Klaus Mueller: Research Statement and Achievements

My research spans three principal domains, **Medical Imaging, Computer Graphics,** and **Visualization**, which is demonstrated by a wealth of publications, scholarly activities, and community service. The following document is just a narrated summary – more detail can be obtained from my full CV.

RESEARCH DOMAIN #1: MEDICAL IMAGING

In Medical Imaging I am best known for my pioneering work on exploiting commodity graphics hardware (GPUs), as commonly used in PC-based computer games, to accelerate computed tomography reconstruction (CT) algorithms, in particular the especially compute-intensive class of iterative CT reconstruction algorithms. My 2005 paper in *IEEE Transactions on Nuclear Science*, co-authored with my student Fang Xu, has become the classical reference on the theoretical underpinnings of GPU-based parallel computation for diagnostic imaging. This paper alone has 140 citations according to Google Scholar and is the 2nd most cited paper in the journal since 2005. Another, more recent paper with Fang Xu on a more focused topic, published in the journal *Physics in Medicine and Biology* in 2007, has 100 citations and is the 8th most cited paper in this journal since its publication. Along with my student Erik Papenhausen, I currently lead the worldwide ranking of top performing CT reconstruction implementations (http://www.rabbitct.com) – an achievement that won us the *2011 High Performance Reconstruction Workshop* Audience Best Paper Ward.

In this domain and in iterative CT in particular, it is testimonial for my scholarly esteem that even my doctoral dissertation, available on the web, has garnered 55 citations and apparently has been studied by many more. I have been invited to give keynotes at international conferences, such as the *IEEE International Symposium on Biomedical Imaging (ISBI)* in 2009 and the *SPIE International Electronic Imaging Conference* in 2007, and have given invited talks at various universities around the world, such as Yale University, Hong Kong University of Science and Technology, Vienna Technical University, Peking University, Zhejiang University and Ochanomizu University, Japan. In addition, I have lectured widely on the topic of GPU-accelerated CT, leading or participating in 5 full- and half-day tutorials at major international conferences, such as *IEEE Medical Imaging* and *SPIE Medical Imaging*. I have further served as co-chair at 4 high-performance computing workshops held in conjunction with major international medical imaging conferences. It is possibly due to these various educational efforts and also my continued publications on this topic that GPU-acceleration has now become a major computational platform in the medical imaging industry.

More recently, with my student Wei Xu, I have achieved new groundbreaking results by applying principles from computer vision and machine learning to aid in ill-conditioned CT reconstruction scenarios, for example when the projections are sparse in number or of poor quality. One of these works, reporting on an efficient GPU implementation of the framework, earned a Best Paper award at the 2009 *International High Performance Workshop on Image Reconstruction*, held in Beijing, China. The developed algorithmic framework holds great promise for reconstructing images with just a few projections and therefore at lower radiation doses to the patient, making medical imaging safer for large-scale clinical use.

RESEARCH DOMAIN #2: COMPUTER GRAPHICS

Along with my student Tomihisa Welsh and colleague Dr. Michael Ashikhmin, I have founded the flourishing research area of image and video colorization. My 2002 landmark paper, published in *ACM Transactions on Graphics* and presented at the flagship computer graphics conference *SIGGRAPH*, has thus far earned 296 references according to Google Scholar. Colorization has remained a major topic of my research but more as an application in the domain of visualization.

I have also made numerous contributions in the field of computational fluid dynamics where I, in collaboration with Prof. Arie Kaufman, devised efficient lattice structures that lend themselves well to GPU acceleration. This allows the simulation of amorphous phenomena, such as gas, fluid, and fire, at interactive frame rates to enable on-the-fly engineering design. A respectable set of publications have resulted from this effort, with two papers in *IEEE Visualization* (2002) and *IEEE Transactions on Visualization and Computer Graphics* (2004) cited 56 times each.

RESEARCH DOMAIN #3: VISUALIZATION

I am a leading expert on volume rendering, the prime technique to visualize medical data in 3D. In 2000, I authored, with 4 colleagues, the standard reference work comparing the four main algorithmic paradigms for volume rendering, both from a theoretical and from a practical standpoint. This paper, presented at the *IEEE/ACM Volume Visualization Symposium*, has 169 citations according to Google Scholar.

I have a prolific publication record in the field of visualization. Just in the flagship conference *IEEE Visualization* alone I have authored or co-authored 15 papers, as well as 2 poster extended abstracts. Of special note here is my groundbreaking work on an efficient and high-quality volume visualization algorithm called splatting, which has resulted in numerous papers, three of these are each cited 80 times or more. More recently, my research has expanded to Visual Analytics, which is the science of visualization-enabled reasoning with data – typically high dimensional. Here I have produced, with my students, a number of landmark results, such as my 2008 *Computer Graphics Forum* journal paper on *Illustrative Parallel Coordinates* (21 citations) that exploits rules gleaned from illustrative design for more intuitive high-dimensional data exploration, as well as my 2008 and 2010 *Visual Analytics Science and Technology (VAST)* conference papers (12 citations) on *Model-Driven Visual Analytics* that, via highly interactive visual interfaces, allows domain experts to interactively steer machine learning methods to build better models from data more efficiently.

Further evidence for my strong scientific recognition and stature in the field of visualization is the fact that journal and book editors have routinely relied on my scientific expertise and vision to provide up-to-date surveys or viewpoints on new developments. Thus far, I have co-authored (all with Prof. Arie Kaufman) three invited survey articles in major academic handbooks, such as the *Visualization Handbook* published by *Academic Press*, the *Handbook of Medical Imaging* published by *Elsevier*, and the *Computer Science and Engineering Handbook* published by *CRC Press*. Further, I was invited to provide a viewpoint article, focused on human-computer interface aspects in Visual Analytics, for the journal *IEEE Computer Graphics & Applications* (to appear in June 2011), and I have been invited five times to the prestigious invite-only research retreat week at Schloss Dagstuhl in Germany, which is globally recognized to be "the world's premier venue for computer science – a venue where world-class scientists, promising young researchers and practitioners come together to exchange their knowledge and to discuss their research findings" (from their website). Finally, I was recently invited to give the keynote at the *International Symposium of Visual Computing 2011* in Las Vegas.

Finally, apart from these more scholarly aspects, I have also played a leading role in the international community, serving as the *IEEE Visualization* conference's general chair in 2009 – this conference has more than 800 attendees from academia and industry – and in various other leading functions, such as serving twice (2001 and 2003) as program co-chair for the *International Volume Graphics Workshop* and once (2005) as the workshop's co-chair, twice (2005 and 2006) as tutorial co-chair for *IEEE Visualization*, once (2002) as program co-chair for the *Symposium on Volume Visualization and Graphics*, and now (2011) as program co-chair for *IEEE Visualization*. Further, I serve on the editorial board of the top journal of the field of visualization, *IEEE Transactions on Visualization and Computer Graphics* and have been on numerous conference program committees, such as *IEEE Visualization* and the European Visualization conference *EuroVis.*

ACADEMIC COLLABORATIONS

I have been a very active collaborator with international researchers and institutions. These have resulted in numerous publications in both computer and domain sciences, as well as substantial research funding:

Pacific Northwest Laboratories

At PNL, I maintain a longstanding collaboration – now going into its 10-year anniversary – with **Dr. Alla Zelenyuk**, a world-class aerosol and climate researcher with whom I have been developing a substantial portion of my high-dimensional data visualization research. This effort has just culminated in a $500k NSF grant. Several papers have resulted, both in visualization and in mass spectroscopy.

Brookhaven National Laboratories

At BNL, I have recently joined forces with **Dr. Yangang Liu** and his research group. Dr. Liu is also a climate researcher but more focused on mathematical climate modeling and is now heading the highly visible FASTER project. This collaboration has resulted in a recent $200k LDRD grant awarded by the BNL lab director, meant to catalyze a larger DOE grant in the near future. Past collaborations with BNL include research on brain functional imaging with **Dr. Nora Volkow** (now head of NIH NIDA) and **Dr. Thomas Ernst** (now at the University of Hawaii) on high performance computing for MRI imaging.

University of Florida

Here I have maintained – for the past 5 years -- an active research relationship with robotics expert **Dr. Scott Banks** and medical physicist **Dr. Frank Bova**, with the goal to develop a robotic CT scanner for intra-operative surgery. We have sent multiple joint NIH proposals on this topic but so far luck has evaded us. Nevertheless, the application has inspired several research threads with my students and has yielded papers both in visualization and medical imaging.

UC San Francisco

At this institution I have collaborated with **Dr. David Agard**, one of the world authorities in 3D transmission electron microscopy which he uses primarily for cancer microbiology. Over the three years of our collaboration, I developed a GPU-accelerated CT reconstruction framework for his lab that cut the processing time from an order of days into an order of minutes. Dr. Agard had supported a student of mine and a paper with the results was just published at the reputed Journal of Structural Biology.

Loma Linda Hospital, CA

At LLH, I have been working with **Dr. Reinhard Schulte** on advanced CT reconstruction algorithms for Proton CT. Proton Therapy is an up-and-coming cancer treatment modality for which Loma Linda has enjoyed a world-class reputation. Several papers in the medical imaging area have resulted from this effort – some also jointly with Prof. Jerome Liang (Radiology).

University of Jena, Germany

Here I have been enjoying a very productive research collaboration with **Prof. Joachim Giesen**, an expert in algorithms and computational theory, mainly on using machine learning for constructing perceptual user models to devise better paradigms for data visualization. This collaboration has given rise to several journal papers, and one of Prof. Giesen's students has just been visiting my laboratory to make progress on a related joint research project.

Memorial Sloan-Kettering Cancer Institute, New York

At MSKCI, I have worked with a team of medical physicists on a project that aimed to exploit imagery obtained with the megavoltage radiation used for cancer therapy also for online CT reconstruction. This

pioneering work was a first step towards better patient treatment planning directly in the scanner. It gave rise to quite a few highly-cited papers in the field of medical physics and radiation oncology.

Stony Brook University

At my home institution I have maintained strong research collaborations with a great many of faculty members. With **Prof. Wei Zhu** (AMS) I worked on *BrainMiner*, a framework for functional brain analysis. With **Dr. Don Harrington** (Radiology) I have been pursuing new CT reconstruction techniques for low-radiation dose CT. With **Prof Erez Zadok** (Computer Science) I have had a joint NSF grant on using visual analytics techniques for file system analysis. I am presently collaborating on new research with him that seeks to use sophisticated data analytics techniques for the design of more energy efficient compute clusters and data centers. This is a joint project with **Profs. Scott Smolka**, **Scott Stoller**, and **Radu Grosu** and my role is the development of new visual analytics techniques to assist in the design of the required numerical controllers. A mid-size NSF proposal to fund this research is currently underway. With **Prof. Arie Kaufman** (Computer Science) I have maintained many funded and unfunded research projects over the years, with the most recent one being the NSF-funded $1.2M project to develop and construct the first GigaPixel immersive visualization facility of the world. Finally, with **Prof. Tamara Berg** (Computer Science) I have collaborated on an interactive framework, called *Iconizer*, which utilizes web-scale image databases to design semantically more powerful visual representations for information concepts.

Stony Brook Hospital

A promising collaboration at Stony Brook is with **Prof. IV Ramakrishnan** (Computer Science) and **Dr. Peter Viccellio** (Department of Emergency Medicine), staked out to develop the next generation of health care informatics systems. Here we seek to exploit the emerging Electronic Health Record (EHR) infrastructure to vastly improve the capabilities of emergency room physicians to reason with patient health data in the busy routine of clinical practice. My interest is to devise a visual interface that best supports the clinical workflow to enable more accurate and faster patient diagnostics. We have published several workshop papers on this topic and also applied for NIH NLM funding.

Stony Brook Center for Wireless and Information Technology (CEWIT)

A number of my research projects of mine are located at Stony Brook CEWIT and this is also where my research lab – the *Visual Analytics and Imaging Laboratory* – is located. I am particularly indebted to **Profs Shmuel Einav** and **Rong Zhao** for introducing me to Dr. Viccellio who is a tremendously energetic and resourceful collaborator.

Stony Brook Advanced Energy Research and Technology Center (AERTC)

Finally, I have also been working with **Profs. Eugene Feinberg** (AMS), **Samir Das**, **Rob Johnson**, and various other computer science faculty -- as well as National Grid / LIPA and SUNY Farmingdale – on the highly visible multi-million dollar Smart Energy Corridor project. My specific role is to test and develop visual interfaces that allow customers, both private and commercial, to better assess and plan their use of energy. I maintain a laboratory at AERTC as well.

INDUSTRIAL COLLABORATIONS

My research has also shown tremendous commercial potential and impact. Numerous companies, mainly in the medical and security imaging arenas, have relied on my expertise in high-performance GPU-based computing. One of my PhD students (with me as an academic supervisor) is currently sponsored by **Medtronic Inc.** to devise new CT reconstruction algorithms for low-dose imaging with their flagship O-arm CT scanner. An earlier outcome of these industrial collaborations was a seminal paper I authored in 2000, which has set the stage for high-performance computing in imaging using commodity hardware.

Published in IEEE Transactions on Medical Imaging, this paper has earned 103 citations so far. Other evidences for the commercial reach of my research efforts are the two (recent) patent applications and numerous technical disclosures I have filed over the years.

FUNDING

My funding record is substantial. I have won the most prestigious NSF grant – the CAREER award – and have secured funding, as PI and co-PI, from all major agencies, DOE, NSF, NIH, and from various industry and national labs for a total of over $8.2M ($3.2M my share, many single PI efforts). Of special note is my role as co-PI on the recent $1.2M NSF MRI grant mentioned above tasked to establish the world's first GigaPixel display, called the *RealityDeck*, to be installed at CEWIT. Another noteworthy effort is my participation in the Stony Brook research team that earned, in 2010, the multi-million ($12.5M) DOE grant for a Smart Grid demonstration project with LIPA and SUNY Farmingdale.

STUDENTS

This funding has been put to good use by supporting numerous successful PhD students. I have continuously maintained funding for 5-6 PhD students every year and graduated 8 so far. I have currently 8 PhD students of which one (Wei Xu) has passed her prelim exams, four (Ziyi Zheng, Hyunjung Lee, Zhiyuan Zhang, Nafees Ahmed) have passed their Research Proficiency (RPE) Exams, and two (Bing Wang, Sungsoo Ha, Puripant Ruchikachorn) are in good standing after their first year as PhD students. I also entertain active research activities (and co-authored submitted papers) with currently two MS students (Ambuj Thacker, Eric Papenhausen) and one BS Honors Student (Anuresh Mittal). I am proud to say that the great majority of my papers were co-authored together with my students, which is I find testimonial for my striving to nourish and provide mentorship to young talent.

RESEARCH PAPERS

Specifically, just in 2010 alone I have published 15 peer-reviewed papers, 5 of which in journals. In total, I have published 145 peer-reviewed research papers with 38 of these in journals and 30 in prestigious conferences of journal status. The great majority of these journals have 5-year impact factors of 2 and higher, and many of the conferences have acceptance rates of 30% or less. I have earned 3 Best Paper awards and 1 Best Hot Topic award at *IEEE Visualization* and 3 further such awards. Please see my CV and also the CV Supplement document for more detail on the research papers I have published.

SCIENTIFIC IMPACT

These research publications have shown tremendous impact. My current standing in the Citeseer index of the world's most cited CS authors is #6,101[1], and according to Google Scholar my papers have been cited around 3,000 times. This impact is further demonstrated by my remarkably high h-index of 31[2]. According to Wikipedia "The h-index is an index that attempts to measure both the productivity and impact of the published work of a scientist or scholar. The index is based on the set of the scientist's most cited papers and the number of citations that they have received in other people's publications." In essence an h-index of 30 means that a scholar has 30 papers with 30 citations or more. According to Hirsch who invented the index: "...for physicists, a value for h of about 10–12 might be a useful guideline for tenure decisions at major research universities. A value of about 18 could mean a full professorship, 15–20 could mean a fellowship in the American Physical Society, and 45 or higher could mean membership in the United States National Academy of Sciences."[3]

Coincidentally, my h-index of 31 is near the average h-index (30.7) of the Stony Brook computer science full professor faculty (standard deviation 10.7). On the other hand, the average h-index of the Stony Brook computer science associate professor faculty is 23.1 (standard deviation 6.8). To measure the h-index I used Harzing's Publish or Perish software (http://www.harzing.com) which has evolved into the industry standard. This freely-available software package uses Google Scholar and has become quite sophisticated

in recent years. It reproduces the scores quite accurately without much editing. For example, my score without editing the list of returned papers is 34 – there are just a few false-positive papers that are not authored by me. In my comparative study I did not edit the other lists beyond eliminating the most obvious false positives. I also used various possible spellings of the respective names. Hence there might be some inaccuracies, but the numbers should be in the ballpark. Finally, as a another piece of evidence might serve that I also rank at position #14 among all top-cited Stony Brook University authors according to Microsoft Academic Research [4]. More detail on the academic impact of my work can be obtained from my CV Supplement document.

[1]Generated from documents in the CiteSeer[x] database as of September 14, 2010
[2]As of June 18, 2011
[3]However, the same source also notes that little systematic investigation has been made on how academic recognition correlates with h-index over different institutions, nations and fields of study
[4] http://academic.research.microsoft.com/Detail?entitytype=5&searchtype=1&id=15818