robust and flexible multi-robot coordination. The approach is to exploit the features of wireless communications and mobility of robots to significantly increase the coverage and reliability of communications and the chance of forming communication links among mobile robots. This research addresses the critical challenge in the establishment and maintenance of communication connections among collaborative mobile robots, which is imposed by unstable communication conditions and the dynamic nature of robots and environments. The contribution of this work lies in the transformative development and integration of novel distributed control and communication techniques for controlled communications, multi-robot motion, and wireless signal search under unstable communication conditions and robot/task constraints. The integrated application of these techniques will lead to robust and efficient communication networking, high-freedom task operation and exploration, and thus highly robust and flexible coordination among multiple collaborative robots. (NSF)

Modeling and Understanding Complex Influence in Social Networks

Jie Gao

ygao@cs.stonybrook.edu

Social interactions constitute a crucial part of everyday life. Behavior changes, similar to rumors or viruses, spread in the social network and become a contagion. Some of these contagions can be beneficial (e.g., adopting healthy lifestyles) or profitable (e.g., viral marketing), and we would like to encourage or promote them. Some of these contagions can be destructive (such as teenager smoking, obesity, or alcohol abuse), and we would like to discourage or stop them. It is therefore of great importance and urgency to understand how these contagions naturally spread in social networks and how to effectively encourage or discourage a contagion with our available (limited) resources.

Diseases and information can spread through a single contact. Thus they spread fast in social networks with the small world property. However, in most of the realistic settings when agents' actions and behavioral changes are involved, it often takes multiple activated neighbors to spread a contagion. We denote this type of contagion as a complex contagion. The requirement of synergy between neighbors, intuitively, makes the spreading of a complex contagion to be more unlikely, slower, and more delicate. Enabling the successful spreading of a complex contagion requires special graph structures.

This project will answer both the scientific question of what factors enable a complex contagion to spread in a social network, as well as the engineering question of how to design intervention schemes to guide complex contagions. It will provide the fundamental understanding of the interplay of structural properties of the social network and the behaviors of social processes operating on the network. Specific questions to be addressed include: how fast do complex contagions spread in a social network, on both model networks and real world networks; how does the choice of

initial seeds matter; how does the degree distribution of the network, the 'rich club' property, or community structure relate to the spread of complex contagions; how do we stop complex contagions from spreading; and finally how do external environmental factors affect complex contagions. We will conduct theoretical, rigorous analysis of the behavior of complex contagions under various models of the social network and the diffusion process, as well as simulations on real world networks to calibrate our model construction.

The results of this project will have deep and long-lasting social impact. Rigorous understanding of complex contagions will provide much needed theoretical guidance to real world applications, ranging from healthcare to word-of-mouth advertising, from influencing cultural trends to political campaigns. The technical content of this project is inherently interdisciplinary and will have direct applications to related fields such as probability, economics, sociology, and statistical physics. (NSF, DARPA)

Multicast in Data Center Networks

Yuanyuan Yang

yuanyuan.yang@stonybrook.edu

Data center networks (DCNs) interconnect tens of thousands of servers to form the infrastructure for today's ubiquitous collaborative computing, which are the backbone of the clouds. Most of today's DCNs provide high bandwidth using inexpensive commodity switches. Also, many online applications (e.g., web searching) and backend infrastructural computations (e.g., distributed file system and database) hosted by data centers require one-to-many group communication. Network-level multicast can greatly benefit such group communication through reducing network traffic and releasing the sender from duplicated transmission tasks, thus significantly improving application throughput and increasing network capacity. Several unique features in data centers facilitate the implementation of multicast. First of all virtualization of machines and networks provides plenty of flexibility. For



computing tasks submitted to data centers, physical resources can be logically divided and each task can be provided a virtual network according to its scale and requirements. We exploit the freedom of resource allocation to meet several goals: (1) to deliver satisfied user experience, i.e., the communication between members of a virtual network is nonblocking for multicast; (2) to achieve efficient resource utilization, i.e., hardware cost of the data center network can be reduced. The second valuable feature in data centers is server redundancy. Server redundancy is originally proposed in order to deliver high availability and fault tolerance in data centers. Nevertheless, this feature can be utilized to improve the design of data center networks and reduce network cost. For example, the number of core switches in a fat-tree data center network can be dramatically reduced if the server redundancy is taken into consideration in the overall network design. Third, the recently developed OpenFlow framework gives endless possibility to design various network control strategies based on different needs. We develop several online multicast scheduling algorithms to balance the traffic load in data center networks. These algorithms not only improve the utilization of network bandwidth, but also ensure that no link in the network is oversubscribed. (NSF)

Research on Wireless Rechargeable Sensor Networks

Yuanyuan Yang
yuanyuan.yang@stonybrook.edu

In this project, we introduce the novel wireless charging technology to power wireless sensor networks. The objective is to achieve perpetual network operations as well as improve network performance. Traditional battery-powered sensor networks usually have limited lifetime that pose great challenges to meet a variety of energy-demanding applications. Energy harvesting from environmental sources can sustain network operations. However, dynamics from the energy sources may cause interruptions in network services and degrade performance greatly. The novel wireless charging technology has opened up a new dimension to replenish energy in sensor networks without wires or plugs. A charging vehicle equipped with resonant coils can move around the field to recharge nodes conveniently. Several important issues are studied in this project.

The first question is how to gather real-time node's energy information in a scalable manner. To achieve this, we propose an NDN (Named Data Networking)-based real time communication protocol for gathering real-time energy status that divides the network into hierarchies. We leverage concepts and mechanisms from NDN to design a set of protocols that continu-

ously gather and deliver energy information to the vehicle to cope with a variety of recharge requests. Analytical results based on the energy neutral conditions that give rise to perpetual operation are also derived.

The second question is how to schedule vehicle(s) to achieve perpetual operations. For a single vehicle, we formulate the problem into an Orienteering problem, which is NP-hard. It aims to maximize the total energy recharged in a given time. Since the problem is NP-hard in nature, we take reasonable approximations to simplify the problem into a Knapsack problem so we can develop polynomial-time solutions. Further, the problem on how to schedule multiple vehicles that has more scalability and robustness immediately follows. Our focus is to minimize the vehicles' total traveling cost while ensuring all nodes are functional. We formulate the problem into a Multiple Traveling Salesman Problem with Deadlines (m-TSP with Deadlines), which is also NP-hard. To accommodate energy dynamics and reduce computational overhead, we develop an on-line algorithm that selects the node with the minimum weighted sum of traveling time and residual lifetime. Our scheme not only improves network scalability but also guarantees the perpetual operation of networks.

The third problem is how to integrate wireless charging with traditional sensing applications such as mobile data gathering. We can combine wireless charging and data gathering utilities on a single vehicle to improve spatial-temporal efficiency. Our objective is to maximize the network utility. First, a set of sensor nodes with minimum residual energy are selected for the vehicle to perform recharge. The algorithm ensures a bounded traveling time under a given threshold. Upon traversing each node for recharging, the vehicle collects data messages in the neighborhood by multi-hop transmissions. To maximize network utility, we divide the original problem into several sub-problems to find the optimal data rates, flow routing and vehicle's stopping time. Distributed algorithms to solve these problems are proposed and convergence properties are examined. (NSF)

Center for Dynamic Data Analytics (CDDA): A National Science Foundation Industry/University Cooperative Research Center

Arie Kaufman, Rong Zhao
ari@cs.stonybrook.edu

The Center for Dynamic Data Analytics (CDDA), an NSF supported Industry/University Cooperative Research Center (I/UCRC), was established in 2011 to conduct pre-competitive research to manage, analyze, and visualize massive, complex, multidimensional and multi-scale dynamic data. The CDDA is a multi-institution consortium located at Rutgers University and Stony Brook University. Its mission is to turn chaos into knowledge and to unleash the transformative potential of big data in a wide range of application domains, such as information technology, healthcare, pharmaceutical, biotechnology, commerce, retail, finance, insurance, media, entertainment, transportation, logistics, manufacturing, defense, security, education, and public administration. Dynamic data pose new challenges in algorithm design for analysis and visualization that traditionally have not been addressed.

Types of data that are considered dynamic in nature include data captured from unmanned vehicles; sensor networks; telecommunications networks; utility grid and other critical infrastructure; mobile and interactive Web 2.0 applications; robotics and intelligent agents; biomedical and healthcare informatics; computer simulations and modeling; real-time financial applications; social media and entertainment; just to name a few. Our research will strategically focus on design and evaluation methods, algorithms, architectures, software, visualization techniques, mathematical and statistical foundations, and benchmarking of complex systems that facilitate large-scale, dynamic data analytics. The ultimate goal of the CDDA is to develop new technologies that can be applied by its industry partners to create value across a wide range of sectors. For more information please visit cdda.cs.stonybrook.edu. (NSF)

An Ontology and Reasoning System for Access Policies to Software Services

Michael Kifer and Paul Fodor
michael.kifer@stonybrook.edu
pfodor@cs.stonybrook.edu

Rule-based models are commonly used to govern access rights to software services and systems. As administrators attempt to define increasingly sophisticated rules that take into account company policies, user context (e.g., mobile vs. fixed devices and geo-location), and the characteristics of the system, service or files, the rule sets can become complex. Moreover, administrators increasingly want to use available data analytics about users and the usage of the system's, services, and files. A security-domain-based reasoning system can help administrators design more robust, consistent rules and can help applications interpret those rules as intended by the administrators. The development and management of the basic building blocks for security mechanisms, which can be integrated into a wide variety of mission-critical information systems, also needs such a domain specific knowledge base system.

In this project, we develop a knowledge base infrastructure to help security domain experts to build ontologies of security knowledge components. This knowledge base infrastructure is intended to not only provide an interactive knowledge entry facility, but also to support reasoning capabilities to answer various types of queries about rule sets. It is also intended to help administrators construct more robust policy rules. *Figure 1.* (Center for Dynamic Data Analytics and CA Technologies)

Scalable Multilingual NLP through Deep Learning

Steven Skiena
skiena@cs.sunysb.edu

Accurate interpretation of natural language, including search queries, requires an understanding of how word meanings change with time and group. The word "gay" meant something quite different thirty years ago than it does today, a cycle of reinterpretation that can happen within days or even hours due to external events (storm names like Sandy or Katrina) or people in the news. Twitter hashtags are constantly recycled, the same string now with a different meaning. Regional slang implies that certain words mean different things in different places.

The key with all of these examples is capturing the semantics of what words actually mean. Recently re-introduced techniques in unsupervised feature learning make this possible, by acquiring common features for a specific language vocabulary from unlabeled text. These features, also known as distributed words representations (embeddings), have been used by us and other groups to build a unified NLP architecture that solved multiple tasks; part of speech (POS) tagging, named entity recognition (NER), semantic role labeling and chunking. We have built word embeddings for one hundred of world's most frequently spoken languages, using neural networks (auto-encoders) trained on each language's Wikipedia in an unsupervised setting, and shown that they capture surprisingly subtle features of language usage. We seek to build better systems for large-scale analysis of text streams regardless of language, and explore new methods for training word embeddings and their applications. (NSF)

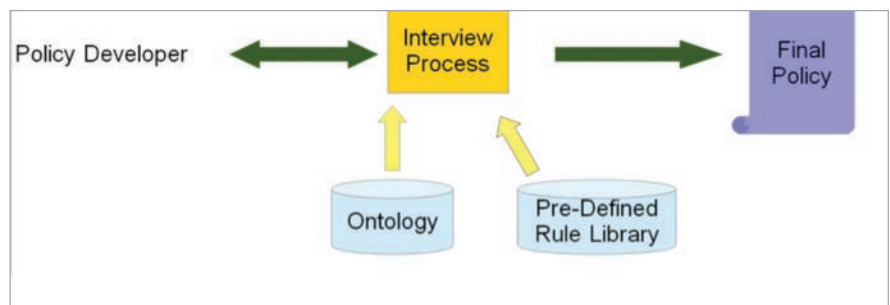


Fig.1

The ND-Scope: Visual Analytics for High-Dimensional Data

Klaus Mueller
mueller@cs.stonybrook.edu

The growth of digital data is tremendous. Any aspect of life and matter is being recorded and stored on cheap disks, either in the cloud, in businesses, or in research labs. We can now afford to explore very complex relationships with many variables playing a part. But for this data-driven research we need powerful tools that allow us to be creative, to sculpt this intricate insight from the raw block of data. High-quality visual feedback plays a decisive role here. And the more people participate in the effort the more progress we can make. We have created a framework and software package, called the ND-Scope, which incorporates various facilities for high-dimensional data exploration and reasoning with high-dimensional data. It couples powerful data analysis with artistic illustration techniques to help users to only show those aspects of the data they deem relevant. This clears the clutter of current data displays and so fosters better insight. (NSF)

Capturing Word Meanings by Language, Place, and Time

Steven Skiena
skiena@cs.sunysb.edu

Building multilingual processing systems remains challenging, particularly for resource-poor languages. But recent advancements in unsupervised feature learning present a promising alternative. Instead of relying on expert knowledge, these approaches employ automatically generated task-independent features, acquired for a specific language vocabulary by training on unlabeled text corpora.

These features, known as distributed word representations or embeddings, have been used by our lab and other groups to build unified NLP architectures that solve multiple tasks, including part of speech (POS)

tagging, named entity recognition (NER), semantic role labeling and chunking. We have built word embeddings for over one hundred of world's most widely-used languages, and shown that they capture surprisingly subtle features of language usage.

We will improve and generalize word embeddings to new problems and domains. Our research here includes:

(1) Better Word Embeddings — We will construct improved language representations based on our experiences building the first-generation of Polyglot embeddings. We will develop effective techniques to capture representative lexicons in the face of diverse language morphology, phrase co-locations, and major sources of out-of-vocabulary (OOV) words such as names/numbers/locations. The trick here is to do these things within an unsupervised learning framework, so we can build embeddings without labeled training data and domain experts.

(2) Applications of Word Embeddings to Multilingual NLP — We seek to apply our improved embeddings to construct a full stack of NLP processing tools -- including POS tagging, named entity recognition and classification, anaphora resolution, sentiment analysis, parsing, and even gross language translation -- for all 100+ languages. We anticipate these resources will continue to interest a broad research community. We also will use them ourselves, to establish a sentiment/volume trend analysis system to measure changes in the world condition, in its native tongues.

(3) New Domains for Word Embeddings— We believe that word embeddings have considerable promise in scientific and cultural domains far removed from traditional language processing. By appropriately training embeddings on scientific corpora like Pubmed/Medline abstracts, we can reduce the “bibliome” to concise numerical features about diverse entities like drugs and diseases, suitable for building models for a wide variety of applications. We can even use embeddings to study philosophical issues of what words really mean, by relating embeddings to traditional definitions found in dictionaries. (NSF)

Application of Parallel Computing to the Analysis of Next-generation Sequencing Data

Song Wu
songwu@ams.sunysb.edu

The focus of my lab is on developing and implementing novel statistical/bioinformatics methodologies for the analysis of large-scale genetic and genomic data.

In the past few years, high-throughput next generation sequencing (NGS) technology has quickly emerged as the dominating method in biomedical research, replacing the once-prevalent microarray technology. With NGS, many problems traditionally considered arduous or impossible have now become feasible. However, along with the rapid development of NGS technology, the data magnitude and complexity for analyzing NGS data far exceed the capacity and capability of traditional small-scale computing facilities, such as standalone workstations. On a separate note, massively parallel processing (MPP) systems have undergone dramatic development in the past decades. In spite of significant advances in both fields of NGS and MPP, cutting-edge research in applying supercomputing to NGS analysis is still at its infancy. Therefore, the enormous computational demand of the NGS data analysis along with the huge supply of processing power by MPP systems presents a unique and substantial opportunity for developing highly efficient analysis strategies for the NGS data. Specifically, in this project, we aim to address the two most fundamental problems in NGS data processing: (1) how to quickly align billions of short reads to a reference genome and (2) how to assemble billions of short reads into a genome in a timely manner. We are currently developing a scalable hierarchical multitasking algorithm for importing classical sequencing algorithms for these tasks to modern parallel computers.

The bioinformatics and parallel computing schemes developed in this project would be a valuable contribution to the genomic community for comparative genomics based on ultra-fast computing, which could potentially have a much broader impact,

especially on public health and clinical research. By tackling these challenges in handling the massive NGS data generated, our parallel software tool will offer unparalleled opportunities for targeted biological discovery and more accurate clinical diagnostics to accelerate personalized medication. (Stony Brook University)

Develop Novel Statistical Methods for Multi-loci Genetic Mapping

Song Wu

songwu@ams.sunysb.edu

The overall goal of this project is to develop novel statistical methods for detecting genetic risk factors by modeling multiple linked genetic loci in genome-wide association studies (GWAS). GWAS have achieved great success in moving forward our understanding of genetic contributions in many complex traits, however, current GWAS analyses experience a bottleneck of the so-called “missing heritability”, a phenomenon commonly observed in many GWAS that the detected genetic variants can explain only a small proportion of phenotypic heritability, while the majority part remains mysterious. Part of the reason is that the traditional GWAS analyses are mainly based on the single marker association, which is often inefficient in detecting small effect sizes and complicated gene interactions involved in complex traits. Therefore, developing more powerful genetic mapping methods has become increasingly important. Capitalizing on the observation that the effect of a causal genetic locus may be carried by its neighboring marker loci due to the structure of linkage disequilibrium blocks, methods that integrate multiple linked markers are presumably more powerful to detect the causal locus. However, the main difficulty for multi-loci genetic mapping is that when the number of linked SNPs is large, like 100 or more, high collinearity and model validity become serious issues. We aim to address these issues and extend current genetic mapping practice to incorporate the linkage information of multiple linked loci, and develop a series of statistical methods for multi-loci genetic mapping with a block of SNPs.

Most economically, biologically and clinically important traits, such as grain yield, reproductive behavior and cancer risk, are inherently complex. This project will greatly advance the discovery of novel genes and their interactions to facilitate the identification of drug targets to enhance public health, or to help animal and plant breeders to improve trait quality. (CEWIT)

Cache-efficient Parallel Algorithms for Fast Flexible Docking and Molecular Dynamics

Rezaul Chowdhury

rezaul@cs.stonybrook.edu

Proteins are one of the major structural and functional building blocks in our cells, and they often realize their functions through mutual interactions. The problem of computationally determining the relative transformation and conformation of two proteins that form a stable complex, reproducible in nature, is known as “protein-protein docking”. Docking is an important step towards understanding protein-protein interactions, and has applications in drug design, structure function analysis, and studying molecular assemblies. In spite of recent advancements the imaging of macromolecular complexes remains a difficult task, and the need for fast and robust computational approaches to predicting the structures of protein-protein interactions is growing.

We have already developed F2Dock -- a rigid-body protein-protein docking program based on cache-efficient multicore algorithms and data structures. F2Dock is the first docking program to use tunable approximations, and partly builds on our work on the fast estimation of the total free energy of bio-molecules in almost linear work and linear space. While F2Dock performs rigid-body docking using static octree-based data structures, the goal of this project is to develop cache- and communication-efficient dynamic data structures and algorithms for fast docking of flexible molecules and performing molecular dynamics on modern parallel machines. (CS/IACS Startup)

Efficient Resource-oblivious Parallel Algorithms

Rezaul Chowdhury

rezaul@cs.stonybrook.edu

Parallelism is now ubiquitous. From the tiny multicore smartphones in our pockets to the gigantic supercomputers with thousands of multicore and manycore compute nodes, parallel processing is supported at all scales. Unfortunately, however, writing correct parallel programs is hard, and making them efficient is even harder because of two apparently conflicting reasons: (1) abundance of resources, and (2) constraints on resources. To understand resource abundance consider a cluster of nodes each containing multicore processors and manycore coprocessors. In order to fully utilize these resources one must exploit internode distributed-memory parallelism as well as intranode shared-memory parallelism, and even inside a single node, keep both the processors and coprocessors busy with task parallel computations and data parallel computations, respectively. On the other hand, each level of the cluster comes with its own set of resource constraints, e.g., number of available compute nodes and internode communication bandwidth at the network level; number of available processing cores and coprocessors, and cache/RAM/external-memory sizes inside a compute node; and the number of cores and cache/memory sizes inside a coprocessor. Structures of the internode network and intranode memory-hierarchies also impose significant resource constraints. Finally, there could be a limited energy budget for the computation. Thus writing efficient parallel codes for state-of-the-art machines remains the job of a few experts who also must work hard to keep their skills up-to-date as the state-of-the-art changes so frequently. As a result, being able to design efficient “resource-oblivious” algorithms, i.e., algorithms that do not use the knowledge of resource parameters but still can perform with reasonable efficiency across machines, has become a highly desirable goal. We have already shown how to design efficient core- and cache-oblivious parallel algorithms for the multilevel cache-hierarchy of a multicore machine. The objective of this project is to improve known results

for multicores, and extend the notion of resource-obliviousness to real-world hybrid computing environments. If successful, this research will enable programmers to easily produce highly efficient codes for state-of-the-art parallel computing platforms. A wide variety of scientific applications — ranging across physics, biology, chemistry, energy, climate, mechanical and electrical engineering, finance, and other areas — will become easier to develop and maintain, benefiting these application areas, as well as society at large. (CS/IACS Startup)

FTFS: A Read/Write-optimized Fractal Tree File System

*Michael Bender, Rob Johnson, and Don Porter
rob@cs.stonybrook.edu*

Today's general-purpose file systems offer poor performance on microdata operations, such as file creation and destruction, small writes to large files, and metadata updates, such as inode or atime updates. Microdata operations are pervasive in file system workloads. This performance artifact is the result of a tacit, but demonstrably incorrect, assumption underlying file system designs: that one must tradeoff between data locality on disk and small-write performance.

The proposed research observes that a recently-discovered class of data structures called write-read-optimized (WRO) data structures, such as B ϵ -trees and fractal tree indexes, can bridge this gap. WRO data structures give comparable asymptotic behavior to a B-tree for queries and bulk updates, and support small updates with performance close to logging. Prior work demonstrates that these asymptotic benefits translate to real performance improvements—up to two orders of magnitude faster than a traditional B-tree for some operations.

The organizing principle of this research is the creation of FTFS, a general-purpose, Linux kernel implementation of a file system designed around WRO data structures. Unfortunately, operating systems have ossified certain assumptions about how file systems are designed, such as inducing extra lookups during update operations (called

cryptoreads). Cryptoreads cause update operations to block on lookups, thus throttling the faster updates that WRO data structures provide. The proposed work will investigate OS support for WRO data structures, as well as redesigning WRO data structures to support the operations of a fully-featured file system. (CEWIT)

High-performance Rule Engine for Intelligent Web Information Systems

*Michael Kifer
michael.kifer@stonybrook.edu*

Rule languages have been a popular choice for research on intelligent information systems for over three decades now. However, until recently, the lack of necessary computing power and of suitable networking infrastructure prevented wide adoption of this paradigm in software development. With the advent of the Semantic Web and following wide recognition of the limitations of its current standards, such as RDF and OWL, rule-based systems were put firmly back on the agenda. As proof of this renewed interest, a host of projects are expected to release new rule-based engines in the near future, and the World Wide Web Consortium is busy standardizing the necessary infrastructure, the Rule Interchange Format.

Most of the aforesaid projects, however, aim at a low-hanging fruit by augmenting existing Semantic Web languages with simple rule-based capabilities. Their aim is to provide for the most basic today's needs of the developers of semantic content on the Web. In contrast, our project targets future intelligent information systems, which require very expressive and high-performance rule-based knowledge programming languages. In this work, we develop new and integrate existing technologies ranging from higher-order, frame-based, and defeasible knowledge representation, to reasoning about processes, to truth maintenance, databases, and logic programming. (Vulcan Inc., NSF) Potential markets: Financial regulations, medical informatics, security

Eliminating the Data Ingestion Bottleneck in Big Data Applications

*Michael Bender and Rob Johnson
rob@cs.stonybrook.edu*

“Old-school” tools, such as RDBMS and file systems, are unable to keep up with the data ingestion rates of big data applications. The result has been the proliferation of specialty solutions that make compromises on generality, data freshness, result accuracy, or hardware cost. RDBMS and file systems have been orphaned at the largest scales by a mismatch between their underlying algorithms and hardware trends. At this point they run up to three orders of magnitude behind the capabilities of the hardware.

This proposal aims to deploy write-optimization techniques to yield a disruptive improvement in general-purpose analytical tools. By recouping two of the orders of magnitude available on modern hardware, this project will enable RDBMS and file systems to scale to orders-of-magnitude larger speeds and data sizes. This research will bring large-scale data analytics to a much broader, less specialized audience. A key component of this proposal is to rethink how storage-system hardware gets deployed for large data. By shifting the bottleneck on indexing away from I/Os per second to bandwidth, it is possible to replace small disks with large disks, at a savings of more than an order of magnitude per byte. Similarly, write-optimization can be used to reduce the number of relatively expensive solid state disks (SSDs), and to use them more effectively, for a further cost savings. A twin goal of this proposal is to increase the scale of data amenable for analysis while cutting storage-hardware costs.

To achieve these goals, this project will redesign storage mechanisms to avoid crypto-searches that can throttle write-optimized data structures down to the speed of B-trees, develop SSD-resident approximate membership query (AMQ) data structures, explore AMQs for range queries, design write-optimized data structures that minimize write-amplification when used on SSDs, and investigate I/O models for capacitated maintenance and online index creation. (NSF)

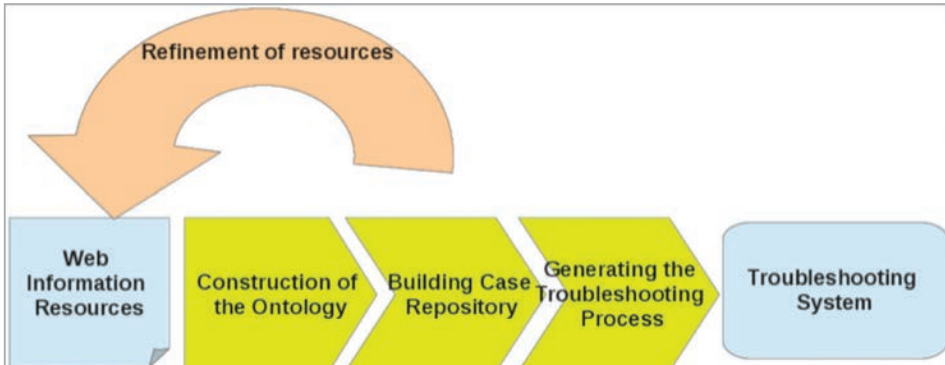


Fig.2 Diagnosis system development process

Automatic Generation of Virtual Troubleshooting Systems Using Ontologies

Paul Fodor, Michael Kifer, and Steve Greenspan (CA Technologies)
pfodor@cs.stonybrook.edu
michael.kifer@stonybrook.edu

Development of troubleshooting software is an attractive area of research for agent based system developers. In this project, we use ontologies extracted from textual representations to automatically construct a troubleshooting virtual expert. In our solution, we verify the information about the structure of the system extracted from the textual document, then generate a conversation with the user in order to identify the problem and recommend appropriate remedies. To illustrate the approach, we have built knowledge base for a simple use case and developed a special parser to generate conversations that can help the user solve software configuration problems. *Figure 2.* (Center for Dynamic Data Analytics and CA Technologies)

Visual Correlation Analysis of Numerical and Categorical Multivariate Data

Klaus Mueller
mueller@cs.stonybrook.edu

Correlation analysis can reveal complex relationships in multivariate data. It touches a diverse set of application domains: science, medicine, finance, business, and many more. Knowledge about correlations can enable the discovery of causal links and enables predictive analysis. However, as the number of variables grows, it can be difficult to gain a good understanding of the correlation landscape and important intricate relationships might be missed. To address this significant problem, we have devised an interactive visual interface to assist in correlation analysis. Our interface visualizes the data in two linked spaces: data space and variable (attribute) space. For the data space, a data-centric plot such as parallel coordinates of low-dimensional embedding, visualizes correlations among the variables in terms of their patterns. And for the variable space, a network-based correlation map directly visualizes the relationships among all variables—a multi-scale semantic zooming method provides scalability for high-dimensional and large data. (NSF, DOE)

High Performance Big Data Analytics with the User in the Loop

Klaus Mueller
mueller@cs.stonybrook.edu

In data mining and especially in big data, preprocessing consumes a large portion of the workflow, as much as 80-90%. Preprocessing includes data preparation, integration, cleaning, reduction, and transformation. As big data analysis can often be mission critical, preprocessing should be done expediently. The massively parallel architecture of GPUs offers an effective platform to accomplish high speed data processing. However, as GPU technology has been developing steadily and rapidly, users have trouble keeping up. And even if they do, the largeness of big data requires not just one GPU but multiples of these in conjunction with large memory. These needs are best addressed in a cloud-based platform. Our framework utilizes both a single GPU as well as a multi-GPU cloud server and it supports clustering, redundancy-based data decimation, outlier detection, data fusion, and so on. Using a progressive refinement scheme, users are given immediate visual feedback as partial results become available, allowing them to focus and steer the data processing. (NSF, DOE)

High Performance Computing for Medical Imaging

Klaus Mueller
mueller@cs.stonybrook.edu

We have devised solutions based on massively parallel commodity graphics hardware (GPUs) to accelerate compute-intensive applications in medical imaging, such as iterative algorithms used for low-dose X-ray Computed Tomography (CT). For example, we have pioneered a novel streaming CT framework that conceptualizes the reconstruction process as a steady flow of data across a computing pipeline, updating the reconstruction result immediately after the projections have been acquired. Using a single PC equipped with a single high-end commodity graphics board, our system is able to process clinically-sized projection data at speeds meeting and exceeding the typical flat-panel detector data production rates, enabling throughput rates of 100 projections/s for the reconstruction of clinically sized data volumes. Apart from enabling fast patient throughput and diagnosis, the streaming CT framework also represents an excellent platform for image-guided surgery and diagnostic imaging of transient phenomena. Multi-GPU solutions can achieve even higher throughput rates. A further, non-medical, application can be inside rapid security scanners for luggage and freight. (NIH)

A Visual Framework for Health Care Analytics

Klaus Mueller, IV Ramakrishnan, Rong Zhao, Asa Vicellio

High costs, lack of speed, non-intuitive interfaces, and inefficient, fragmented display of patient information have hindered the adoption of the Electronic Health Record (EHR, EMR). Critical factors inhibiting adoption of the EMR include the time spent by the health care providers in accessing and also documenting patient information during clinical encounters. We have devised a visual analytics system that unifies all EMR information fragments, such as current symptoms, history of pres-

ent illness, previous treatments, available data, current medications, past history, family history, and others into a single interactive visual framework. Based on this information the physician can then follow through a medical diagnostics chain that includes requests for further data, diagnosis, treatment, follow-up, and eventually a report of treatment outcome. As patients often have rather complex medical histories this visualization and visual analytics framework can offer large benefits for the navigation and reasoning with this information. (CEWIT)

Detecting Data Visualization Preferences Using Games

Lori Scarlatos
lori.scarlatos@stonybrook.edu

In visualizations of large multivariate data sets, discrete data can be effectively represented using glyphs. Glyphs have the advantage of allowing for rapid visual comparison, using differing visual dimensions to represent the different variables in the data. Some types of glyphs accommodate even more variables by using shape to represent the data. Yet the characteristics of these shapes may have underlying perceptual meanings. The purpose of this study was to determine whether certain shape characteristics are commonly viewed as good or bad. We conducted a study using two methods to gather data: a traditional survey, and a casual game. The results of this study strongly suggest that there are certain shape characteristics that are generally perceived as positive/negative, although they are not necessarily what might be expected. Research is continuing on how to effectively use games as survey instruments. (CEWIT)

Visual Analytics for Open-Ended Educational Multiplayer Games

Lori Scarlatos
lori.scarlatos@stonybrook.edu

Open-ended multiplayer games have the potential to promote collaborative, active, and inquiry-based learning. Yet the data gathered by such games can be massive, with a great number of variables to consider. Assessing the learning from this activity, without knowing which variables to focus on, can be a daunting task. We address this problem by presenting guidelines for visualizing data collected from an open-ended educational multiplayer game. An implementation of these guidelines, resulting in a system for visualizing data from a socio-scientific multiplayer game, was created. The system was pilot tested to determine how changes to the gameplay could provide more instructive feedback to the students and increase learning. Students playing the game also used the visual analytics system to better understand how the collective players' choices affect the outcome of the simulation. We are continuing this research by investigating other open-ended educational games, and how their data can be effectively visualized. (CEWIT)

Energy Choices: Using an Agent-Based Modeling Simulation and Game to Teach Socio-Scientific Topics

Lori Scarlatos
lori.scarlatos@stonybrook.edu

In our modern world, where science, technology and society are tightly interwoven, it is essential that all students be able to evaluate scientific evidence and make informed decisions. Energy Choices, an agent-based simulation with a multiplayer game interface, was developed as a learning tool that models the interdependencies between the energy choices that are

made, growth in local economies, and climate change on a global scale. We have pilot tested Energy Choices in two different settings, using two different modes of delivery. In our research, we are continuing development of the simulation - to increase the number of parameter choices - and the game - to make it more engaging for student players. We are also investigating the creation of a general interface framework that can be applied to other games built upon socio-scientific agent-based simulations. (CEWIT)

Compressive Sensing Approach to Intraoperative Mass Spectrometry for Tumor Margin Delineation

Allen Tannenbaum
Allen.tannenbaum@stonybrook.edu

Tumor resection is a key procedure for the treatment of a number of tumors including those of the brain and breast. The accurate determination of tumor boundaries is of crucial clinical importance since the surgical resection of tumors requires a delicate balance between a maximal resection of diseased tissue and minimal injury to the surrounding healthy tissue. Medical imaging techniques such as computed tomography and magnetic resonance imaging (CT and MR) are currently used in diagnosis as well as in image-guided surgical procedures, but studies show that these methods fail to accurately identify the full extent of malignant tumors and their microscopic infiltration. This highlights the need for a procedure that allows microscopic inspection of the tissue in real-time, and a framework to collectively analyze a limited set of local measurements and assist the surgeon in delineating the tumor margin intraoperatively.

Accordingly, at CEWIT we are developing the use of compressive sensing methods for the reliable and robust detection of tumor margins using a small number of measurements, in particular for the breast and the brain. These measurements will

consist of desorption electrospray ionization mass spectrometry (DESI-MS) imaging data. Compressive sampling/sensing has been utilized to show that a very large class of signals can be accurately (or in some cases exactly) reconstructed from far fewer samples than suggested by conventional sampling theory. Classical signal processing techniques lead to sufficient sampling by employing the band-limitedness of signals. In the compressive sensing approach, one defines sufficient sampling conditions based on the compressibility of a signal relative to a given dictionary designed for the problem at hand. Via models based upon the tumor cell concentrations derived from compressive sampling, our framework requires no assumptions about tumor shape or tumor cell density concentration. Further, we are developing certain statistical filtering techniques (e.g., particle and Kalman) that may be performed in the compressed domain to better approximate tumor boundaries.

Compared with scalar CT/MR imagery where each voxel holds a single value, and a diffusion weighted MR imagery where about 100 values are associated with a sample location, for DESI-MS imaging, roughly 10⁴ numbers are obtained at a single voxel. The treatment of such big data provides a great opportunity for data analysis and classification, but also poses a challenge for carrying out the analysis in a near real-time intraoperative fashion. Hence, the necessity of selecting and extracting the important information from such big data sets, and performing the computations in a sparse manner. In addition to reducing the computational load, identifying a sparse subset in the rich spectrum of DESI-MS would naturally link with the exploration of the biochemical significance of the MS data. Indeed, among the 10⁴ values at a single voxel, only less than 1000 render themselves as key chemical substances that differentiate the normal and cancerous tissue. Thus, determining the sparse spectral features provides indications which mass/charge ratios are significantly different among various tissues. (NIH)

Reality Deck – Immersive Gigapixel Display

Arie Kaufman (PI) and Klaus Mueller, Dimitris Samaras, Hong Qin (co-PIs)
ari@cs.stonybrook.edu

Large, interactive, high-resolution displays have been demonstrated as a valuable tool for the exploration of massive amounts of data. Due to their size, they allow for physical navigation (walking around in space) rather than virtual navigation (manipulating a virtual camera with a controller). Such displays were limited to 300 megapixels in aggregate resolution. Additionally, they presented themselves as a single planar surface, reducing the potential physical navigation benefits that users can enjoy. We have built the Reality Deck, which is the next-generation immersive, large, interactive, super resolution display. It is a unique 416 tiled display visualization environment that offers a total resolution of 1.5 billion pixels and 100dpi in a 4-wall horizontally immersive layout, while providing 20/20 visual acuity for the visualization space. It is the first facility of its kind and improves the resolution by a factor of 5 compared to the next largest tiled display wall and by a factor of 15 compared to other immersive environments, such as the CAVE. The high-resolution tiled LCD displays are driven by an efficient 20-node visualization cluster that utilizes four AMD FirePro V9800 GPUs per node with 6 displays connected to each GPU. The cluster provides aggregate 2.3 teraflops CPU performance, 220 teraflops GPU performance and 1.2 TB of memory. The Reality Deck is a one-of-a-kind facility, which serves as a platform for core visualization and usability research, systems-level research for enabling the visualization of new types of data (such as gigapixel video) and finally as an exploration platform for real-world visualization problems. We have implemented a number of interactive applications that leverage the high resolution and deal with a variety of large datasets, including gigapixel panoramic images, global GIS data, molecular models, medical data and art collections. (NSF)

NYRISE Visualization for Climate Simulation Data

Arie Kaufman
ari@cs.stonybrook.edu

The New York Resiliency Institute for Storms & Emergencies (NYRISE) is a cross-institutional and cross-disciplinary effort to enable the state of New York to be one step ahead of superstorm Sandy-like disasters. The goal of the Institute is to bring together the expertise of various researchers at Stony Brook (lead by researchers at the School of Marine and Atmospheric Sciences) and other universities in the fields of marine sciences, climatology, road network simulation, emergency response planning and visualization. We are focusing on the visualization aspect of the NYRISE and are developing novel and scalable visualization techniques for climate simulation data that will be utilized in the work of the institute. Such visualizations would merge the simulation data with underlying road networks, elevation and other GIS sources, enabling emergency planners to be better prepared for future storms and emergencies. These technologies will be deployable on traditional desktop computers but also scale up to gigapixel resolution facilities, such as the Reality Deck. (NYS)

Gigapixel Video Streaming

Arie Kaufman
ari@cs.stonybrook.edu

Gigapixel-resolution image capture is an active research area, with proof of concept devices already in the field. The AWARE family of cameras, developed at the Duke Imaging and Spectroscopy Program, can capture color gigapixel resolution images in a single “shot” by utilizing arrays of microcameras. The natural next step is the acquisition of gigapixel resolution video and the Reality Deck facility is a perfect fit for displaying and inspecting such content. However, the sheer bandwidth requirement for the delivery of such contents from a remote camera to the visualization facility

poses a challenging problem. Additionally, the microcamera data needs to be pre-processed (effectively “stitched” together) prior to the final gigapixel image being available, a process that currently is not real-time. We are developing solutions to this challenging technical problem. First, techniques for optimizing data transfers between the camera and the Reality Deck facility are being developed. Specifically, the position of a user within the Reality Deck can be used in order to adaptively select the image quality that should be delivered to each display (reducing in this way the overall bandwidth requirement for data streaming). Additionally, novel techniques for the reconstruction of these gigapixel images are being investigated. These techniques will leverage the parallel processing abilities of the graphics pipeline on modern GPUs, with a goal of achieving real-time reconstruction of the gigapixel image. (NSF, internal funding)

Natural Interaction with VR Environments


Arie Kaufman
ari@cs.stonybrook.edu

Traditional user interface devices, such as mice and keyboards, impose a “tethering” constrain to the user, anchoring him to a stationary working surface. Consequently, these interfaces do not apply themselves to a facility such as the Reality Deck, which encourages users to physically navigate the data by walking within the space. This physical navigation enables users to leverage their spatial memory when attacking data analysis tasks within an immersive visualization system and should not be discouraged by using tethered interaction devices. We are developing the next generation of user interfaces, fit for usage within large immersive environments. These user interfaces are hand-driven and gesture centric, allowing users to explore the visualization from any point within the facility, simply by wearing a pair of low-cost tracked gloves. (NSF, Samsung)

Immersive Virtual Colonoscopy

Arie Kaufman
ari@cs.stonybrook.edu

Virtual Colonoscopy (VC) is an established alternative to optical colonoscopy for colorectal cancer screening. A major advantage compared to the traditional procedure is that it is non-invasive and may be easier on the patient. There are also significant advantages in terms of completeness of the examination as VC provides complete coverage of the surface of the colon. We have developed an Immersive Virtual Colonoscopy (IVC) visualization system that further improves the efficiency of the examination through the use of virtual reality technologies in the Immersive Cabin (IC), also known as the CAVE, at Stony Brook University. The IC is a 5-sided enclosed visualization environment that uses pairs of high-resolution projectors to create stereoscopic images around the user. Stereoscopic rendering provides the radiologist with enhanced shape and depth perception, which improves the visual detection of polyps. The immersion further gives the ability to look around obstacles in the data, such as the haustral folds or sharp bends in the colon, simply by looking around in the IC. We have also developed a conformal visualization technique to support facilities that are only partially immersive (e.g., missing ceiling surface such as the IC) while preserving important 3D shapes in the data, such as the polyps. The combined effect of IVC is reduced examination time and improved sensitivity for the screening procedure. (NIH, NSF)



Interactive Immersive Visualization of Computed Microtomography Data

Arie Kaufman
ari@cs.stonybrook.edu

Computed microtomography (CMT) is widely used at synchrotron facilities for characterization of samples in many different fields such as energy, environment, materials, biomedicine, and plants. The experiments produce 3D data sets of several GB that need to be analyzed and visualized. With the BNL NSLS II facility that will become operational shortly, the resolution and data sizes will be significantly larger due to the increased x-ray intensities. This will require novel visualization tools to allow for the efficient analysis of the data. We have developed interactive visualization techniques for CMT data that leverage the immersive visualization facilities at Stony Brook University. The Immersive Cabin (IC), also known as the CAVE, is a 5-sided enclosed visualization environment that uses pairs of high-resolution projectors to create stereoscopic images around the user. This allows for interactive virtual flight through the pore structure of the sample, which can be used as a stage in the analysis framework. While the stereoscopic rendering is particularly suitable for studying intricate 3D structures, the pixel density in the IC (and CAVEs in general) is too low to handle larger datasets. We plan to implement the novel visualization techniques recently developed Reality Deck, which uses 416 high-resolution professional LCDs to provide more than 1.5 billion pixels of aggregate resolution in an immersive 4-sided and significantly larger working space with an order of magnitude increase in pixel density compared to the IC. (BNL)



Fusion of GIS and Procedural Modeling for Driving Simulation in New York

Arie Kaufman
ari@cs.stonybrook.edu

We have developed a mixed modeling pipeline for the creation of virtual 3D urban environments that approximate the real-world environments with sufficient fidelity to be suitable for driving simulations. One of the target applications is the simulation of eco-driving on the Manhattan to JFK route, where the area is too large to be modeled effectively by traditional modeling tools. Our approach combines accurate GIS data for large areas with procedural tools for the semi-automatic generation of the 3D data based on mathematical models. From sources such as elevation and aerial maps, transportation networks and building footprints, our pipeline produces compelling and varied models that approximate the target location. We have further developed visualization tools for the exploration of these models leveraging our virtual reality expertise and the Immersive Cabin and Reality Deck facilities at Stony Brook University. (Hyundai Motors, SUNY)

Medical Volume Rendering on the Gigapixel Reality Deck

Arie Kaufman
ari@cs.stonybrook.edu

We have developed a novel visualization system based on the reconstruction of high resolution and high frame-rate images from a multi-tiered stream of samples that are rendered framelessly as opposed to the traditional rendering on a pixel grid. The sample generator can be a GPU cluster that is separate from the display cluster, or even a cloud-based service. This decoupling of the rendering system from the display system is particularly suitable when dealing with very high resolution displays or expensive rendering algorithms, where the latency of generating complete frames may be prohibitively high for interactive applications. We specifically address the application of medical volumetric rendering on a gigapixel display, such as the Reality Deck at Stony Brook University, where the traditional visualization pipeline cannot produce complete images at interactive frame-rates. (NSF)

Lattice Simulation and Rendering

Arie Kaufman and Klaus Mueller
ari@cs.stonybrook.edu

In this project we have developed an innovative lattice of the computation and rendering domain, the hexagonal lattice, such as BCC (Body Centered Cubic) and FCC (Face Centered Cubic), for a variety of 3D and 4D wave phenomena simulations, such as lighting, acoustics and tomography. The hexagonal lattice discretizes the

domain into lattice sites, where every site is connected to 12 neighbors in contrast with the cubic cell in Cartesian grids. Photons, phonons, wavelets, or wavefronts are traced on the lattice links and interact at the lattice sites (e.g., scattering, absorption) using cellular automata like processes. The proposed work investigates and develops structures, constructions, sampling, algorithms and applications of the hexagonal lattices. The hexagonal lattice can be applied in and has the potential to revolutionize a variety of applications, such as simulation and rendering of light, acoustics and tomography, and has the potential to a very broad impact by extending the framework to many other particle based wave phenomena, such as those in particle physics. (NSF, CEWIT)

Plume Modeling Simulation and Visualization

*Arie Kaufman and Klaus Mueller
ari@cs.stonybrook.edu*

We have adopted a numerical method from computational fluid dynamics, the Lattice Boltzmann Method (LBM), for real-time simulation and visualization of flow and amorphous phenomena. Unlike other approaches, LBM discretizes the micro-physics of local interactions and can handle very complex boundary conditions, such as deep urban canyons, curved walls, indoors, and dynamic boundaries of moving objects. Due to its discrete nature, LBM lends itself to multi-resolution approaches, and its computational pattern is easily parallelizable. We have accelerated LBM on commodity graphics processing units (GPUs), achieving real-time or even accelerated real-time on a single GPU or on a GPU cluster. Another key innovation of LBM is its extension to support input from pervasive sensors, influencing the simulation so as to maintain its faithfulness to real-time live sensor readings. We have implemented a 3D urban navigation sys-

tem and have tested it with a 10-block GIS in the West Village of New York City, and with an 851-building area in Times Square of NYC. In addition to a pivotal application in simulation of airborne contaminants in urban environments, our approach will enable the development of other superior prediction simulation capabilities for physically accurate environmental modeling and disaster management, visual simulations for computer graphics, and has the potential to revolutionize the way scientists and engineers conduct their simulations. (NSF)

Visual Simulation of Thermal Fluid Dynamics in a Water Reactor

*Arie Kaufman
ari@cs.stonybrook.edu*

We have developed a simulation and visualization system for the critical application of analyzing the thermal fluid dynamics inside a pressurized water reactor of a nuclear power plant when cold water is injected into the reactor vessel for a possible thermal shock to the vessel. We employ a hybrid thermal lattice Boltzmann method (HTLBM), which has the advantages of ease of parallelization and ease of handling complex simulation boundaries. For efficient computation and storage of the irregular-shaped simulation domain, we classify the domain into nonempty and empty cells and apply a novel packing technique to organize the nonempty cells. We further implement this method on a GPU (graphics processing unit) cluster for acceleration.

It demonstrates the formation of cold-water plumes in the reactor vessel. We also develop a set of interactive visualization tools, such as side-view slices, 3D volume rendering, thermal layers rendering, and panorama rendering, which are provided to collectively visualize the structure and dynamics of the temperature field in the vessel. To the best of our knowledge, this is the first system that combines 3D simulation and visualization for analyzing thermal shock risk in a pressurized water reactor. (NRC, ISL)

Volumetric Segmentation of Computed Tomography Angiography

*Arie Kaufman, Rong Zhao and
Satya Sharma
ari@cs.stonybrook.edu*

The resolution and size of the Computed Tomography Angiography (CTA) data is rapidly growing and can easily reach up to 6GB per patient, especially when multi-slice CT scanners are used, such as the 320-slice CT scanner. However, the essential diagnostic information is contained only within the heart envelope and coronary vessel structures. This project develops robust volumetric segmentation algorithms of the whole heart envelope, as well as the coronary arteries for the purpose of data size reduction. Segmented data is significantly smaller in size and hence can be delivered, preferably compressed using either lossless or lossy compression, to the doctor at any location far away from the hospital for rapid preliminary diagnosis in the cases of acute chest pain, which is ultimately a lifesaving scenario. (CEWIT)

Automatic Spleen Segmentation for Non-Invasive Lymphoma Diagnosis

Arie Kaufman
ari@cs.stonybrook.edu

Diagnosis of spleen disorders, especially lymphoma, is a common and challenging issue in the clinical practice. Changes in the size of the spleen under the therapy provide a good evidence to assess the course of lymphoma malignancy. The initial step for the diagnosis is automatic spleen segmentation. In clinical setting it is still often performed manually, which is time consuming and highly observer-dependent. We focus on developing algorithms for robust automatic spleen segmentation and evaluate them against available ground truth results. Additionally, we automatically evaluate the volume of the spleen that is to be compared across repeated scans of the patient for evaluating changes and lymphoma diagnosis. (Stony Brook University)

Mobile-based Volume Rendering Pipeline for m-Health

Arie Kaufman
ari@cs.stonybrook.edu

Current mobile devices such as smart-phones or tablets have a number of unique characteristics that make them suitable platforms for medical applications. Their portability and always-on connectivity allows a medical doctor or health care provider to conduct the diagnostic process and follow up without being constrained to the workstation computer in the hospital facility.

We develop a pipeline for visualization of medical imaging, such as Computed Tomography (CT) and Magnetic Resonance Imaging (MRI) data that is common for a variety of applications, such as CT angiography, virtual colonoscopy, and brain imaging. In our work we concentrate on two main architectures for volumetric rendering of medical data: rendering of the data fully on the mobile device, when the data is already transmitted to the device, and a thin-client architecture, where the entire data resides on the remote server and the image is rendered on it and then streamed to the client mobile device. As mobile devices have been establishing new ways of interaction, we explore and develop 3D User Interfaces for interacting with the volume rendered visualization. These include touch-based interaction for improved exploration of the data. (Samsung)

Video Streaming for Interactive Visualization on Mobile Devices

Arie Kaufman
ari@cs.stonybrook.edu

We have developed a client-server streaming architecture to address the computational challenge of large data visualization on modern mobile devices, such as in medical imaging cases. In particular, tablets are used extensively in many fields, including the medical field, however, they do not provide sufficient performance to run applications such as volume rendering that traditionally require a workstation-class computer. The problem is further complicated by the ever-growing dataset sizes in medicine, driven by advancements in scanner and sensor technologies. In our system, the user's (e.g., doctor's) input on the tablet device is sent to a dedicated GPU cluster that executes the visualization algorithms, encodes the resulting video using high quality hardware video encoders, and streams the video back to the tablet. Full resolution high quality results can be achieved with minimal latency over a range of wireless technologies, enabling faster and more efficient diagnosis. (Samsung)



Volumetric Mesh Mapping

*Arie Kaufman and Xianfeng David Gu
ari@cs.stonybrook.edu*

With the rapid development of volumetric acquisition and computational technologies, vast quantities of volumetric datasets exist in numerous applications, such as industrial inspection and medical imaging. The demands for processing such volumes are pressing to analyze the topology and geometry, such as volumetric mapping to canonical structures, volumetric registration, volumetric feature extraction, geometric database indexing, volumetric parameterization, and so on. This project focuses on developing rigorous algorithms for computing the topology and geometry for general mesh volumes. Specifically, we have been developing computational algorithms for Ricci flows. On the other hand, it is highly desirable to map one or more volumes to a canonical domain, to support database indexing and volume registration.

We have built a concrete set of software tools for computing and visualizing the topology and geometric structures for mesh volumes, including volumetric parameterization, volumetric registration, volumetric mapping to canonical structures, fundamental groups computation, and topological and geometric feature extraction. Engineering, science, medicine, computer graphics, vision, scientific computing, and mathematics will directly benefit from these tools, the research and education. These tools can be further used in: (1) industry: in CAD/CAM/CFD simulation and analysis, non-destructive testing of scanned parts, reverse engineering, and large geometric database indexing; (2) medical imaging for volumetric registration and fusion, comparison, shape analysis, abnormality and cancer detection; (3) computational fields, for weather prediction, air flow around vehicles, and toxin prediction, using volumetric computed datasets; and (4) other fields for confocal volume microscopy for cellular research, seismology for earthquake prediction and gas and oil exploration, radar and underwater sonography for terrain mapping and object detection, both civilian and military. (NSF)

Conformal Mapping for Medical Imaging

*Arie Kaufman, Xianfeng David Gu,
Dimitris Samaras, and Wei Zhu
ari@cs.stonybrook.edu*

It is paramount in medical imaging to measure, compare, calibrate, register and analyze potentially deformed organ shapes with high accuracy and fidelity. However, this is extremely difficult due to the complicated shape of human organs. Different organs have different topologies and curvature distributions, and furthermore, the shape may deform due to disease progression, movement, imaging, surgery and treatment. We have used conformal geometry, a theoretically rigorous and practically efficient and robust method, to tackle this challenge. The broad objective of this project is to develop conformal geometry as a primary tool in the vast biomedical applications of medical imaging. We have been developing application of conformal surface flattening to a variety of organs, use of conformal mapping for volumetric parameterization, registration and fusion using conformal geometry, and statistical analysis and feature extraction using conformal geometry. The research design and methodology include developing and validating techniques to conformally flatten 3D organ surfaces to canonical parametric surfaces for colonic polyp detection, bladder cancer screening, and endovascular surgical planning for aortic aneurysm. We have further extended flattening to implement volumetric parameterization based on Ricci flow and then apply it to brain and colon structure segmentation, tumor evaluation, diffuse tensor field study. In addition, we have implemented shape registration and data fusion using a common canonical parameter domain. Brain data sets have been fused between and within subjects and modalities, as well as colon supine and prone have been registered for cancer screening. Finally, we have conducted statistical analysis and feature extraction using conformal geometry for drug addiction and Alzheimer's disease, where Fourier analysis on the canonical domains have transformed them to frequency domain. (NIH)

Automatic Detection of Colon Cancer

*Arie Kaufman
ari@cs.stonybrook.edu*

This project involves a method for computer aided detection (CAD) of colorectal cancer. The CAD pipeline automatically detects polyps while reducing the number of false positives (FPs). It integrates volume rendering techniques and conformal (i.e., angle preserving) colon flattening with texture and shape analysis. Using our colon flattening method, the CAD problem is greatly simplified by converting it from 3D into 2D. The colon is first digitally cleansed, segmented, and extracted from the CT dataset of the abdomen. The colon surface is then mapped to a 2D rectangle using conformal mapping. This flattened image is colored using a direct volume rendering of the 3D colon dataset with a translucent transfer function. Suspicious polyps are detected by applying a clustering method on the 2D flattened image to locate regions with irregularity and high density. The FPs are reduced by analyzing shape and texture features of the suspicious areas detected by the clustering step. Compared with shape-based methods, ours is much faster and much more efficient as it avoids computing curvature and other shape parameters for the whole colon wall. We tested our method with 88 datasets from NIH and Stony Brook University Hospital and found it to be 100% sensitive to adenomatous polyps with a low rate of FPs. The CAD results are seamlessly integrated into a virtual colonoscopy system, providing the physician with visual cues and likelihood indicators of areas likely to contain polyps. This, serving as a second reader, allows the physician to quickly inspect the suspicious areas and exploit the flattened colon view for easy navigation and bookmark placement. (NIH)

Reconstruction and Registration Framework for Endoscopic Videos

Arie Kaufman
ari@cs.stonybrook.edu

We develop a general reconstruction and registration framework (RnR) for endoscopic videos. We focus on and demonstrate RnR for optical colonoscopy (OC) videos and underline a novel way of visualizing endoscopic images and videos by using high-quality depth maps. To recover these depth maps from any given image captured from an endoscopic video, we employ a non-parametric machine learning technique which uses a dictionary of RGB images captured from virtual colonoscopy (VC) datasets and their corresponding depth patches. Given a video sequence, we can use shape-from-motion (SfM) and these inferred depth maps to recover a complete 3D structure defined in that sequence, along with the camera parameters. This allows us to register the recovered structure from the OC video with the one captured in the corresponding VC model, using quasi-conformal mapping. As a result, our framework can document the OC procedure using reconstruction and can localize the polyps found in VC, via registration. This framework can also be used for registration of two OCs in different times. In addition, we can study the quality of our reconstruction using the registration process, keeping the VC model as ground truth. Due to the non-rigidity of the colon, some components (e.g., folds) can sometimes have large deformations in successive frames which cannot be handled by SfM. To deal with these cases, our framework can also reconstruct a complete colon segment from a single colonoscopy image, assuming the intrinsic parameters of the camera are known. Our RnR general framework may be used for other endoscopic procedures (e.g., bronchoscopy, cystoscopy, etc.). (NIH)

Registration of Volumetric Prostate Scans using Curvature Flow

Arie Kaufman
ari@cs.stonybrook.edu

Radiological imaging of the prostate is becoming more popular among researchers and clinicians in searching for diseases, primarily cancer. Scans might be acquired with different equipment or at different times for prognosis monitoring, with patient movement between scans, resulting in multiple datasets that need to be registered. For these cases, we introduce a method for volumetric registration using curvature flow. Multiple prostate datasets are mapped to canonical solid spheres, which are in turn aligned and registered through the use of identified landmarks on or within the gland. Theoretical proof and experimental results show that our method produces homeomorphisms with feature constraints. We provide thorough validation of our method by registering prostate scans of the same patient in different orientations, from different days and using different modes of MRI. Our method also provides the foundation for a general group-wise registration using a standard reference, defined on the complex plane, for any input. In the present context, this can be used for registering as many scans as needed for a single patient or different patients on the basis of age, weight or even malignant and non-malignant attributes to study the differences in general population. Though we present this technique with a specific application to the prostate, it is generally applicable for volumetric registration problems. (NIH)

Brain Parcellation

Arie Kaufman
ari@cs.stonybrook.edu

Establishing correspondences across structural and functional brain images via labeling, or parcellation, is an important and challenging task in clinical neuroscience and cognitive psychology. A limitation with existing approaches is that they (i) are based on heuristic manual feature engineering, and (ii) assume the validity of the designed feature model. In contrast, we advocate a machine learning approach to automate brain parcellation. We are developing a novel shape-based approach for automatically labeling the anatomical features (folds, gyri, fundi) of the brain. We show that this approach can be considered a viable alternative to the atlas-based labeling approach. Also, this method can be used as an initialization process for atlas-based approach to provide for more accurate labeling of anatomical features. We create a machine learning framework to learn from the manually-labeled anatomical dataset, Mindboggle. This framework takes into account multiple features and shape measures to achieve up to 91% accuracy, without taking into account any spatial information. Moreover, this work is the first in the domain of 3D anatomical labeling of the brain. (internal funding)

Multimodal and Multivariate Visualization of the Brain

Arie Kaufman
ari@cs.stonybrook.edu

Current connectivity diagrams of human brain image data are either overly complex or overly simplistic. We attempt in this work to introduce simple yet accurate interactive visual representations of multiple brain image structures and connectivity among them. We map cortical surfaces extracted from human brain magnetic resonance image (MRI) data onto 2-D surfaces that preserve shape (angle), extent (area) and spatial (neighborhood) information for 2-D (circular disc or square) mapping or optimal angle preservation for 3D (spherical) mapping, break up these surfaces into separate patches, overlay shape information in these patches, and cluster functional and diffusion tractography MRI connections between pairs of these patches. The resulting visualizations are computationally easier to compute on and more visually intuitive to interact with than the original data, and facilitate simultaneous exploration of multiple data sets, modalities, and statistical maps. Additional specific contributions include two new techniques to map the brain surface on the sphere, and a novel insight regarding shape-based delineation of brain feature boundaries. (internal funding)

Meshless Point Cloud Registration by Conformal Mapping

Arie Kaufman and Xianfeng David Gu
ari@cs.stonybrook.edu

With the mass-production of home-use 3D sensors, such as the Microsoft Kinect, 3D Scanning is becoming more popular. However, due to limitations of this scanners (e.g., low resolution, 2.5D) advanced registration is required for generating higher quality models, such as the user's avatar. These scanners create 3D point cloud models, which we need to register in order to build the high quality model or to search for a match in a database of point clouds. The classical Iterative Close Point (ICP) method assumes that there is a large percentage of overlap between the two models to be registered. On the other hand, traditional conformal mapping requires meshes for boundary estimation and does not extend well to point clouds. We introduce a mesh-less conformal mapping framework capable of registering noisy point clouds captured by Kinect without any mesh data using model component segmentation and skeleton tracking. (internal funding)

Volumetric Shape Analysis

Arie Kaufman
ari@cs.stonybrook.edu

This project involves development of novel algorithms and techniques for geometric shape analysis in 3D volumetric data. Unlike in the case of 3D surfaces, the shape analysis in volumes is still in a nascent stage due to several challenges such as immense data sizes, increased complexities, etc. This project aims to analyze shapes in 3D volumes, including laying down theoretic foundations, designing computational algorithms, implementing software systems and applying to computer graphics and visualization fields. The aim is to use both deterministic methods such as heat kernel, Laplace-Beltrami operator and stochastic methods such as random walks, Markov Random Fields, Markov Chain Monte Carlo Methods. These methods are used to generate unique shape signatures and spectra which in turn are used in various applications such as matching, registration, tracking and analysis. The focus is also on developing fast and efficient computational methods using GPUs in order to facilitate real-time shape analysis. The results obtained using all the aforementioned methods are compared to provide an evaluative survey. In addition to computer graphics and visualization, the project results can be applied to other fields such as, computer vision, geometric modeling, medical imaging, manufacturing and architecture. Finally, as a byproduct, a general purpose volumetric geometric database is built up, with the set of geometric analysis tools, which allows volumetric geometric processing, searching, registration, fusion, comparison and classification. (NSF)



Ambienizer: Turning Digital Photos into Ambient Visualizations

Arie Kaufman and Klaus Mueller
ari@cs.stonybrook.edu

Although traditional information visualization systems are of great aid to experts, they are typically not appropriate for common citizens because these main stream users are often not familiar even with the most traditional techniques, such as bar charts or line charts, let alone advanced paradigms such as parallel coordinates. In this project we have developed a framework that allows users to turn any digital photo into an ambient visualization which seeks to convey information in a casual non-technical way using metaphors accessible to any common sense viewer. In our ambient visualization, called Ambienizer, the levels of the variable(s) to be monitored are mapped to suitable image processing operations whose effects reflect the level of these variables or retrieved images with certain level of features corresponding to value of data. Our approach is attractive since users can either choose any image of their preference and map any of the available visual effects to the variables they wish to monitor or view an image with preferred image features from their photo achieve or the web. Although it shows various types of data such as bank account, stock portfolio, weather information, and so on, we demonstrate the use of our system in the context of energy use monitoring in a household. (DOE)

Saliency-Aware Compression of Volumetric Data

Arie Kaufman
ari@cs.stonybrook.edu

With a constantly growing size of the volumetric data, the demand is increasing for its efficient compression, storage, transmission and visualization. Often only a certain part of the large volume is of interest to the user, and significant amount of the data contains no relevant information. In this project, we have developed a framework designed for processing and visualizing a large amount of volumetric data on a target device, possibly a mobile device with limited computational resources. Our system involves a 3D block-based saliency and transfer function guided compression scheme of volumetric data that achieves content and spatial scalability. Saliency of each volumetric region is computed from the coefficients of the 3D discrete cosine transform on the 3D volume. The identification of the salient volumetric blocks and weighting them by the transfer function help to schedule the ordered transmission of the regions in the volume. Additionally, our method is integrated into a resolution scalable coding scheme with integer wavelet transform of the image, so it allows the rendering of each significant region at a different resolution and even lossless reconstruction can be achieved. In addition, in order to further reduce the compressed data, our method provides an option to remove less salient blocks obtained by 3D seam carving from a 3D minimal cut before the compressing data. At the target device the received data is rendered progressively based on its saliency. Prioritized streaming of the data helps to achieve data reduction by rendering regions based on their saliency and disregarding less essential components. (CEWIT)

Virtual Dressing Room Project

Arie Kaufman
ari@cs.stonybrook.edu

Development of virtual reality head-mounted display kits, such as Oculus Rift, and their popularity among the consumers, calls for integration of them into the daily routines, for example shopping experiences. In our scenario the user can shop in any store in a fully immersive way, yet all in the comfort of her/his home. In this project we integrate immersive 3D headset to create fully immersive shopping environment that can be explored by the user similar to any real life department store. We structure the shopping environment as a physical store, with which user can interact by walking around, browsing items, and trying them on, all done in the first-person view. The user is provided with an opportunity to import an avatar, which resembles his/her body shape. Creation of this avatar is an intricate process and utilizes state-of-the-art algorithms developed in our lab. Alternatively, we provide a set of synthetic avatars with set of body measurements that correspond to traditional clothing sizes. In order to facilitate virtual fitting in our environment we develop custom geometric clothing models and methods of virtual fitting. The user is provided with a feedback from trying on such clothing as it is simulated according to its physical properties. The project is implemented on the basis of the 3D game engine to provide real time interaction and feedback to the user. (Omni-Scient)




Figure 3

3D Facial Recognition

Xianfeng Gu
gu@cs.stonybrook.edu

3D facial recognition has fundamental importance for homeland security. This project focuses on 3D facial surface recognition based on modern geometry and machine learning method. The human facial surfaces are captured using dynamic 3D camera based on phase shifting principle in real time with high resolution and high accuracy. The 3D facial surfaces are mapped onto the planar unit disk via Riemann mapping. The Riemannian metric on the original surface is encoded by the conformal factor on the disk. Then prominent geometric features are automatically selected by machine learning method. A diffeomorphism of the planar disk is computed by optimizing a special functional, which describes the elastic deformation and bending of the shape. This optimal registration also induces a distance between two 3D facial surfaces. By using the distance among the faces, we can compare their similarities. Different geometric features are weighted in order to improve the recognition rate, the weights are obtained by machine learning method automatically. Current system beats the state of the art. *Figure 3.* (NSF)

Geometric Manifold Theory for Higher-Dimensional Data Modeling, Analysis, and Visualization

Hong Qin
qin@cs.stonybrook.edu

Continuous functions defined over arbitrary manifold such as piecewise polynomial-centric splines enable compact data representation, analysis (especially quantitative analysis), simulation, and digital prototyping. They are relevant throughout the entire higher-dimensional, multi-attribute data processing pipeline including raw data acquisition and organization, visual data modeling and interactive manipulation, synthesis, analysis, and visualization. The main thrust of this project is to promote manifold functions as a new powerful data modeling, analysis, and simulation tool that not only continues to enjoy their popularity in conventional application fields such as geometric and shape representation, but also aims to broaden its widespread acceptance in a general higher-dimensional data modeling and analysis framework. Consequently, our goal is to transcend the traditional and current boundary of

splines' application domains and realize their full scientific potential. We systematically trailblaze a novel geometric manifold theory founded upon continuous polynomial representations, and to apply this mathematically-rigorous theory to both shape geometry and higher-dimensional, multi-attribute data modeling, analysis and visualization, with a special emphasis on visual computing applications. We explore new geometric manifold theory at the interface of differential geometry, numerical approximation theory, computational topology, and linear algebra. We conduct the comprehensive study of new and important geometric manifold theory that can enable the accurate and effective modeling of higher-dimensional, multi-attribute volumetric datasets which are of complicated geometry, arbitrary topology, and with rich geometric features. This research is expected to significantly improve the current theory and practice of information integration by potentially enabling a more accurate, more efficient, and easier-to-use data representation and structure in the future for processing geometric and higher-dimensional scientific datasets. (NSF)




Figure 4

Wireless Sensor Network Routing

Xianfeng Gu
gu@cs.stonybrook.edu

This project focuses on design efficient, reliable and secure routing strategies for wireless sensor network routing. Given a sensor network, according to their distance, one can build a planar graph. By graph embedding method, the graph can be assigned geometric coordinates, such that all cells are convex. The convexity of all cells guarantees the delivery using greedy routing method. The embedding can be transformed to achieve load balancing. Furthermore, the curvature of the network can be designed to avoid congestion. *Figure 4.* (NSF)

Multivariate Spline Theory, Algorithms, and Computational Techniques for Shape Modeling and Graphics Applications

Hong Qin
qin@cs.stonybrook.edu

Conventional multivariate splines (being defined as piecewise polynomials while satisfying certain continuity requirements) had found their mathematical root in approximation theory. In this digital era, splines are becoming ubiquitous, with widespread and deep penetration into many scientific and engineering fields including functional data analysis, reverse engineering, medical image analysis (e.g., image registration of non-rigid shapes), finite element analysis and simulation (e.g., numerical solutions of PDEs), data visualization (e.g., digital ocean), data manipulation and deformation (e.g., scalar fields and heterogeneous volumetric datasets for solid modeling), etc. In visual computing, traditional multivariate splines have been extremely popular in geometric modeling, graphics, and visualization, and they are commonly defined over an open, planar domain. Nevertheless, real-world volumetric objects are oftentimes of arbitrarily complicated topology. In addition, modern surface and volumetric scanning technologies enable data acquisition for not only geometric shapes but also multi-scale, multi-attribute material properties. In order to bridge the large gap between conventional spline

formulations and the strong demand to accurately and efficiently model acquired datasets towards quantitative analysis and finite element simulation, our research efforts center on the unexplored mathematical theory of manifold splines, that will enable popular spline schemes to effectively represent objects of arbitrary topology. We conduct the comprehensive study of new and important theoretical foundations that can transform the spline-centric representations to the accurate and effective modeling of surfaces of arbitrary topology. The primary and long-term thrust is to promote splines as a powerful data modeling, analysis, and simulation tool that not only continues to enjoy its popularity in conventional application fields such as geometric and shape modeling, but also aims to broaden its widespread acceptance in a general data modeling and analysis framework. The flexible and effective construction of splines defined over arbitrary manifold has immediate impact in geometric design, visual information processing, and graphics. It will provide a sound theoretical foundation for rapid product design and data analysis and have a potential to streamline the entire virtual prototyping processes for reverse engineering in the future by expediting data conversion from discrete samples to continuous geometry and spline-centric finite element models. (NSF)

Human Cortical Surface Morphological Study

Xianfeng Gu
gu@cs.stonybrook.edu

This project develops a general approach that uses conformal geometry to parameterize anatomical surfaces with complex (possibly branching) topology. Rather than evolve the surface geometry to a plane or sphere, we instead use the fact that all orientable surfaces are Riemann surfaces and admit conformal structures, which induce special curvilinear coordinate systems on the surfaces. Based on Riemann surface structure, we can then canonically partition the surface into patches. Each of these patches can be conformally mapped to a parallelogram. The resulting surface subdivision and the parameterizations of the components are intrinsic and stable. To illustrate the technique, we computed conformal structures for several types of anatomical surfaces in MRI scans of the brain, including the cortex, hippocampus, and lateral ventricles. We found that the resulting parameterizations were consistent across subjects, even for branching structures such as the ventricles, which are otherwise difficult to parameterize. Compare with other variational approaches based on surface inflation, our technique works on surfaces with arbitrary complexity while guaranteeing minimal distortion in the parameterization. It also offers a way to explicitly match landmark curves in anatomical surfaces such as the cortex, providing a surface-based framework to compare anatomy statistically and to generate grids on surfaces for PDE-based signal processing. *Figure 5.* (NIH)

Volumetric Modeling and Shape Design in Virtual Environments

Hong Qin
qin@cs.stonybrook.edu

IT-enabled engineering design is an innovative and iterative practice that consists of a variety of complex, challenging, and creative processes, ranging from conceptual design, interactive shape modeling, quantitative test/evaluation, rapid prototyping, manufacturing, assembly, to production. The essence of design is rapid, effective, and creative change towards optimal solutions, which can be accomplished through the iterative modification of numerous design parameters spanning across multi-dimensions. To ameliorate the entire CAE processes and enhance designers' creativity through IT-supported tools, we develop a creativity-enhancing virtual environment that can greatly facilitate human-computer interaction through physics-based modeling of real-world objects and force-enabled haptic sculpting. Our technical approach aims to broaden the accessibility of volumetric modeling of real-world objects by combining haptic sculpting with computational physics, thus offering novel interactive methodologies towards more creative and intuitive engineering design. Our research activities are concentrating on: (1) to bridge the large gap between physical objects and digital models, and (2) to improve the communication and interaction between human beings and computerized virtual environments (VEs), with a special emphasis on creativity enhancement. In

particular, our strategy aims to develop a haptics-enabled platform for volumetric data modeling, design, analysis, and relevant applications, which includes a suite of new, solid models for representing volumetric datasets as well as haptics-based physical interactions for manipulating and managing such representations. We take a unique and integrated approach that aims to incorporate within a single IT system disparate research thrusts that span volumetric, physical, and material modeling for novel data representation; geometric and physics-based algorithms for interaction, analysis, and visualization; and software tools for various applications. The confluence of modeling, visualization, and haptics is both imperative and uniquely valuable, not only for advancing the state of the knowledge in distinct disciplines but also for collectively enhancing current technologies with more advantages towards creativity enhancement. Ultimately, the anticipated outcome is a new IT platform for modeling, visualizing, and interacting with CAD-based, volumetric datasets that have arbitrary topology, volumetric structure, heterogeneous material, and dynamic behavior.



Figure 5

Shape Analysis with Teichmüller Shape Space

Xianfeng Gu
gu@cs.stonybrook.edu

Shape indexing, classification, and retrieval are fundamental problems in computer graphics. This work introduces a novel method for surface indexing and classification based on Teichmüller theory. Two surfaces are conformal equivalent, if there exists a bijective angle-preserving map between them.

The Teichmüller space for surfaces with the same topology is a finite dimensional manifold, where each point represents a conformal equivalence class, and the conformal map is homotopic to Identity. A curve in the Teichmüller space represents a deformation process from one class to the other.

In this work, we apply Teichmüller space coordinates as shape descriptors, which are succinct, discriminating and intrinsic, invariant under the rigid motions and scalings, insensitive to resolutions.

Furthermore, the method has solid theoretic foundation, and the computation of Teichmüller coordinates is practical, stable and efficient.

This work develops the algorithms for the Teichmüller coordinates of surfaces with arbitrary topologies. The coordinates which we will compute are conformal modules represented as the lengths of a special set of geodesics under this special metric. The metric can be obtained by the curvature flow algorithm, the geodesics can be calculated using algebraic topological method. We tested our method extensively for indexing and comparison of large surface databases with various topologies, geometries and resolutions. The experimental results show the efficacy and efficiency of the length coordinate of the Teichmüller space. *Figure 6.* (NSF)

EmoteControl: An Interactive Multimedia Stress Management System

Tony Scarlatos
tony@cs.stonybrook.edu

EmoteControl is an interactive multimedia application, developed in collaboration with Stony Brook University's Center for Prevention and Outreach, that provides exercises and simulations to relieve stress, while measuring brain wave activity for subsequent analysis and assessment. EmoteControl uses a consumer-grade electroencephalogram (EEG) headset, called the MindWave, to measure a user's level of concentration, anxiety, and overall brain activity. Four activities, including doodling and creating music sequences, provide exercises to help users relax. Three guided meditations, using video and narration, instruct users on breathing and meditation techniques. Users see their relative level of concentration or anxiety as a color-coded icon onscreen in real time. Their brain activity is recorded over several sessions for subsequent analysis and assessment. The purpose of this prototype is to study whether real time feedback helps users to manage their stress levels over time, and to examine which exercises or simulations have the greatest efficacy for stress relief. (CEWIT)




Figure 6

Shelterware: An Integrated Multi-Platform System to Automate the Services of the Smithtown Animal Shelter

Tony Scarlatos
tony@cs.stonybrook.edu

Shelterware was created to assist the Smithtown Animal Shelter (SAS), a publicly funded agency, update its system from an antiquated paper-based system to a modern digital system, so that the shelter's limited staff resources could be better optimized. Presently, the shelter houses 120 to 150 ready-to-adopt animals at any given time. The Shelterware system will allow SAS staff to manage adoption and volunteer information through a web-based interface, and to find and identify lost animals. The Shelterware system melds various state of the art client-server technologies from both the web and from mobile devices to automate the services provided by SAS. We have developed an iPad application that allows volunteers and staff to process incoming animals at the shelter by taking photographs of the animals, entering data about the animals, and automatically generating QR codes for identification, all of which are uploaded to a database. This database will generate a catalog of pets for adoption, which can be accessed by potential adopters on the web. Furthermore, potential adopters will now be able to fill out an online application form. Additionally, we have developed a field application, which allows staff to use iPhones for the identification of lost animals. Staff can scan QR codes attached to the collars of lost animals, or photograph the animals. This data will then be uploaded to a database with a time and date stamp along with the lost animals' location data. (CEWIT)



Figure 7

Conformal Wasserstein Shape Space

Xianfeng Gu
gu@cs.stonybrook.edu

Surface based 3D shape analysis plays a fundamental role in computer vision and medical imaging. This project uses optimal mass transportation maps for modeling shape space. The computation of the optimal mass transport map is based on Monge-Brenier theory, in comparison to the conventional method based on Monge-Kantorovich theory, this method significantly improves the efficiency by reducing computational complexity.

This project develops the framework of Conformal Wasserstein Shape Space. Given a Riemannian manifold, the space of all probability measures on it is the Wasserstein space. The cost of an optimal mass transportation between two measures is the Wasserstein distance, which endows a Riemannian metric of the Wasserstein space. In this work, all metric surfaces with the disk topology are mapped to the unit planar disk by a conformal mapping, which pushes the area element on the surface to a probability measure on the

disk. This gives a map from the shape space of all topological disks with metrics to the Wasserstein space of the disk, the pull back Wasserstein metric equips the shape space with a Riemannian metric. The conformal Wasserstein shape space framework is applied for expression classification.

3D human faces with dynamic expressions are captured using real time structure light 3D camera based on phase shifting principle. The facial surfaces with expressions form a shape space. By using Wasserstein distance, they can be clustered and classified, each cluster represents an expression. This technique is applied for expression analysis and autism diagnosis. *Figure 7.* (US Air Force)

BrainRank: A Multimedia Game to Increase Fluid IQ

Tony Scarlatos
tony@cs.stonybrook.edu

The Dual N-Back game is a memory game of recalling sequences of audio and visual stimuli (Dual), with the increasing difficulty of remembering sequences several steps ago (N-Back). Recent studies have confirmed that persistent practice with Dual N-Back games increases working memory capacity, and improves Fluid IQ. But most versions of the game are tedious for users, and they quit playing before much benefit to their working memory has occurred. BrainRank will study the impact of a kinesthetic version of the Dual N-Back game, using a dance pad as the input device for an application that records performance and time on task. Users can then take a version of Raven's Progressive Matrices test to gauge any improvement in their Fluid IQ scores over time. The study, to be conducted through ES BOCES, will measure the impact of physical activity on cognitive function in elementary school students. Additionally, BrainRank introduces "gamification" techniques to the design of the software, such as leader boards and accomplishment badges, to evaluate if such motivators increase users' time on task. (CEWIT)


Integrating Humans and Computers for Image and Video Understanding

Dimitris Samaras
samaras@cs.stonybrook.edu

Visual imagery is powerful; it is a transformative force on a global scale. It has fueled the social reforms that are sweeping across the Middle East and Africa, and is sparking public debate in this country over questions of personal privacy. It is also creating unparalleled opportunities for people across the planet to exchange information, communicate ideas, and to collaborate in ways that would not have been possible even a few short years ago. Part of the power of visual imagery comes from the staggering number of images and videos available to the public via social media and community photo websites. With this explosive growth, however, comes the problem of searching for relevant content. Most search engines largely ignore the actual visual content of images, relying almost exclusively on associated text, which is often insufficient. In addition to the growing ubiquitousness of web imagery, we also notice another kind of visual phenomena, the proliferation of cameras viewing the user, from the ever present webcams peering out at us from our laptops, to cell phone cameras carried in our pockets wherever we go. This record of a user's viewing behavior, particularly of their eye, head, body movements, or description, could provide enormous insight into how people interact with images or video, and inform applications like image retrieval. In this project our research goals are:

- 1) behavioral experiments to better understand the relationship between human viewers and imagery and
- 2) development of a human-computer collaborative system for image and video understanding, including subject-verb-object annotations,

- 3) implementation of retrieval, collection organization, and real world applications using our collaborative models. People are usually the end consumers of visual imagery. Understanding what people recognize, attend to, or describe about an image is therefore a necessary goal for general image understanding. Toward this goal, we first address how people view and narrate images through behavioral experiments aimed at discovering the relationships between gaze, description, and imagery. Second, our methods integrate human input cues – gaze or description – with object and action recognition algorithms from computer vision to better align models of image understanding with how people interpret imagery. Our underlying belief is that humans and computers provide complimentary sources of information regarding the content of images and video. Computer vision algorithms can provide automatic indications of content through detection and recognition algorithms. These methods can inform estimates of "what" might be "where" in visual imagery, but will always be noisy. In contrast, humans can provide: passive indications of content through gaze patterns - where people look in images or video, or active indications of content through annotations. Finally, we expect that our proposed methods for human-computer collaborative applications will enable improved systems for search, organization, and interaction. (NSF).



Machine Learning for the Analysis of fMRI Images

Dimitris Samaras
samaras@cs.stonybrook.edu

Analysis of functional Magnetic Resonance Images has allowed numerous insights in the way the brain functions. A large amount of data is being collected in numerous studies, however significant challenges remain towards automatic analysis of such data. These challenges include high levels of noise in the acquired signal, difficulties in registering scans between different subjects and different modalities, relatively small number of subjects per study and differences in protocol design and scanning equipment between studies. We hypothesize that unique patterns of variability in brain function can assist in identification of brain mechanisms rooted in neuroscientific prior knowledge. Such patterns will increase our understanding of brain connectivity and circuitry as we move iteratively between a-priori and exploratory means of describing functional brain circuits. We propose an integrated machine learning framework for the joint exploration of spatial, temporal and functional information for the analysis of fMRI signals, thus allowing the testing of hypotheses and development of applications that are not supported by traditional analysis methods. While our major focus is drug addiction and disorders of inhibitory control, we are testing our methods in datasets for other disorders such as major depression and autism. In recent work in my group with fMRI data we have demonstrated that it is possible to classify different groups of human subjects performing the same tasks based on the observed 3D fMRI BOLD images, through

discovery of the appropriate features that capture the most discriminative differences in activation levels. We have validated our technique in multiple fMRI data sets. In prediction of behavior based on neuro-imaging data, we search for the features which can unambiguously relate brain function with behavioral measurements. At the same time we have been exploring functional connectivity between brain regions by searching for conditional probabilistic dependencies between such regions, described by Gaussian Graphical Models, suitable to high-dimensional datasets. The method has guaranteed global minima, and does not require a-priori brain segmentation or selection of active regions. As a structure learning technique, the effect of confounding variables of brain region activations is removed, giving a clearer interpretation of possible dependencies between brain regions. Structure learning approaches will be useful in our quest to integrate spatiotemporally heterogeneous data sets such as MRI and PET. In immediate future work we plan to explore how domain knowledge and assumptions about common underlying neuropsychological processes can be used as constraints that allow us to combine similar (but not equal) features or dependencies across related datasets in multitask learning frameworks. Leveraging such related datasets will improve the stability of machine learning methods that are often encumbered by the difficulties posed by the high dimensionality in brain imaging datasets. In order to study larger cohorts of subjects we are also investigating resting state data for a number of disorders. (NIH)

Learning Models for Illumination, Shadows and Shading

Dimitris Samaras
samaras@cs.stonybrook.edu

Pictures result from the complex interactions between the geometry, illumination and materials present in a scene. Estimating any of these factors individually from a single image leads to severely under-constrained problems. Having a rough geometric model of the scene has allowed us to jointly estimate the illumination and the cast shadows of the scene. We solve the joint estimation using graphical models and novel shadow cues such as the bright channel. We further develop our method so as to be able to perform inference on not only illumination and cast shadows but also on the scene's geometry, thus refining the initial rough geometrical model. Scene understanding methods that want to account for illumination effects require reasonable shadow estimates. General shadow detection from a single image is a challenging problem, especially when dealing with consumer grade pictures (or pictures from the internet). We investigate methods for single image shadow detection based on learning the appearance of shadows from labelled data sets. Our methods are the state of the art in shadow detection. In our work we take advantage of domain knowledge in illumination and image formation for feature design, as well as extensive use of machine learning techniques such as graphical models (MRFs, CRFs) and classifiers (SVMs). (NSF)

Platform Comparison via Errors in Variables Models with or without Replicates

Wei Zhu and Jinmiao Fu
Wei.Zhu@stonybrook.edu

Platform or instrument comparison is a critical task in many lines of research and industry. In biotechnology, for example, with the rapid development gene sequencing platforms – we now have the first and the next generation sequencing platforms, each with several brands manufactured by different companies, co-sharing the market – at the same time, the third generation sequencing method built upon the sequencing of a single molecule of DNA, has already emerged (<http://genomena.com/technology-how-we-look-at-dna/third-generation-genome-sequencing/>). An accompanying critical task for the statisticians is how to best compare and combine results from different platforms. Previously, we have demonstrated the advantages of the errors-in-variable (EIV) model as an optimal instrument comparison and calibration device. However, one limitation to the traditional EIV modeling approach is its reliance on the availability of repeated measures of the same sample. Such replicates can be expensive and at times unattainable. Two methods by Wald (1940) and Kukush (2005) respectively are applicable for estimating the EIV model in the absence of replicates -- both relying heavily on the pre-cluster of data points. In this work, we aim to combine and improve these two methods through a better clustering strategy. (Simons Foundation, NIH)

Biomarker Agreement and Integration Analyses across Measurement Platforms

Wei Zhu and Jinmiao Fu
Wei.Zhu@stonybrook.edu

With the rapid development of biotechnology, an increasing number of platforms for measuring gene expression levels are co-sharing the market – with the older technology such as the gene microarray being phased out, while the newer ones based on the RNA sequencing (RNAseq) being brought in, generation after generation. A crucial question to the entire biomedical community is whether biomarkers detected through these different platforms are consistent or not? In this work, we present a theoretical framework on biomarker consistency study based on the errors in variable (EIV) model. First, we calibrate the measurements between two platforms through an EIV model featuring indices for the constant bias and the proportional bias. Subsequently we demonstrate how different biomarker detection algorithms including the fundamental fold change and Z - test, T - test, will be influenced by such biases. Finally, we discuss strategies to combine measurements from different platforms for better biomarker detection. (Simons Foundation, NIH)

Comparing the RNA-seq and Microarray Platforms with Generalized Linear EIV Model

Ellen Li, Wei Zhu, and Yuanhao Zhang
Ellen.Li@stonybrook.edu;
Wei.Zhu@stonybrook.edu

The discovery of gene biomarkers and related pathways is critical in the studies of disease and therapeutic treatment. The rapid evolution of modern biotechnology has generated several waves of measurement platforms – with some phasing out, while others co-sharing the market. Two front runners are the more traditional gene microarray technology and the newly arrived RNA-Sequencing (RNA-Seq) platform. An intensive literature review revealed that the prevalent statistical method for the comparison of the Microarray and the RNA-Seq platforms is the Pearson product moment correlation coefficient. However, the Pearson correlation is unable to provide a calibration formula to convert the expression levels between the platforms. It also fails to account for the fixed and proportional biases between the two platforms. To fill this void, we have developed a method based on the generalized linear errors-in-variable (EIV) model that is able to provide both a calibration equation as well as a statistical measure of the two biases. (Simons Foundation, NIH)



Bootstrap-based RANOVA for microRNA Panel Data Analysis

Ellen Li, Wei Zhu, Yuanhao Zhang,
Jennie Williams
Ellen.Li@stonybrook.edu;
Wei.Zhu@stonybrook.edu

Gene microarray is widely-adopted to identify differentially microRNAs (miRNAs) across different levels of potential prognostic factors such as phenotype, genotype, race, gender, etc. The significance analysis of microarray package (SAM) is arguably the most common analysis procedure which has been included as part of the Agilent miRNA analysis pipeline. However, SAM can only analyze data with one prognostic factor. When the experimental design is factorial with multiple factors, Zhou and Wong (2011) provided a nice non-parametric bootstrap ANOVA pipeline. Both this and SAM are nonparametric in nature as the miRNA data are often not normal. When the experimental design is in a panel data format containing both the within- subject factor (repeated measures) and the between-subject factors, however, the most suitable method to compare group means is the repeated measures ANOVA (RANOVA) rather than the factorial ANOVA. To our knowledge, the available methods are either parametric (Li, 2004; Conesa et al., 2006) or only one-way repeated measures analysis without between-subject factors (Elbakry, 2012). To rectify the lack of a non-parametric RANOVA with multiple test correction via FDR for microarray panel data analysis, we have developed a novel analysis pipeline combining the modern RMANOVA (Seco et al., 2006) and the empirical FDR (Zhou and Wong, 2011) for multiple test comparison, and implemented which in an efficient computational paradigm. In addition, our method is also applicable to the traditional messenger RNA (mRNA) panel data analysis. (Simons Foundation, NIH)

Integrative Multi-Scale Biomedical Image Analysis

Joel Saltz, Tahsin Kurc, and
Allen Tannenbaum
joel.saltz@stonybrookmedicine.com

Understanding the interplay between morphology and molecular mechanisms is central to the success of research studies in practically every major disease. At disease onset and through disease progression changes occur in tissue structure and function at multiple spatial and temporal scales. Quantitative study of disease through synthesis of multi-scale integrated structure, function, and molecular information has tremendous potential to significantly enhance disease diagnosis and therapeutic strategies in a wide range of disease types.

The need for deep, quantitative understanding of biomedical systems is crucial to unravel the underlying mechanisms and pathophysiology exhibited in cancer biology. Our goal is to develop detailed, models of multiscale tissue morphology and molecular characterization to 1) enable high precision predictions of tumor change and evolution, and 2) develop an ability to anticipate treatment response to combined regimens consisting of radiation therapy, traditional chemotherapy and targeted therapies. This research is an interplay between systems biology, multi-scale image analysis and population

health, combining “omics”, Pathology, Radiology and patient outcome. We leverage data from integrative patient studies including the Cancer Genome Atlas study, from the Stony Brook Cancer Center and from several other collaborating New York area cancer centers including Yale and Rutgers. Analyses of brain tumors using whole slide tissue images, for example, have shown that correlations exist between the morphological characterization of micro-anatomic structures and clinical outcome and genomic profiles.

Our research and development is in three inter-related areas: (1) Analysis methods capable of extracting descriptions of tissue morphology from massive 2-dimensional, 3-dimensional, and temporal whole slide tissue imaging datasets for multi-scale morphological characterization of tissue; (2) Analytical methods and tools to explore combined imaging, molecular, and clinical features to reveal the interplay between tissue structure and genomic and clinical profiles; and (3) a high performance middleware platform, which can take advantage of high end computing systems and emerging architectures with multi-core CPUs and hardware accelerators, to process and manage very large volumes of high resolution microscopy imaging datasets. (NCI, NLM, NIH)

Feature-based Exploration of Extremely Large Spatio-Temporal Scientific Datasets

Tahsin Kurc and Joel Saltz
Tahsin.Kurc@stonybrook.edu

Advances in sensor technologies make it possible to rapidly collect vast quantities of low-dimensional, spatio-temporal datasets, in which data elements are associated with coordinates in a multi-dimensional space with low-dimensionality and potentially obtained at multiple time steps. Analysis and characterization of features (e.g., spatial structures, their properties, function of the properties over space and time) in these datasets are important in many scientific

domains, including weather prediction and climate modeling, earth systems science, biomedicine, and materials science. In order to fully exploit the potential of spatio-temporal sensor datasets in scientific research, high performance computing capabilities are needed to rapidly extract and classify various features from large volumes of data, ranging from multiple terabytes to petabytes, using data and computation intensive analysis pipelines. The primary objective of our research is to develop and evaluate a suite of novel data and processing abstractions and optimizations within an integrated framework to enable analysis of extremely large low-dimensional spatio-temporal data for scientific, biomedical, and clinical studies. The methods and software systems developed in this research support novel data representations and runtime optimizations to be able to 1) ingest and manage large volume of diverse data sets, 2) stage datasets using resources in the data path, such as clusters and GPU accelerators, and 3) rapidly process datasets using a repertoire of analysis operations. The research also investigates the interplay between spatio-temporal data analysis applications, middleware software, and hardware configurations, targeting high end parallel machines/clusters containing hybrid multicore CPU-GPU nodes; extreme scale machines which consist of hundreds of thousands of CPU cores; and machines with deep memory and storage hierarchies linked to machines that acquire and aggregate data. (NLM, DOE, NIH)



Figure 8

Home Sleep Monitoring Using Bluetooth LE Biosensors to Facilitate Sleep Health

Helmut H. Strey
Helmut.Strey@stonybrook.edu

An estimated 50–70 million adults in the United States have chronic sleep and wakefulness disorders. Sleep difficulties, some of which are preventable, are associated with chronic diseases, mental disorders, health-risk behaviors, limitations of daily functioning, injury, and mortality. The goal of this project is to enable patients with sleep disorders to improve their sleep health by providing long-term sleep quality and quantity feedback while maintaining the accuracy of a clinical sleep study. We are developing a Bluetooth low-energy wireless sensor suite equivalent to a Type-3 sleep study that is suitable for home use. The quality of our measurements is verified by side-by-side testing in a clinical sleep laboratory. Furthermore, we are developing mobile device (iPhone, iPad, iPod) software that captures and analyzes the sensor data to provide feedback to the user about sleep quantity and quality. Combined with professional questionnaires before and after sleep, the patients will be able to correlate their life-style to sleep quality and through this achieve long-term changes in behavior to improve their sleep health. Optionally, the patients can link their data to social media (Facebook, Twitter) to improve their motivation and to add a competitive edge to their efforts. Similar approaches in the fitness area (Flexbit, Fitbit, Polar Heart Rate Monitor, etc.) have been tremendously successful in improving user motivation and ultimately leading to positive behavioral changes.

Figure 8.

Unified Statistical Models for Integrative Omics

Pei Fen Kuan
peifen.kuan@stonybrook.edu

Large-scale epigenomics public resources such as the Cancer Genome Atlas (TCGA) and NIH Roadmap Epigenomics Mapping Consortium offer an exciting research direction to integrate analysis results from multiple omics experiments generated from different technologies/platforms (tiling arrays, next generation sequencing). This is an important step to unlocking the secrets of complex biological systems. A couple of examples include studying gene-gene and gene-transcription factor interactions via co-expression and regulatory network analysis; and integrating mutation, copy number and gene expression data in identifying candidate driver genes in cancer. Such analyses usually require large data sets and implicitly assume independence among the different data sets. However, in most cases, these data sets are generated from biologically dependent experiments (e.g., multiple time series gene expression data) and ignoring this fact could lead to incorrect inferences on biologically significant pathways. In addition, most of the existing softwares work by first analyzing each data source separately, followed by declaring a set of candidate regions based on user-specified error rate threshold, and finally constructing regulatory network from these reduced data. These approaches have multiple pitfalls including difficulties in overall error rate control and information loss due to data reduction. Our lab aims to develop unified statistical methodologies for integrating omics data to identify functionally associated regulators and improve error rate control, by incorporating the dependence among these experiments at the raw data level. Such data integration encompasses both combining multiple omics data (1) of similar types across studies and (2) of different types across and within studies. Softwares implementing our proposed methodologies will be developed as Bioconductor/R packages. (CEWIT)

Mobile Healthcare Device Platform

Wei Lin
wei.lin@stonybrook.edu

Our research focuses on the development of wearable healthcare devices. We've developed a hybrid platform that utilizes ARM Cortex-M series processor as the embedded system and the open source FreeRTOS as its real time operating system. The ARM processors are designed for mobile applications with low power consumption. The FreeRTOS provides a small foot print kernel for real time, multitask applications. In order to integrate the FreeRTOS with the selected processor, a set of drivers were designed to bridge the FreeRTOS and the mobile processor. They provide a uniformed software interface that allows the application code running on the FreeRTOS to easily control the hardware resources such as analog digital converter, I2C bus and universal asynchronous receiver/transmitter (UART). Application code can be divided into small modules called tasks for design simplicity. The software architecture maximizes the code reusability and enables the quick switch of hardware design with little impact on the existing code. Its flexibility is very attractive for the development of mobile healthcare applications. This platform has been adopted in the mobile infant monitor device for the prevention of infant sudden death syndrome in collaboration with the School of Health Technology and Management at Stony Brook University. The device

collects the vital signals from infants and performs on-site analysis. It will send instant warning to mobile device of parents and caregivers if abnormalities are observed. The platform has significantly reduced the develop time of the prototype. *Figure 9.*

Study of How IT Affects Patient Access, Quality of Care and Administrative Costs in a Newly Formed ACO

Eugene Feinberg
eugene.feinberg@stonybrook.edu

This project studies efficiency of implementing particular information technology solutions for medical offices and hospitals in order to achieve the following goals:

- 1) Enable secure, digital Patient-to-Provider and Provider-to-Provider communications that are traceable and reportable.
- 2) Reduce phone call volume in medical offices and hospitals.
- 3) Achieve cost savings and productivity improvements. (ClipboardMD)

Developing Computer Games and Virtual Reality Environments for Use in Physical Rehabilitation

Rong Zhao
rong.zhao@stonybrook.edu

Computer augmented rehabilitation combines interactive computer technology with various devices and apparatus to enhance physical rehabilitation. The degree of interaction can range from simple motion biofeedback to virtual reality (VR) simulations. We propose to have applications that are rich with rehabilitation strategies, protocols and applications. The two areas we focus on with the greatest market potential are 1) balance training and fall prevention solutions, and 2) gait training solutions. Innovative computer games and virtual reality environments are developed and integrated into rehabilitation processes to deliver high quality care to the aging population and injured veterans returning from the two wars. We partner with Biodex Medical Systems, a leading provider of a wide range of physical rehabilitation and training systems, which offers specific products for balance rehabilitation that provided an excellent platform for interactive gaming. With these products the patient's motion provides a joystick output that can "work" or navigate based on the patient's motion through an appropriate program. Biodex also produces an instrumented treadmills that could interface for "walking" simulations as well as a computer based robotic system for upper and lower extremities that would benefit from virtual, interactive computer games. Various competition-type training programs are developed to increase the motivation and further engage patients in the interactive rehabilitation process. Data analytics and pattern recognition techniques are incorporated into such systems to assist the clinical evaluation and decision-support processes. (Biodex Medical Systems)

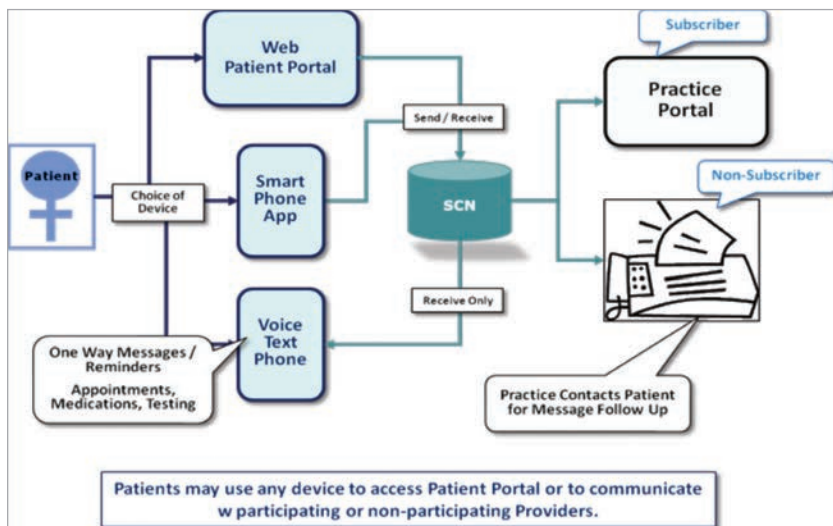


Figure 9

RFID sense-a-tags for the Internet of Things

Petar M. Djuric

petar.djuric@stonybrook.edu

We address the question of the feasibility of a design of backscattering devices that can communicate independently. We investigate the design of a novel component for the Internet of Things. We refer to it as a sense-a-tag tag. The tag may be passive or semipassive, and it has the following functionalities: (a) initiation and generation of query signals for the tags in its proximity, (b) passive detection and decoding of backscatter signals from RFID tags in its vicinity, (c) processing of information that it acquires during its operation, and (d) communication of the collected information by backscattering. The research also involves development of protocols for the new tags. (NSF)

Design of Audio Interfaces in Adverse Acoustic Environments

Miltuin Stanacevic

miltuin.stanacevic@stonybrook.edu

Achieving high speech intelligibility and recognition in noisy and reverberant environment presents a hard problem especially when the size of the microphone array is limited. A time delays between the observed signals are fraction of the sampling time, while the reverberations add attenuated and delayed source signal to the mixture. We devised a new set of blind source localization and separation algorithms in the unique framework that combines the spatial sampling, sub-band processing and independent component analysis to achieve improvement in the localization and separation performance. The advancement in the separation performance leads to better intelligibility of the speech signals and improved speech recognition in the audio interfaces operating in adverse acoustic environments. Our design of miniature size sensor arrays with sensory information processing implemented in mixed-signal VLSI enables a wide range of new applications of the miniature microphone arrays where the power consumption and size of the digital implementation where the limiting factors. (NSF)

Non-isotropic Networked Sensor Deployment for Smart Buildings

Jie Gao

jgao@cs.stonybrook.edu

The recent advancement of ubiquitous computing, sensing, distributed control and embedded systems has enabled a wide spectrum of smart building applications, such as surveillance for public safety, power monitoring and control for energy saving, patient monitoring for health care, and hazard detection for disaster response. The deployments of those sensing and control systems have been tested in a number of small-scale houses and labs. However, due to the use of simplified isotropic sensing model, the initial deployments usually do not yield satisfactory sensing quality in real world settings. As a result, adjustment of initial deployment or even multiple redeployments are typically required, which introduces high overhead and decreased quality of applications. The inefficiency and inaccuracy of existing sensor deployment call for new research solutions. The desired solution shall systematically guide the sensor deployment through the integration of sensing and communication considerations. To accomplish this task, this work aims at (i) addressing a fundamental challenge that is essentially tied with the nature of heterogeneity of sensors: non-isotropic sensing models, (ii) proposing a set of deploying solutions suitable for a wide range of sensor systems, and (iii) designing and implementing an open test-bed and simulation tools to evaluate real world deployments for the community.

This work investigates the non-isotropic networked sensor deployment for localization and tracking of human and activities in smart buildings, with special attention on a new set of problems arising from the deployment of four types of sensors: cameras, tripwire, motion sensors and microphones. The set of deployment problems would strive for full sensor coverage and wireless connectivity with a complex floor plan, and involve one or more of the following constraints: (i) Non-isotropic model using visibility (cameras and tripwires), (ii) Non-overlapping sensing range, (iii) Robust 2-coverage for accurate localization. The sensor deployment problem will heavily involve the geometric properties of the environment (building layout), as well as the geometric properties of the sensors (non-iso-

tropic sensing ranges). Ideas from computational geometry will be applied to tackle these problems. For some of the problems it is not hard to show to be NP-hard. Thus we will focus on approximation algorithms with worst-case guarantee for theoretical rigor as well as practical algorithms suitable for implementations. (NSF, DARPA)

Self-powered Wireless Hybrid Langasite Sensor for Pressure/Temperature Monitoring of Nuclear Reactors

Lei Zuo (SBU), Haifeng Zhang (UNT)
and Jie Lian (RPI)

lei.zuo@stonybrook.edu

The objective of this proposal is to develop a novel self-powered wireless hybrid sensor that can accurately monitor both pressure and temperature using a single device without requiring external electricity, even in the extreme harsh environments of severe nuclear accidents, such as up to 14000C temperature, 10,000 psi pressure, and excessive radiation. This objective is achieved through three innovations: the first innovation is to design a dual-mode langasite (La3Ga5SiO14) resonant sensor to detect extreme high temperature and high pressure simultaneously; the second is to create a multi-source energy harvester to harness intrinsic heat of the reactor and the kinetic energy of the reactor components (such as pump vibration or seismic motion) to provide the electric power needed for the sensors; and the third innovation is to design a self-protected sensor package upon integration of radiation shielding and mechanical support for mitigating severe environmental impacts.

The novel concept of self-powered wireless langasite-based P/T sensor that will operate under high temperature, high pressure, and excessive radiation provides a pathway to significantly improve monitoring technology for current nuclear reactors, and unquestionably support the program of Nuclear Reactor Technologies: Monitoring Technologies for Severe Accident Conditions. Upon success, the technology can also be used during normal operating conditions to provide enhanced monitoring of critical components in a standoff and energy-efficient manner. (DOE)




Figure 10

Smart Grid Regional Demonstration: Long Island Smart Energy Corridor

Eugene Feinberg, Samir Das, Rob Johnson, Ari Kaufman, Klaus Mueller, Erez Zadok
eugene.feinberg@stonybrook.edu

This is a joint project with Long Island Power Authority and Farmingdale College. The goal is to create a Smart Energy Corridor along Route 110 on Long Island. The project demonstrates the integration of a suite of Smart Grid technologies on the distribution and consumer systems, such as smart meters, distribution automation, distributed energy resources, and electric vehicle charging stations.

The project also includes testing cybersecurity systems, using Smart Grid technologies to enhance efficiency and reliability of the distribution network, identifying optimal combinations of features to encourage consumer participation, and educating the public about the tools and techniques available with the Smart Grid. *Figure 10*. (DOE)

Smart Composites for Energy Harvesting

T. A. Venkatesh
t.venkatesh@stonybrook.edu

Smart piezoelectric materials, by virtue of their coupled electromechanical characteristics, have been recognized for their potential utility in many applications as sensors and actuators, from medical ultrasound devices for prenatal care, micro/nano-positioners for atomic force microscopes and sonar hydrophones to non-destructive testers and inkjet print heads. Considerable research efforts in the past years have resulted in the development of several monolithic piezoelectric materials such as lead zirconate titanate (PZT) and barium titanate, with enhanced coupled properties. However, despite the enhancement in their piezoelectric properties, monolithic piezoelectric materials generally exhibit certain limitations. For example, they are mechanically brittle as most of the piezoelectric materials are ceramic-type materials and their functionality is generally unidirectional as the poling characteristics of the piezoelectric material allow them to sense or actuate in one direction (i.e., in the dominant poled direction) only. Because of these limitations, the range of applicability of monolithic piezoelectric materials is limited. A composites' approach to piezoelectric materials can potentially overcome the limitations of

monolithic piezoelectric materials. The overall objectives of our research efforts are: (i) To obtain a comprehensive understanding of the fundamental properties of smart piezoelectric composites; and (ii) To design novel smart materials based devices and structures for sensing and actuating functions as well as for energy harvesting applications. (CEWIT)

Thin Film Solar Cells with Tunable Transparency

T. A. Venkatesh (SBU) and
M. Cotlet (BNL)
t.venkatesh@stonybrook.edu

Solar technologies are currently based on conventional solar cells made out of inorganic materials like silicon which are generally opaque. Organic semiconducting materials, including conjugated polymers, have been investigated as potential replacement materials for solar cells. Conjugated polymer-based organic photovoltaic solar cells have several attractive features such as being made of inexpensive materials, by cost-effective processing methods like printing techniques, dip or spin casting, requiring low amounts of active material (around 100 nm thick layer), and being light weight and mechanically flexible. However, even at thicknesses of around 100 nm, these continuous thin films also tend to be quite opaque, restricting their application for integrated photovoltaics such as power-generating windows. In order to obtain transparent solar cells, recent efforts have focused on reducing the thickness of the active layer to less than 100 nm, which improves the transparency, but drastically reduces the conversion efficiency of the solar cells. Hence, the overall objective of our research effort is to develop conjugated polymer-based thin films with a novel micro-porous structure as active materials for solar cells with tunable transparency and good photovoltaic properties. (DOE)

Enhanced Power System Operation and Control

Eugene Feinberg
eugene.feinberg@stonybrook.edu

This project investigates and implements numerical algorithms and high performance computing software for solving power-flow, state-estimation, and system-stability control problems for electric transmission grids. The main objective is to develop solution methods that are 10 to 100 times faster than the solution times achieved in current control system implementations. (NYSERDA, NYPA)

Smart Grid Android Manager

Radu Sion
sion@cs.stonybrook.edu

The Mobile Smart Energy Manager is an Android/iPhone based Energy Manager that runs on smart phones and tablets. It will connect to local autonomous micro-grid information gateways and components (such as solar panel relays), the macro-grid operators and utility companies' portals, smart home components, and the internet, to provide real-time energy-related information, debit and consumption data, billing and the ability to manage smart-home devices in real-time on-demand.

The Mobile Smart Energy Manager will also allow a unified user control platform for the integration of various external smart grid data processing and visualization plugins. Examples of such plugins include: (1) big data analytics visualization of micro-grid and macro-grid energy data, (2) connectivity conduit to external data sources and (3) visual devices such as the reality deck display and the SCADA smart grid control center, (4) networking plugins to interface with any additional custom wireless protocols designed as part of the SGRID3 project. (DOE)

An Interactive User Interface for the Smart Grid

Arie Kaufman and Klaus Mueller
ari@cs.stonybrook.edu

In the traditional system, customers just purchase the energy from suppliers and consume it. However, smart grid is two-way communication channel between suppliers and consumers. The roles of consumers are to reduce their total consumption and shift their usage to off-peak time. However, it is difficult to encourage consumers to change their behavior with simple visualization. In this project, we have developed an interactive system to help customers gain better understanding of their energy consumption. In our system, since customers hardly understand their energy consumption of their own electric devices, customers could configure their own virtual house with electric devices to estimate their current energy consumption. Customers could choose what kind of devices they have by dragging and dropping an electric device into their virtual house. Customers can easily select a specific model of each device. Our system also provides a tool to analyze their energy consumption pattern in order to control their energy usage efficiently. Given their total energy consumption from their smart meters, it shows their daily, weekly, and monthly energy usage patterns. In addition, it enables customers to predict their energy consumption once they replace their current electric devices with new ones. (DOE)



Non-Visual Skimming: Improving the Usability of Web Access for Blind People

*Yevgen Borodin and
I.V. Ramakrishnan
Yevgen.Borodin@stonybrook.edu*

In our information-driven web-based society, we are all gradually falling “victims” to information overload. However, while sighted people are finding ways to sift through information faster, visually impaired computer users are experiencing an even greater information overload. These people access computers and the Web using screen readers, the software that narrates the content on a computer screen using computer-generated speech, and allows users to navigate with keyboard shortcuts or gestures.

While sighted people can learn how to quickly skim web content to get the gist of the information and find the information they need, people who are blind have to process volumes of content narrated through a serial audio interface, which does not allow them to find out what content is important before they listen to it. So, they either listen to all content or listen to the first part of each sentence or paragraph before they skip to the next one. Dr. Borodin and Dr. Ramakrishnan are developing novel interfaces and algorithmic techniques for non-visual skimming that will empower people with vision impairments to access digitized information significantly faster than they currently can with the state-of-the-art screen readers. In this way, skimming will help reduce the cognitive load associated with non-visual browsing.

In the process of skimming, sighted people quickly look through content while picking out words and phrases that are emphasized visually and/or carry the most meaning. Dr. Borodin and Dr. Ramakrishnan are exploring ways to emulate this process and enable a computer-assisted skimming experience for screen-reader users by designing interfaces and algorithms

for non-visual skimming. The research and development of non-visual skimming interfaces in this project can lead to drastic improvement of Web accessibility for blind people. Non-visual skimming will not only enhance webpage navigation, but it may also address the usability of non-visual browsing interfaces by reducing information overload. (NSF)

Divisible Load Scheduling

*Thomas Robertazzi
thomas.robertazzi@stonybrook.edu*

Ongoing work seeks to look at novel ways in which very large partitionable data loads can be processed in a minimal amount of time with the biggest parallel processing advantage. Divisible or partitionable loads are loads that can be chopped arbitrarily into smaller fragments for processing on a networked collection of processors. One would like to make the best temporal choice and assignment of fragments to processors and links so as to minimize solution time and maximize parallel processing gain (known as a speedup). A number of uses are being investigated. One is processing huge volumes of data looking for characteristic patterns (i.e. signatures) that signify the presence of an interesting event or object. Applications include network security, sensors, bio-medical problems and image processing. Work to date has examined establishing the performance of systems doing such processing. A second use is in systems where there is a monetary cost for doing processing and transmission. The goal is to tradeoff solution time against the cost of problem solution. This is relevant in an era of cloud computing. Finally, using divide and conquer strategies to speed problem solving has been examined. Such techniques can yield impressive savings in time. Since the original work on analytical solutions for partitionable loads by Cheng and Robertazzi, and also by researchers at Bell Labs, almost 150 journal papers have been written on divisible load scheduling. It remains an active and relevant field, especially in an age of big data. (CEWIT)

Approximation Algorithms for Geometric Optimization

*Esther Arkin and Joe Mitchell
Joseph.Mitchell@stonybrook.edu*

We apply the methodologies of computational geometry to design, analyze, implement, and test algorithms for problems that arise in several application areas, including geometric network optimization, air traffic management, sensor networks, robotics, geometric modeling, and manufacturing. The main goal is the development of fundamental advances in approximation algorithms for geometric problems. Additionally, the project will strive to foster and deepen collaborations with researchers and domain experts in application areas and industry, in order to formulate their algorithmic needs precisely and to make available algorithmic tools, insights from theoretical results, and software from experimental investigations. The specific repertoire of problems includes geometric network optimization (optimal routing and network design in geometric contexts, including TSP variants, vehicle routing, constrained spanning trees, minimum-weight subdivisions, optimal route planning with various constraints, and survivable network design); air traffic management (optimal use of airspace in the face of dynamic and uncertain constraints induced by weather and traffic congestion, sectorization (load balancing), and optimization of flow management structures for the National Airspace System); and sensor networks and coverage (sensor deployment, localization, data field monitoring, and coverage for stationary or mobile (robotic) sensors). (NSF)

Geometric Networks

*Esther Arkin and Joe Mitchell
Joseph.Mitchell@stonybrook.edu*

We study geometric networks, which represent interconnections between entities that arise in physical domains or geometric spaces. Networks are all around us and are an important part of the technology in our daily lives. Examples of geometric

networks include wired/wireless communication networks, transportation systems, power grids, sensor networks, and geometric graphs that arise in information visualization. Geometric networks often have special structure that allows their analysis and optimization to be done more efficiently than is possible in general (non-geometric) networks. We study the design and analysis of energy-efficient wireless communication networks. In particular, we investigate wireless communication networks involving directional antennas and/or network improvement methods. We also study several related optimization problems in geometric networks that arise in other applications, including sensor networks, transportation science, air traffic management, vehicle routing in robotics, covering tours, and exploration/mapping. The primary objective of the project is to develop new algorithmic solutions to a cohesive collection of geometric optimization problems that fall in the common category of network problems. The goal is to develop theoretically sound, provable methods, but with a strong awareness towards practicality and implementation. (US-Israel Binational Science Foundation)

Algorithms in Support of Flight Trajectory Analysis

Joe Mitchell

Joseph.Mitchell@stonybrook.edu

This project applies geometric algorithms to the analysis of large databases of trajectory data for moving agents. Given the exploding amount of time trajectory data collected, there is a need to be able to process, analyze, mine, and understand trajectories that arise from vehicles (cars, buses, trains, aircraft), pedestrians, animals, autonomous vehicles/sensors, etc. Using methods of computational geometry and geometric data structures, we design methods and tools to process trajectory data in the space-time domain, allowing a high level of understanding of structures in patterns and rapid query access to large databases. (Sandia National Labs)

Using Evolutionary Computations to Understand the Design and Evolution of Gene and Cell Regulatory Networks

Alexander V. Spirov

Alexander.Spirov@cewit.stonybrook.edu

A central challenge to understanding gene expression is to recreate the regulatory networks controlling expression; or at the least to generate models of these networks which capture essential characteristics of their connectivity and control, and which can be quantitatively analyzed. By developing such a quantitative theory of gene control, we can hope to develop far more powerful experimental tests and applications. A great deal of diverse work has gone into developing and testing models of gene regulatory networks (GRNs) in recent decades. As GRN models become more developed, allowing us greater understanding of the particular 'wiring' involved in particular cases, it becomes natural to ask how these GRNs evolved, to uncover general principles of evolutionary processes. Such questions can be approached computationally by evolution in silico. Complementary to these biologically-focused approaches, a now well-established field of computer science is Evolutionary Computations (EC), in which highly efficient optimization techniques are inspired from evolutionary principles. In our project we focused on the considerations that must be taken with respect to level of detail to model (from solely gene-gene interactions to the DNA sequence level); and efficiency of computation. With respect to the latter, we argue that developments in computer EC offer the means to perform more complete simulation searches, and will lead to more comprehensive model predictions. While the some of the projects on GRN evolution in silico begin to account for multiple binding of transcription factors, they do not address the architecture of genes'

cis-regulatory regions. Beginning to model and understand the role of this architecture on gene regulation requires what we call a 'mid-grained' level of approach. Mid-grained modeling involves treating functional regions of the cis-regulatory modules (CRMs) as discrete units subject to evolution. The choice of fitness function is crucial for GRN evolutionary design. Particularly, a solution can be decomposed into a number of building blocks (BBs, e.g. biological cis-regulatory modules, CRMs), which can be searched for independently and afterwards be combined to obtain good or optimal solutions. To implement BBs in GRN evolution, we have developed a crossover recombination operator which maintains meaningful BB sequences, such as CRMs. Systematic comparisons of the cost/benefit between coarse-grained and finer-grained approaches are needed for more test cases, to better understand the appropriate level for addressing particular questions. Our conclusion at this stage is that the mid-grained level of GRN modeling (CRM level) is the best tradeoff between highly expensive calculations (which impact the extent of computational evolution that can be performed) and biologically reasonable simplification of the gene regulatory organization. The biological and computer science approaches to evolution have operated somewhat independently, but have much to offer each other with increased crossover of ideas. We feel, especially, that simulations of GRN evolution, which are computationally intensive, could benefit from increased use of EC optimization techniques and the analytical tools that come with them. The progress in modeling biological evolution will be greatly enhanced by incorporation of recent developments in EC optimization techniques, many of which have already been developed on 'test case' biological type problems. At the same time, the computer science stands to benefit greatly from increased appreciation and use of the complexity of biological evolutionary dynamics. (CEWIT)

Adaptive Runtime Verification and Recovery for Mission-Critical Software

Radu Grosu, Scott A. Smolka, Scott D. Smolka, Erez Zadok
sas@cs.stonybrook.edu

Runtime verification (RV) refers to the use of lightweight yet powerful formal techniques to monitor, analyze, and guide the execution of programs at run-time. RV is becoming increasingly important for at least two reasons. First, software systems are becoming more complex and more adaptive, making it difficult to statically understand all of their possible behaviors; this is especially true of mission-critical software on autonomous unmanned vehicles, where completion of mission goals depends upon adaptive responses to changing conditions. RV thus plays a valuable role in helping users monitor and understand system behavior during testing, debugging, and deployment. Second, to increase the reliability of complex adaptive software, RV must monitor and analyze the behavior of the software, its environment, and their interaction, in order to trigger adaptive responses when necessary.

To fill these needs, runtime verification itself must become more flexible and adaptive, and it must be equipped with a recovery framework that will help ensure

mission completion in the face of runtime violations. We propose to develop Adaptive Runtime Verification and Recovery (Arrive), a novel extension of runtime verification in which the runtime verification itself is adaptive. Arrive will dynamically allocate more runtime-verification resources to high-criticality monitored objects, thereby increasing the probability of detecting property violations within a given overhead budget. Moreover, when a violation is imminent, Arrive will take adaptive and possibly preemptive action in response, thereby ensuring recovery.

We are investigating three related aspects of Arrive: overhead control, incomplete monitoring, and predictive analysis. We are developing a Simplex architecture for cyber-physical systems by extending Simplex to monitor the software state as well as the physical-plant state. We will evaluate the performance and utility of the Arrive framework through significant case studies, including the runtime monitoring of the command-and-control and energy-management infrastructure of a fleet of UAVs.

It is anticipated that this project will be funded by the Air Force Office of Scientific Research (AFOSR) starting in the first quarter of FY 2014.

Computational Modeling and Analysis for Complex Systems

Radu Grosu, James Glimm,
Scott A. Smolka
sas@cs.stonybrook.edu

The CMACS project is a 5-year \$10M multi-institution NSF Expedition in Computing project focused on far-reaching and transformative research into techniques based on Model Checking and Abstract Interpretation (MCAI) for analyzing the behavior of complex embedded and dynamical systems. Model Checking and Abstract Interpretation have a 30-year record of success at verifying properties of the behavior of discrete systems automatically. The techniques have been fruitfully used, both independently and in combination, to establish properties of systems containing thousands of variables and inputs and several 100,000 of lines of code (for example, the Airbus 380 flight control software), and to detect subtle bugs in a variety of hardware and software applications, ranging from microprocessor designs and communication protocols to railway-switching systems and satellite-control software.

The purpose of the CMACS project is to extend the MCAI paradigm to reasoning about the behavior of models of physical systems that include continuous and stochastic behavior, such as those found in biological and embedded-control areas. Specific research is being undertaken in model discovery/system identification for stochastic and nonlinear hybrid systems; methods for generating sound model abstractions to simplify the reasoning process; and next-generation algorithms for analyzing the behavior of these models. Challenge problems in the areas of pancreatic-cancer modeling, atrial-fibrillation detection, distributed automotive control, and aerospace control software are being used as technology drivers and testbeds for the results obtained in the course of the project. (NSF)




Algorithmic Foundations for Hybrid Mobile Sensor Networks

Jie Gao
jgao@cs.stonybrook.edu

The objective of this project is to develop efficient distributed algorithms for integrating static, pervasive sensor nodes and mobile, autonomous agents that altogether monitor and act on the environment. The project leverages on the recent advance and maturity of static wireless sensor networks for environment monitoring and projects the vision of seamlessly integrating physical users and/or autonomous robots into a hybrid intelligent framework. The focus is on the joint information processing, optimization and coordination of both static and mobile nodes to comprehend and act on the environment. The project covers the following application categories: mobile robots as aids to sensor network operation, tracking, searching and planning mobile identities, online resource management and allocation, maintaining group communication and coordination of mobile agents. A common theme in these problems is the tight coupling of and the frequent information flow between static monitoring nodes and mobile actionable agents.

To enable mobile agent coordination and ensure smooth and frequent information flow between sensor nodes and the agents, we ask how sensor data should be stored and processed in the network, and what type of distributed structures should be maintained, to best serve mobile interactive users. There are two central intellectual questions that transcend the algorithm and system design. (NSF, DARPA)



- How to make use of the data stored in the network to best serve user requests? The biggest challenge is to solve the information brokerage problem, in which sensors detecting interesting local events and the mobile users seeking such information are not aware of one another. Non-trivial brokerage methods need to be designed to 'match' the data suppliers and data consumers, without communication expensive routines such as flooding.

- How to best make use of the continuity and coherence in mobility, either in the presence of mobility enabled data (e.g., continuous trajectories for mobile targets), as well as the locations of mobile users? We are looking for solutions that smoothly adapt to the system configuration with low worst-case or amortized update cost. We want to avoid sudden changes or any level of reconstruction. Smoothly deforming communication structures will be helpful for continuous, steady data flow from the network to mobile users and between multiple mobile users.

The embedded sensors integrated smoothly with other mobile embedded devices would provide real-time data collection, knowledge extraction, environment understanding and eventually evolve into an intelligent component of a smart environment. (CEWIT)

From Clarity to Efficiency for Distributed Algorithms

Y. Annie Liu and Scott D. Stoller
liu@cs.stonybrook.edu

Distributed algorithms are at the core of distributed systems, which are increasingly indispensable in our daily lives. Yet, developing practical implementations of distributed algorithms with correctness and efficiency assurance remains a challenging, recurring task.

This project develops formal semantics and powerful optimizations for a very high level language for distributed algorithms. The language, DistAlgo, allows the algorithms to be expressed easily and clearly, making it easier to understand them precisely and to generate executable implementations. The semantics are for exact understanding of the language. The optimizations are for generating efficient implementations, without manually coding the algorithms and applying ad hoc optimizations.

We evaluate the language and the translation and optimization methods by applying them to important and difficult distributed algorithms. Based on our experience and success teaching concurrent and distributed algorithms in the past four years, we also propose to systematically apply our methods to major distributed algorithms in textbooks. We will make the results publicly available on the Web. (NSF)

Demand-Driven Incremental Object Queries

Y. Annie Liu and Scott D. Stollerliu@cs.stonybrook.edu

High-level queries such as relational database queries have provided significant productivity improvement in the development of database applications. Relational queries are supported by powerful query optimization engines and are significantly easier to write, understand, and maintain than efficient low-level programs. Similar queries, in the form of comprehensions, have also been used increasingly in programming languages, from SETL to Python to Haskell to C# with LINQ and many more.

Object queries can provide even more productivity improvements, because they allow easier and more natural modeling of real-world data and easier and clearer expression of complex queries. Object-oriented modeling and programming have indeed been widely adopted for improved productivity in application development. However, efficient implementation of object queries remains challenging, especially to support queries on demand, on dynamically determined parameter values, and to make the queries incremental, under updates to the values that the queries depend on, while ensuring the desired correctness.

We study a systematic method, based on relational invariants, for efficient implementation of demand-driven incremental object queries. The method translates demand-driven object queries into relational queries, and incrementalizes and optimizes the resulting relational queries by exploiting constraints from objects and demands. The method applies to queries

over arbitrarily nested objects and sets of objects, for any dynamically demanded parameter values, under all possible updates to the values that the query depends on.

Transformation into relational queries allows clean and uniform treatment of objects, sets, and demands, all of which are captured precisely by constraints on relations, resulting in systematic consideration of the trade-offs in incrementalization and optimization. This then allows our method to provide precise complexity guarantees for the generated implementations. Our transformations that exploit demands systematically reduces not only time complexity drastically but also needed space significantly. Our prototype implementations have been applied successfully to complex query problems in a variety of application domains, including social network analysis, role-based access control, query-based debugging, and distributed algorithms. (NSF and ONR)

Performance Analysis and Optimization for Logic Rule Engines

*Michael Kifer, Y. Annie Liu, and David Warren
liu@cs.stonybrook.edu*

As the range of applications of rule-based systems has expanded and W3C is finalizing standards for rules on the Semantic Web, understanding of the execution performance of such systems has become more critical. Rule-based systems are now widely being applied to the Semantic



Web, system security and access control, program analysis and verification, cognitive radio, disruption-tolerant networking, sensor networks, and a number of other domains, each with its own demand for efficiency and scalability. Various implementations of rule systems have been developed and shown to have widely varying performance characteristics. While some systems generally have better performance than others, no single system dominates for all types of rule programs. This project aims to take one successful rule system, the XSB Logic Programming System, and use techniques from other implementations that for some programs have better performance to guide the development of a program optimization strategy for XSB that will make it perform uniformly well.

The approach will be through global program transformations, drawing on techniques from relational database query optimization, compiler optimization and recursive query optimization. Our approach is to 1) explore and collect techniques for generating (possibly more efficient) logically equivalent programs; 2) develop cost models and analysis techniques to generate parameterized cost formulas that approximate the running time (and space) of given rule programs; 3) explore and develop techniques to estimate the parameters required to evaluate the cost formulas; 4) develop heuristic search techniques to approximately search an exponentially large space of rule programs logically equivalent to a given one, to find the one with the least predicted running time (and/or space); and 5) implement the strategy to perform as a program optimizer for the open-source, widely used, XSB Logic Programming System. (NSF)

Secure Provenance in High-End Computing Systems

*Radu Sion and Erez Zadok
sion@cs.stonybrook.edu
zadok@cs.stonybrook.edu*

Data provenance documents the inputs, entities, systems, and processes that influence data of interest—in effect providing a historical record of the data and its origins. The generated evidence supports essential forensic activities such as data-dependency analysis, error/compromise detection and recovery, and auditing and compliance analysis.

This collaborative project is focused on theory and systems supporting practical end-to-end provenance in high-end computing systems. Here, systems are investigated where provenance authorities accept host-level provenance data from validated provenance monitors, to assemble a trustworthy provenance record. Provenance monitors externally observe systems or applications and securely record the evolution of data they manipulate. The provenance record is shared across the distributed environment.

In support of this vision, tools and systems are explored that identify policy (what provenance data to record), trusted authorities (which entities may assert provenance information), and infrastructure (where to record provenance data). Moreover, the provenance has the potential to hurt system performance: collecting too much provenance information or doing so in an inefficient or invasive way can introduce unacceptable overheads. In response, the project is further focused on ways to understand and reduce the costs of provenance collection. (NSF)

Eco Hadoop: Cost and Energy-Aware Cloud and HPC Computing

*Radu Sion
sion@cs.stonybrook.edu*

The project examines novel services built on top of public cloud infrastructure to enable cost-effective high-performance computing. Eco Hadoop explores the on-demand, elastic, and configurable features of cloud computing to complement the traditional supercomputer/cluster platforms. More specifically, this research aims to assess the efficacy of building dynamic cloud-based clusters leveraging the configurability and tiered pricing model of cloud instances. The scientific value of this proposal lies in the novel use of untapped features of cloud computing for HPC and the strategic adoption of small, cloud-based clusters for the purpose of developing/tuning applications for large supercomputers.

Through this research, we expect to answer key research questions regarding: (1) automatic workload-specific cloud cluster configuration, (2) cost-aware and contention-aware data and task co-scheduling, and (3) adaptive, integrated cloud instance provisioning and job scheduling, plus workload aggregation for cloud instance rental cost reduction. If successful, this research will result in tools that adaptively aggregate, configure, and re-configure cloud resources for different HPC needs, with the purpose of offering low-cost R&D environments for scalable parallel applications. (NSF)

An Efficient, Versatile, Scalable, and Portable Storage System for Scientific Data Containers

*Erez Zadok
ezk@fsl.cs.sunysb.edu*

Scientific big data sets are becoming too large and complex to fit in RAM, forcing scientific applications to perform a lot of slow disk and network I/O. This growth also makes scientific data more vulnerable to corruptions due to crashes and human errors. This project will use recent results from algorithms, database, and storage research to improve the performance and reliability of standard scientific data formats. This will make scientific research cheaper, faster, more reliable, and more reproducible.

The Hierarchical Data Format (HDF5) standard is a container format for scientific data. It allows scientists to define and store complex data structures inside HDF5 files. Unfortunately, the current standard forces users to store all data objects and their meta-data properties inside one large physical file; this mix hinders meta-data-specific optimizations. The current storage also uses data-structures that scale poorly for large data. Lastly, the current model lacks snapshot support, important for recovery from errors.

A new HDF5 release allows users to create more versatile storage plugins to control storage policies on each object and attribute. This project is developing support for snapshots in HDF5, designing new data structures and algorithms to scale HDF5 data access on modern storage devices to Bigdata. The project is designing several new HDF5 drivers: mapping objects to a Linux file system; storing objects in a database; and accessing data objects on remote Web servers. These improvements are evaluated using large-scale visualization applications with Bigdata, stemming from real-world scientific computations. (CEWIT)



Workload-Aware Storage Architectures for Optimal Performance and Energy Efficiency

Erez Zadok
ezk@fsl.cs.sunysb.edu

The most significant performance and energy bottlenecks in a computer are often caused by the storage system, because the gap between storage device and CPU speeds is greater than in any other part of the machine. Big data and new storage media only make things worse, because today's systems are still optimized for legacy workloads and hard disks. The team at Stony Brook University, Harvard University, and Harvey Mudd College has shown that large systems are poorly optimized, resulting in waste that increases computing costs, slows scientific progress, and jeopardizes the nation's energy independence.

First, the team is examining modern workloads running on a variety of platforms, including individual computers, large

compute farms, and a next-generation infrastructure, such as Stony Brook's Reality Deck, a massive gigapixel visualization facility. These workloads produce combined performance and energy traces that are being released to the community.

Second, the team is applying techniques such as statistical feature extraction, Hidden Markov Modeling, data-mining, and conditional likelihood maximization to analyze these data sets and traces. The Reality Deck is used to visualize the resulting multi-dimensional performance/energy data sets. The team's analyses reveal fundamental phenomena and principles that inform future designs.

Third, the findings from the first two efforts are being combined to develop new storage architectures that best balance performance and energy under different workloads when used with modern devices, such as solid-state drives (SSDs), phase-change memories, etc. The designs leverage the team's work on storage-optimized algorithms, multi-tier storage, and new optimized data structures. (CEWIT)

Pochoir: A Stencil Computation Compiler for Modern Multicore Machines

Rezaul Chowdhury
rezaul@cs.stonybrook.edu

A stencil computation repeatedly updates each point of a d-dimensional grid as a function of itself and its near neighbors. These computations arise in diverse application areas including physics, biology, chemistry, finance, mechanical and electrical engineering, energy, climate, and recreational mathematics. Such computations are conceptually simple to implement using nested loops, but looping implementations suffer from poor cache performance. Real stencil applications often exhibit complex irregularities and dependencies, which makes it difficult for programmers to produce efficient multicore code for them or to migrate them to other modern hardware platforms. Even simple stencils are hard to code for performance. This project aims to solve the difficult problem of generating high-efficiency cache-oblivious code for stencil computations that make good use of the memory hierarchy and processor pipelines, starting with simple-to-write linguistic specifications. We are developing a language embedded in C++ that can express stencil computations concisely and can be compiled automatically into highly efficient algorithmic code for multicore processors and other platforms. This research will enable ordinary programmers of stencil computations to enjoy the benefits of multicore technology without requiring them to write code any more complex than naive nested loops. (NSF)

Energy-Efficient Superconductor SFQ Processor Design

Mikhail Dorojevets
Mikhail.Dorojevets@stonybrook.edu

New SFQ logics with extremely low power consumption at frequencies of 20-50 GHz make superconductor technology one of the candidates for use in future energy-efficient systems for critical national applications. In the first phase of the project, the PI's SBU team has developed viable design techniques, including asynchronous wave pipelining, and set of design tools and cell libraries for VLSI-scale superconductor design. Several RSFQ chips with the complexity of up to 10K Josephson Junctions (JJs) have been successfully designed using those techniques and demonstrated operation at 20 GHz. The objective of the current phase of the project (2012-2015) is to develop and evaluate an ultra-low-power design of 10-20 GHz 32-bit superconductor processors using a novel Reciprocal Quantum Logic (RQL) with zero static power consumption. The cell-level design and energy efficiency evaluation of a complete set of 32-bit processing and storage components is to be done using SBU RQL VHDL-based design & evaluation tools. The goal of the work is to get a quantitative assessment of how successfully the superconductor technology in its latest form will be able to compete in terms of energy efficiency with future 10 nm CMOS-based designs. (ARO)

Building Metaknowledge Representations in Circuit Design: Symbolic Data Mining for Systematic Modeling of Analog Circuits

Alex Doboli
alex.doboli@stonybrook.edu

Analog circuits are important building blocks in many modern electronic systems used in manufacturing automation, telecommunication, healthcare, defense, infrastructure management, and many more. Analog circuit design tasks are often conducted manually as circuit design has been difficult to automate. Efficient EDA tools have emerged over the years for circuit simulation, transistor sizing and layout generation but other design activities like circuit topology generation, refinement and selection, circuit modeling and verification are still performed mainly by skilled experts. Our recent research work has been addressing the topic of automatically extracting and characterizing the design knowledge embedded in analog circuits, and then using the mined knowledge to produce new circuit topologies through steps like aligning and combining design features, abstracting and generalizing concepts, and inferring new features by similarities and induction. As research in cognitive psychology shows, these are the main steps in human reasoning. We think that an automated, knowledge-based reasoning flow can produce circuit designs that are more similar to those created by human experts, offer superior performance, and are easier to verify. It can also enable new kinds of EDA tools to potentially express (all) knowledge in circuit design, produce innovative designs for new applications, and exploit unstudied opportunities (NSF)

Leveraging Three-Dimensional Integration Technology for Highly Heterogeneous Systems-on-Chip

Emre Salman
emre.salman@stonybrook.edu

Three-dimensional (3D) integration is an emerging integrated circuit (IC) technology that maintains the benefits of scaling by vertically stacking multiple wafers rather than decreasing the size of the transistors in two dimensions. 3D technology not only increases the integration density, but also reduces the length and number of global interconnects. This reduction enhances chip performance (due to reduced interconnect delay) while lowering the power consumption (due to reduced switching capacitance). These advantages have attracted significant attention in the past decade to develop high performance computing systems such as many-core processors with embedded memory.

This research focuses on 3D heterogeneous integration to expand the application domain of 3D integrated circuits from high performance computing to relatively lower power systems-on-chip (SoCs) consisting of, for example, sensors, analog interface circuit, memory, digital processing blocks, and RF wireless transmission circuitry. (NSF, CAREER)

The primary objective is to develop a reliable framework to facilitate highly heterogeneous and tightly coupled 3D SoCs for substantially complex applications, where plug-and-play based approaches are not sufficient. The proposed circuit- and physical-level design and analysis methodologies serve as a novel framework to produce robust 3D SoCs consisting of diverse electrical planes such as analog/RF, digital, and memory. At a time when the fundamental limits of traditional scaling are approaching, the synergy between 3D integration and heterogeneous SoCs brings new opportunities to a variety of applications such as mobile computing, life sciences, and environmental control.

Figure 11. (NSF)

General-Purpose Computation on Graphics Processing Units

Arie Kaufman and Klaus Mueller

A commodity graphics processing unit (GPU) can perform general-purpose computations and not only the specific built-in graphics operations. GPU is a multi-core coprocessor that supports high data parallelism. Its performance has been growing at a rate of squared Moore's law, and its peak performance exceeds that of the CPU by an order of magnitude. We further advocate the use of a GPU cluster for

large scale problems and for high performance computing. We have built several large GPU clusters and implemented several applications on them. The most recent GPU cluster for the Reality Deck is the largest one with an aggregate 2.3 teraflops CPU performance, 220 teraflops GPU performance and 1.2 TB of memory. However, programming a GPU cluster is difficult. Therefore, we have developed the Zippy framework to simplify the GPU cluster programming. Zippy is based on global arrays and stream programming and hides the low-level details. The GPU clusters have the potential to become the peta-scale supercomputers of the near future. We have done much groundbreaking computational science work on GPU and GPU clusters, such as computational fluid dynamics using the Lattice Boltzmann Model (LBM), medical image reconstruction (RapidCT), large-scale complex volume visualization, and hashing and compression on the GPU. A specific example of using the GPU cluster includes real-time plume dispersion simulation in complex urban environments, such as midtown Manhattan, where nearly 1,000 skyscrapers geometries serve as boundary conditions which substantially influence the flow. Example applications using Zippy include LBM, volume visualization, and isosurface extraction. (NSF)





Fig.11 3D heterogeneous IC with four stacked dies



www.cewit.org



Center of Excellence
WIRELESS AND INFORMATION TECHNOLOGY
AT STONY BROOK UNIVERSITY

Center of Excellence in Wireless
and Information Technology (CEWIT)
Stony Brook University Research
and Development Park
1500 Stony Brook Road
Stony Brook, NY 11794-6040
Phone: (631) 216-7000
info@cewit.org