Conduit

DEPARTMENT OF COMPUTER SCIENCE
Brown University // Vol. 25

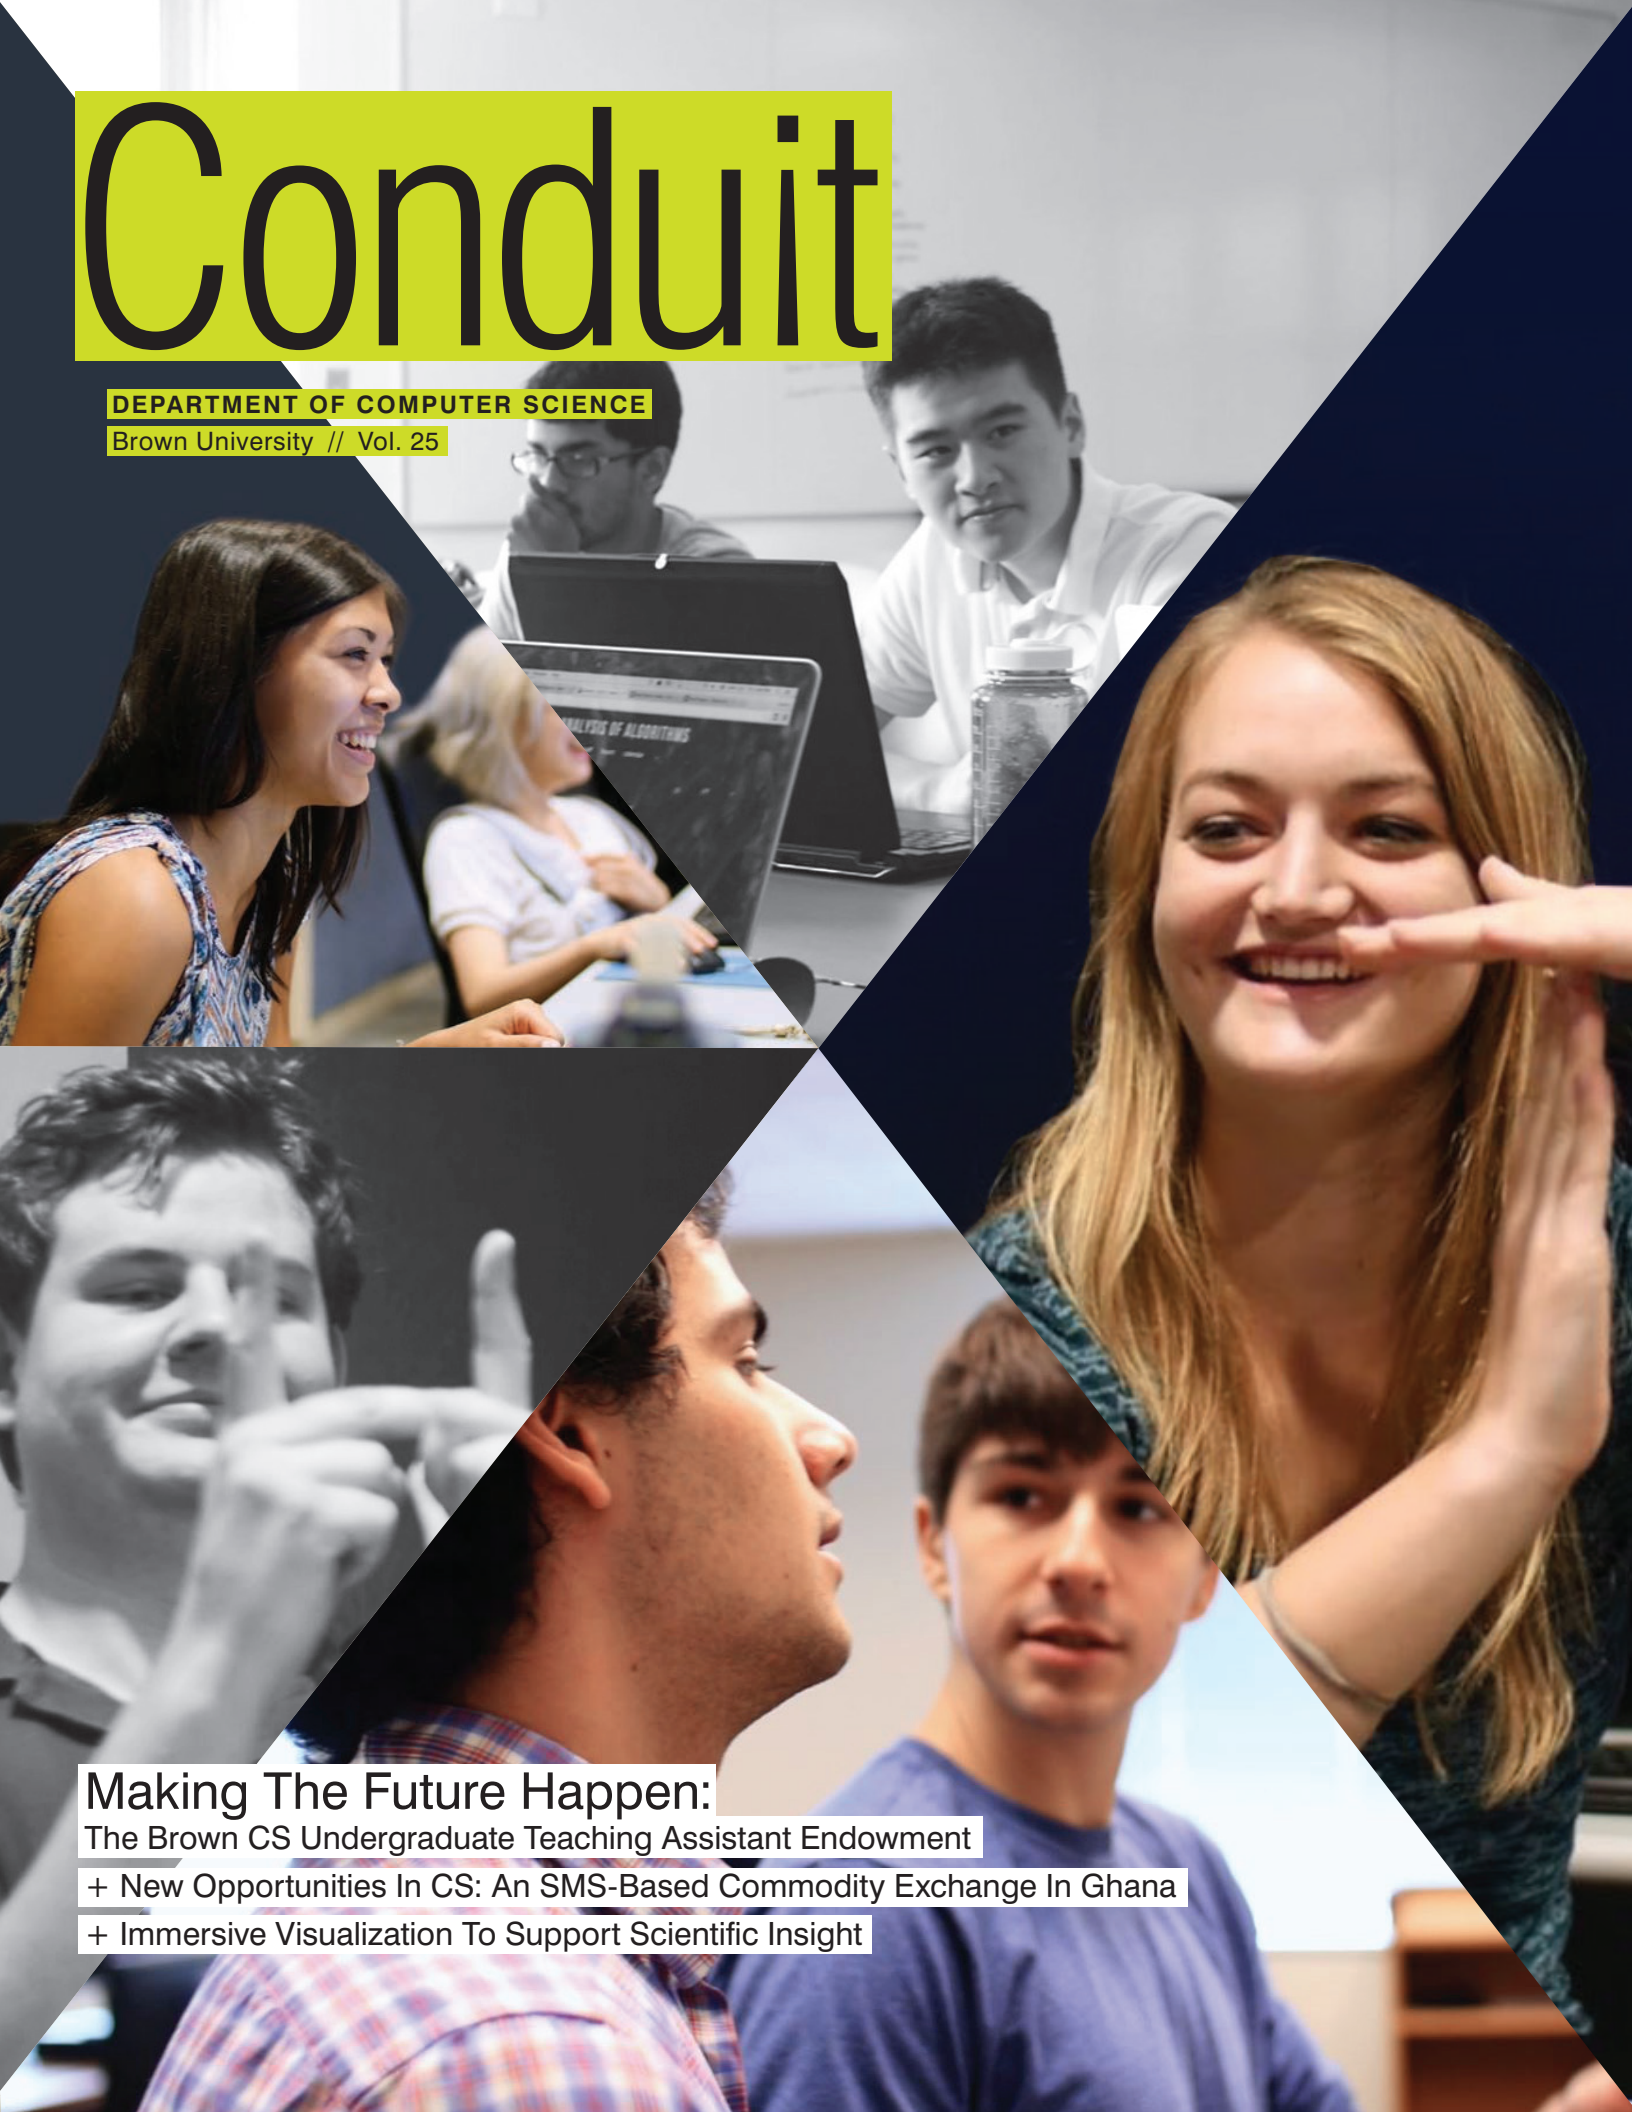

Making The Future Happen:

The Brown CS Undergraduate Teaching Assistant Endowment

+ New Opportunities In CS: An SMS-Based Commodity Exchange In Ghana

+ Immersive Visualization To Support Scientific Insight

Notes From The Chair

BY UGUR CETINTEMEL

I start by welcoming two new additions to our faculty: Seny Kamara and Vasileios Kemerlis. Seny, joining us from Microsoft Research, is an expert on applied cryptography and encrypted search. Vasileios comes from Columbia University after completing his PhD last summer, and his research is on systems and software security. Seny and Vasileios will complement and significantly strengthen our cybersecurity group and activities.

Our junior faculty continue to make us proud with their productivity and the recognition they've earned. For example, just in the first two months of 2016, Jeff Huang received an NSF CAREER award, bringing our total to 13, and Stefanie Tellex received a Sloan Fellowship, giving us a total of six.

We've been working on three major new initiatives, all of which made good progress in the last six months. Our robotics initiative, HCRI, led by Michael Littman, was officially launched in December with seed funding from the Provost, as well as some dedicated space in the SciLi. Our Executive Master in Cybersecurity Program, led by Roberto Tamassia, is officially starting in the Fall of 2016. Finally, our Data Science Initiative has been moving forward, with a new Master's program in Data Science planned to launch in the Fall of 2017. We're very excited about all the interdisciplinary activities that we're pursuing with other units at Brown. We're always considering new ways to engage with the rest of the campus at a time when computing is becoming central to the (r)evolution of almost every scholarly endeavour.

Meanwhile, computer science enrollments are continuing to grow with no signs of slowing down. While we enjoy such unprecedented enthusiasm and interest in our field, we're also pushing our limits to avoid capping classes or lowering the quality of our offerings.

Our UTA program is a key ingredient of the success of our courses, and I'm happy to report that we're making great progress toward the $10M target for a permanent endowment for this unique program. At present, we're more than two-thirds of the way and hope to get to the finish line as soon as possible. I'm truly grateful to our alums who have been giving us their money, time, wisdom, and in many cases, all of the above. In particular, I'd like to thank Norm Meyrowitz, who has shown remarkable leadership skills and resourcefulness in organizing and running our innovative Kickstarter-style campaign. Working with our alums on this campaign has been a real pleasure for me. The defining experience of my time as Department Chair so far has been the growing interactions with so many alums, all of whom have impressed me deeply with their loyalty, skills, hard work, and generosity.

Last November, we held our inaugural alum network reception in the Bay Area, which was hosted by Google. Brown CS faculty members and around 200 alums attended the very well-received event. By the time this goes to press, we'll have held a second reception: April 22 in New York City, to be hosted by Two Sigma. Going forward, we'll continue to organize such events twice a year, in two different regions, to better connect with our far-flung community and help our alums connect with each other across the generations.

We've also been working hard on improving diversity and inclusion in CS, a defining challenge for STEM disciplines and a topic of much recent discussion at Brown University and across the nation. We've been working with various CS student groups and Brown to organize our existing efforts and start new initiatives that will yield a more welcoming and unbiased environment where everyone can succeed. Our newest member of AStaff, Laura Dobler, is already doing great work in this area. Entrepreneurship has been another focus. In addition to events like Startup@Brown and a new lecture series, we're offering our first CS startup course this semester: *csciStartup*, taught by John Jannotti. Finally, I'm happy to mention that Brown CS is playing a substantial role in computer science education outside of Brown with outreach programs like Bootstrap, co-founded by Kathi Fisler and Shriram Krishnamurthi.

You'll hear more about all of these items in the months to come. If you want to stay in touch with us through social media or find out about upcoming events, please go to www.cs.brown.edu/people/alums. We have some exciting news ahead, and I hope you'll contact me with your thoughts and feedback on any of the topics above or anything else that interests you. Thanks for your continuing involvement with Brown CS.

Conduit
Vol. 25

Published by the Department of Computer Science at Brown University, *Conduit* is distributed free to all computer science alums, faculty, staff, students, and industry partners.

Queries about this publication can be directed to:
conduit@cs.brown.edu or
115 Waterman Street
Campus Box 1910
Providence, RI 02912

PUBLISHER
Industry Partners Program,
Department of Computer Science

EDITOR-IN-CHIEF
Jesse C. Polhemus

GRAPHIC DESIGN
Brown Graphic Services

CONDUIT ON THE WEB
John Bazik
Kathy Kirman

PRINTING AND DISTRIBUTION
Brett Turner, Information Services
Brown Graphic Services

Table Of Contents

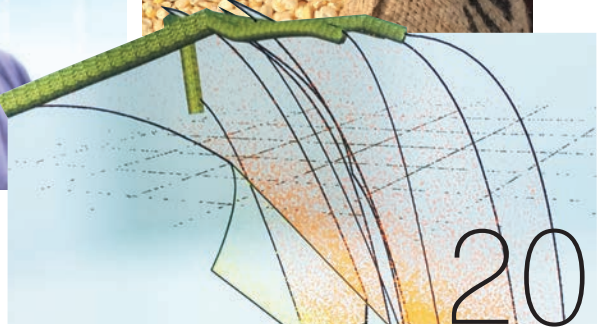

BROWN

MIX
Paper from responsible sources
FSC® C018792

Conduit is printed on Burgo's ChorusArt, an acid free and elemental chlorine free paper containing 50% recycled content and 15% post-consumer waste. You are the difference—reduce, reuse, recycle.

MAKING THE FUTURE HAPPEN:

The Brown CS Undergraduate Teaching Assistant Endowment

"I NEVER WANT TO STOP BEING A BROWN COMPUTER SCIENCE TA."

—Mike Frederickson (Technical Director, Pixar and former UTA)

The Undergraduate Teaching Assistant (UTA) program has been a hallmark of Brown University's Department of Computer Science for more than 50 years. Recently announced by President Paxson, former Provost Colvin, and Department Chair Cetintemel, Brown CS is making a landmark investment in a key element of our educational mission. We're more than two-thirds of the way to creating a $10 million UTA Endowment to celebrate this extraordinary program and make sure that it thrives for generations to come.

The campaign's web site, featuring a video of testimonials from UTAs past and present, is at www.cs.brown.edu/giving/uta.

Andy van Dam has described UTAs as a "collaborative troupe" spanning generations, making the future happen. Brown pioneered the idea of computer science UTAs, and if you work in our field today, you've probably either met one or you were one yourself. Serving as a UTA is an opportunity like no other, and comes at a critical time in a student's life, providing them with a new level of education, self-development, and often much-needed financial support.

- **1 in 9 Brown undergrads is a CS or joint CS concentrator and 1 in 5 took a CS course this year**
- **We employ almost 200 UTAs each semester, with some courses requiring over 30 UTAs**
- **At last count, over 60% of CS concentrators have been a UTA at least once**

Due to soaring enrollment and new federal employment regulations, the UTA program is dramatically more expensive than it once was. At the same time, Brown faces increased budget pressure because of its admirable need-blind financial aid policy.

Using Kickstarter-like grassroots methods usually seen in the technology sector, our campaign is more than two-thirds of the way to the $10M goal. More than 100 gifts have been received, including a $1M donation in honor of Professor Andy van Dam.

"We would like to continue to offer this opportunity for generations to come," says Department Chair Ugur Cetintemel. "There's also a real need now for an expanded UTA presence due to the exploding interest in CS courses. This endowment will be a permanent source to secure the health and future of this unique program. It's one of the best investments we can make in the education of our students, many of whom graduate to become leaders and innovators in information technology and other fields."

We're building this endowment for our UTAs to support their remarkable efforts for all time, and we'd love to have you join us. Today, you can help Brown give the same opportunity to someone who wouldn't otherwise have it. We welcome contributions from former UTAs, CS graduates, families, friends of the department, and all who want to acknowledge and support this invaluable program.

Your donation will fund what we're calling a *UTAship*. Donations can be made individually or be added to a pool named for your graduating class or for a group that you and your friends create. For $50,000, you can name a UTAship, and for $100,000, you can name a Head UTAship. You can also designate the position(s) you've endowed to support UTAships for Women in Computer Science or Underrepresented Minorities, reinforcing Brown's strong ongoing efforts to encourage diversity. You can also add a tribute at the end in someone's memory or honor.

Brown allows the payment for these donations to be spread out over five years, which may be helpful in contributing more, and there may be matching programs at your workplace that can leverage your donation. There are also significant benefits to giving appreciated stock. For more information, please go to https://giving.brown.edu/giving/appreciated-securities/.

Any gift you can make will permanently benefit our students and any industries, disciplines, and communities where they contribute. Please send any questions to utaendowment@brown.edu.

Above //
Ugur Cetintemel, Tom Doeppner, and Andy van Dam

The Virtual Student: Lee-Sien Kao, John Hughes, And CSCI 0170

BY JESSE C. POLHEMUS

THIS IS THE STORY OF A STUDENT WHO WENT TO NEW LENGTHS TO INCLUDE COMPUTER SCIENCE IN HER EDUCATION, OR A PROFESSOR WHO WAS EAGER TO TRY DIFFERENT INSTRUCTIONAL METHODS, OR A FIELD THAT REWARDS CONTINUOUS INVESTIGATION INTO THE WAYS THAT COMPUTATION CAN BENEFIT ANY TYPE OF CAREER. OR MAYBE ALL THREE AT ONCE.

ONE EXPERIMENT SUGGESTED ANOTHER

Lee-Sien Kao, an Economics concentrator who graduated last year, first heard about CSCI 0931 *Introduction to Computation for the Humanities and Social Sciences* in the fall of her junior year. "I was really interested," she says, "because I'd heard that this was a class geared to what I'd use computer science for, but the timing of it interfered with classes I needed to take for my concentration. I was also concerned about how intense CS classes are rumored to be. I didn't know if I could commit the time."

She managed to take the course a year later and was delighted with how "open-ended" it was. "I felt like the philosophy of the class was based on the potential of computing, that it was made for helping people take advantage of the endless opportunities that are out there." Without knowing it, she's almost echoing a favorite quotation from Walt Kelly that Professor John "Spike" Hughes, the course's instructor, lists on his web site: "We are confronted with insurmountable opportunities."

Before long, the opportunities became experiments: Lee-Sien and her classmates were using Google Sheets and then Python to play with data from all sorts of sources. As an early project, the class analyzed where various senators fell on the conservative/liberal spectrum based on their word choice. "Looking at the data like that was social science," says Lee-Sien. "I didn't know a CS course could teach you that." Up next? Determining whether there's a correlation between people searching for information about a country on the web and then traveling to it.

Something else was happening, too: the results of one experiment suggested another, demonstrating the need for further study. Fascinated by a hangman project written in Python, Lee-Sien decided to go a step further and run a function that tallied overused words in the thesis that she was writing. "I had to choose new words to replace 'health' so many times!" she laughs.

"And the assignments always added a new element," she says. "[Spike and the TAs] must have put a lot of time into them. They took you beyond the class: you really learned from the homework. There should be more classes like this, because it's the best way to learn a skill. Students like me, who aren't CS concentrators, need this kind of class, which doesn't go deep into abstractions but gives them tools."

She also credits Spike: "He's so interested in people doing their best. Somehow, his door would always be open when a project was about to be due, and I'd spend six hours in or around his office, asking questions and rethinking what I was doing. It was really helpful."

INSURMOUNTABLE OPPORTUNITIES

Starting with her very first impulse to take a CS course, and continuing with her choice of projects for CSCI 0931, and again with future career considerations, Lee-Sien faced multiple sets of decisions that in many ways paralleled each other. "From the beginning, I was trying to figure out what path I wanted to take," she says, "which skills I hadn't built up yet, and which skills would be useful no matter which direction I went in. There are many pros to being able to pursue your own interests, but the freedom of having so many choices makes the openness of things like studying at Brown and designing your own projects really challenging. With so many possibilities, where do I begin?"

Spike explains the challenges of insurmountable opportunities with a simple analogy: "It's very easy to say that you want to do

Above // Lee-Sien Kao

something for your first project, but not always easy for students to understand if it's something that they can answer with three minutes of web searching or whether it would make a good PhD thesis! 931 addresses this tyranny of choice through lots of student/staff interaction—every project proposal is vetted and refined."

In the spring of her last semester, Lee-Sien wanted to take another CS class, but knew that it was going to require a sizable time commitment to make it worthwhile. In the end, her thesis and a job search took precedence. "But I was thinking what a waste it would be if I didn't go on with computer science," she says. So she contacted John at the end of the semester, asking for ways to continue her CS education after she graduated.

We ask Spike: is this atypical? Absolutely. Was it something about Lee-Sien herself that pushed her to keep going? He remembers, "Lee-Sien was typical in the sense that she had a modest level of preparation for the course: work with spreadsheets, but not programing. She was exceptional in her with-it-ness for the class, following well and always asking questions about the next topic: 'What about this?' Well, just wait five minutes and you'll see!"

He recommended some books on computational principles, web resources, a few relevant online courses...and then had another idea.

What about auditing CSCI 0170 *CS: An Integrated Introduction* remotely, learning the material without credit, and balancing it with her new workload as a health policy research fellow at Dartmouth College?

THE PERFECT [VIRTUAL] STUDENT

"I'd never done an online course before," Spike says, "but I knew if I tried this, I had pretty much the perfect student, so I had some confidence. But I knew that certain student/TA interactions wouldn't be possible and that I'd have to persuade the TAs who are naturally conservative to help out: they want the course like they had it."

"I looked at it and thought I could do it," Lee-Sien says, "so I emailed Spike before the start of the semester and he talked to the TAs." Interestingly, most of the resources she used were available to all students: lecture notes, slides, a guide on how to work from and access Brown CS resources remotely. At design checks, when TAs would normally ask probing questions in person, Lee-Sien answered via email. "It was super-awesome to have it all online. The only limit was that I couldn't drop in for office hours. I was sad about that, but the rest was great. The online discussion board was really useful."

LECTURE NOTES, SLIDES, A GUIDE ON HOW TO WORK FROM AND ACCESS BROWN CS RESOURCES REMOTELY. AT DESIGN CHECKS, WHEN TAS WOULD NORMALLY ASK PROBING QUESTIONS IN PERSON, LEE-SIEN ANSWERED VIA EMAIL.

It wasn't always easy. "Getting lecture content online wasn't difficult," Spike says. "There were a few technical glitches at first, but

Above // John Hughes

nothing that was really a killer." Although the workload varied from week to week, labs and assignments generally took longer than Lee-Sien had expected. At one point, she sent what she calls "a very long email" to Spike and the TAs to say that she'd gotten lost. Before long, they had put her back on track.

SEEING THROUGH COMPLEXITY

But all of the work was interesting, she says, and some of it was very unexpected: "I didn't know we'd learn so much about not just implementation, writing code, but on the theory behind why we do what we do: how an interpreter works, the steps it takes to give you your output. There's no magic. You have to figure out yourself why you didn't get the results you wanted."

Spike makes an analogy to an earlier era of computing: "In the 1980s, with your Apple II, you could know everything: hit some buttons, see a result, and know entirely how you got it. That's changed. In 0931, students use Python, but it's still mysterious to a certain degree. Now, in 0170, with a Python interpreter, you can see through at least one layer of the complexity."

At the end? Success. "Lee-Sien completely rocked the course," says Spike. "It makes you consider what was possible from a distance, with only an occasional paragraph or two of email from the TAs and me."

Of course, a future where even a modest fraction of a Brown CS course's students are virtual isn't exactly around the corner. "This might work for five students per class," says Spike. "Working with Lee-Sien remotely took about the same time as with an average student who was present in the classroom, but this isn't really scalable. An amazing student here in Providence uses minimal resources, and we found that an amazing remote student can use an average amount of resources. But a really weak student off-campus? Disaster."

But that doesn't shut the door on the potential of virtual participation in academia, particularly in continuing education. "Maybe to say that it's useful for continuing education in general is too broad," Spike notes, "but the business of becoming computer literate seems really important to me. For one, it can help bring women into computer science. My mother and sister both learned a few computer skills, but not as much as I think they wish they did."

CHANGING PEOPLE'S THINKING

Tinkering with the delivery of the course changed Spike's view of CSCI 0170 in two ways. "On the mechanical end," he says, "I've been thinking about solo assignments versus pair programming. I'm not sure I've arrived at conclusions yet. But on the curriculum end, this experiment reinforced the idea that this course might be more than just a feeder into a CS degree. It could be your last or second-to-last CS course."

Now immersed in health policy research, in the world of surveys and statistical regression models and evaluating the effectiveness of accountable care organizations, what's on Lee-Sien's mind as she looks back?

"I don't want to forget it all!" she says. "The tools I learned apply to what I'm doing now. They can all help. I don't think I'm consciously aware of using what I learned in 0170 for things like fixing syntax errors in Stata, but I think the main direct impact has been on the efficiency of the code I'm writing: finding better ways to accomplish the same things. We learned a way of thinking that a lot of us hadn't encountered. It was such a great experience."

"In the end," Spike says, "CSCI 0170 teaches a methodical way of interacting with complex problems. Ways of thinking get changed by this course...And this is the sort of thing that happens when you teach at Brown: you can ask, 'Do any of you graduating students want to spend hours and hours in coursework next year and not get any credit?' and someone'll say, 'Sure!' I'm really glad Lee-Sien did this. I don't have any plans to repeat it, but it gives me the confidence to say to the next person who asks, 'Sure...we can try that. Why not? It's worked before.'"

scenes from the Nobel exhibit

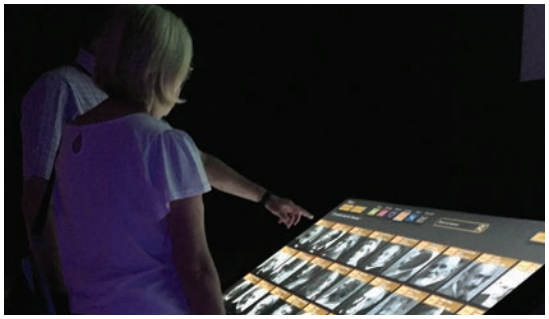

TAG (Touch Art Gallery): Student And Community Education With A Worldwide Impact

BY ROSEMARY SIMPSON

TAG NOBEL SINGAPORE EXHIBIT: CARLENE AND DAVID NIGUIDULA WINTER BREAK VISIT

Created by a group of undergraduate researchers in the Graphics Lab working under the direction of Professor Andy van Dam, Touch Art Gallery (TAG) is a Windows 8/8.1 and Web application that allows content aggregators, including museum curators and educators, to compile, contextualize, and share digital collections of artworks and manuscripts, including artworks that for reasons of size or fragility can not be viewed in person. Recently, TAG provided two state-of-the-art museum experiences at the ArtScience Museum in Singapore that focus on Alfred Nobel's final Will and the 900 Nobel laureates to date.

During Winter break, junior Carlene Niguidula and her family (Carlene's father, David Niguidula '85, is a former student of Andy) were in Singapore, where they went to see the Nobel exhibit at the ArtScience Museum. They went at the beginning of the day so as to avoid the crowds that came later on.

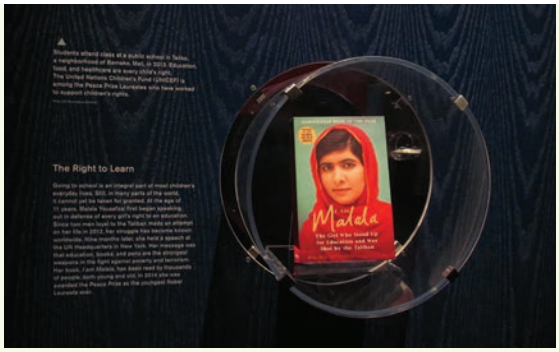

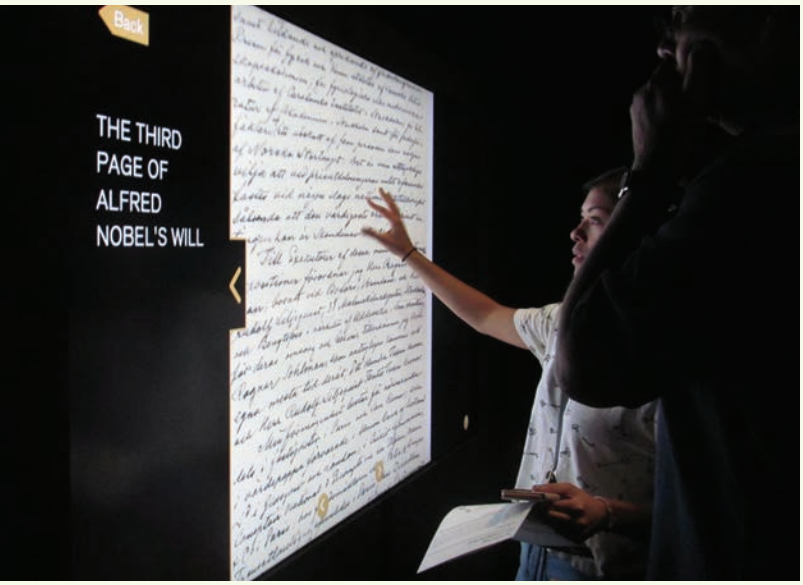

IMAGES THIS PAGE CLOCKWISE FROM TOP LEFT

- the Nobel Laureate snapshot of Malala Yousafzai
- the ArtScience Museum
- a room of Nobel Laureate snapshots
- the Nobel Will TAG exhibit

INTERACTIVE DIGITAL MUSEUMS: JEFFERSON, BOSTON, AND BROWN

Professor Steve Lubar's class AMST 1510 *Museum Collecting and Collections* used Brown's stamp collections and the Computer Graphics Group's TAG software to illustrate the power, problems, and pleasures of designing interactive digital media for physical museum elements. The Massachusetts Historical Society (MHS), located in the Back Bay area of Boston,[1] provides another example through the Coolidge Collection, a surprisingly massive compilation (over 9,500 items) of Thomas Jefferson's personal documents donated by Jefferson descendants living in New England.

MHS has mounted an exhibition ("The Private Jefferson: From the Collections of the Massachusetts Historical Society") that runs from January 29 through May 20, 2016, featuring TAG-based interactive digital experiences in every exhibition room. The designers used standard TAG facilities to explore the collection through explanatory interactive tours and multiple sub-collections, which include journals, correspondence, architectural drawings (Jefferson designed the University of Virginia as well as his own home at Monticello), almanacs, law treatises, a bibliography of his library, his first personal draft of the Declaration of Independence, farm and garden plans and diagrams, and documentation about his slavery practices.

The MHS describes the goal of the exhibit in this way: "Thomas Jefferson has been described as an 'American Sphinx.' As the drafter of the Declaration of Independence and the third President of the United States, he is one of the most famous Americans. Nevertheless, he is an enigmatic figure: an intensely private man who spent more than thirty years in public service; the spokesman for popular democracy who, at the same time, held hundreds of men, women and children as his personal property; an urbane, widely-travelled, and widely-read exemplar of the Enlightenment, who appeared happiest in a meticulously-planned environment that he had created for himself in the back country. The exhibition aims to pull back the veil and uncover the private Jefferson."[2]

[1] **1154 Boylston Street, Boston: http://www.masshist.org**
[2] **http://www.masshist.org/2012/calendar/event?event=1759**

IMAGES THIS PAGE CLOCKWISE FROM TOP LEFT

- the Bhutan recording stamp
- digitization of the Garibaldi panorama
- exploring stamp metadata

TAG@BROWN: STAMPS, PANORAMAS, AND MONASTERIES

Imagine a two-inch stamp from the Himalayan Buddhist kingdom of Bhutan that plays a folk song,[1] or a series of stamps that tell a story of how a country's political institutions have evolved over time.

These stories and the context from which they arose are the subject matter of an exhibit ("Thousands of Little Colored Windows: Brown University's Stamp Collections") that will be on display at the John Hay Library from February 10 through May 13, 2016. Hours are listed at http://library.brown.edu/hay/. The website for the exhibit, which features TAG, is http://library.brown.edu/stamps/.

The exhibition was developed by students in Professor Steve Lubar's course, AMST 1510 *Museum Collecting and Collections*, which focuses on the ways that museums care for their collections, make them available, and use them for research and education. In particular, students worked with approximately 675 stamps from Brown University Library's historic and extensive set of stamp collections, researching their history, cataloguing them, and then creating the exhibition.

Each student chose 30 stamps to ingest into TAG, creating and assigning metadata to each stamp to fill TAG with educational information. Then, some students developed narrative tours with their stamps. "In the case of the Bhutan stamp," explains Teaching Assistant Sarah Dylla, "student Anna Meyer used the tour as a platform to allow visitors to access the audio contained in these unconventional stamps. The tour plays audio from two versions of the stamp: a brief history of the nation of Bhutan, and the Bhutanese national anthem. There is transcribed text and a selection of images of Bhutan to accompany the audio."

The principle focus for TAG in this exhibition was on how stamps reflect political change over time. Students had many sessions discussing the classification issues behind the types of metadata that would be appropriate in this context, learning the hard lessons about the terminology constraints of a targeted exposition versus the development of a resource base whose purpose is general exploration. The final metadata categories provided topics about the countries, stamp colors, and subject matter that are used for search in the collection.

"The project gave the students a good taste of the complications of creating a digital interactive device for a museum," says Professor Lubar. "There's creating the metadata and thinking about how people think and search. Then there are the constraints of the platform itself: what can we put in this and how will our categories and content (i.e. stamps) actually align with how TAG functions. How will that work in the gallery space? Online you often want to have many different ways of exploring. In an exhibit space you want to narrow it down, to present specific experiences: here's the point we're trying to make, here are some objects to look at."

For the last four years, TAG and its predecessor LADS (Large Artwork Display on the Surface), which were developed by undergraduates in Andy van Dam's Computer Graphics Group, have been used in courses at Brown that range from Professor Massimo Riva's two courses on panoramas, Professor Sheila Bonde's *Global Art and Architecture* and *The Medieval Monastery*, and most recently, Professor Steve Lubar's *Museum Collecting and Collections* course described above.

TAG's original design goal was to allow users to explore artworks that were too fragile or too large to be seen in person. Brown's Garibaldi panorama, which is a 4.5' x 260' document painted on both sides that depicts the life and exploits of Giuseppe Garibaldi, is an example of artwork that is both too fragile and too large to be displayed.

[1] **From the George S. Champlin Memorial Stamp Collection from Brown University Library Special Collections**

an image of the University of Virginia from the exhibit

TAG's multimedia overviews

interior of the Massachusetts Historical Society headquarters

TAG NOBEL SINGAPORE EXHIBIT: TRENT GREEN INSTALLATION OF THE EXHIBIT

During early November of 2015, sophomore Trent Green, the TAG team leader, went to Singapore to install the TAG Nobel exhibit at the ArtScience Museum. It was his first trip outside of the US, so he took more than 100 pictures of what he saw in Singapore as well as documenting the installation process within the museum. By the time the exhibit closed, it had been seen by tens of thousands of visitors.

IMAGES THIS PAGE CLOCKWISE FROM TOP LEFT

- the public opening of the exhibit
- Singapore sights
- setup of the exhibit

NEW OPPORTUNITIES IN CS:

AN SMS-BASED COMMODITY EXCHANGE IN GHANA

BY WILSON CUSACK

For my undergraduate thesis, I AM BUILDING AN SMS-BASED COMMODITY AND TRANSPORT EXCHANGE that I am planning to work on full-time in Ghana after graduating. Below is some reflection on my journey to this and the career culture in computer science.

BACKGROUND

My interest in agriculture came when I took a year off before college and spent half of it working in India. When I was in India, agriculture was a problem everywhere I went. I heard about it while talking to the farmers outside of Kolkata, saw it while working in an arid region of northeastern Maharashtra, and experienced it personally while trying to organize the input supply chain for an agribusiness franchise I was helping a company build.

It seemed like some of the core problems were that it was really hard for farmers to connect with buyers, to get fair prices for their goods, and also to know what they should be growing in the first place (in terms of potential profitability). I had a very simple idea: why can't a farmer text a platform saying what she wants to sell, and be matched with a buyer who also texted in saying what she wants to buy?

I told a friend who worked in agriculture about this idea and he told me that what I was talking about is called a commodity exchange and that Eleni Gabre-Madhin was working to build new commodity exchanges in frontier markets. I looked her up, guessed her email, and spent a summer working in Ethiopia. She and her team had built Ethiopia Commodity Exchange, the second commodity exchange in Africa, and were then looking to begin building exchanges in other African countries.

What they had accomplished was impressive. In the six years the exchange had been running before I was there, they had facilitated billions of dollars of trades and there had been zero defaults on payments. The price of every trade was texted out to 800,000 mobile subscribers, announced on the radio, and posted on ticker boards around the country to make sure farmers knew what price they should get for their goods. Because of this price transparency, it is argued, the farmer's share of the final price for goods traded on the exchange went from 38% to 65% in the first three years of the commodity exchange's existence.

There was (and is) dissent, though. Some argued that the commodities being traded outside the exchange had seen similar shifts in shares of final price and that these shifts were actually due to cell phones making prices more transparent. They said the commodity exchange had burdened the agricultural sector, making it easy for the government to tax it, and it had been unnatural to the way things were done before.

These critiques pushed me to again consider a more decentralized model, like trading using cellphones. But the work I did with Eleni and her team helped me see that the problem of agriculture in the developing world is often a problem of markets. This is worth unpacking.

More people die from malnutrition than HIV, tuberculosis, and malaria combined. Nobel Laureate Amartya Sen supposed that people do not die from a lack of food but from an inability to entitle themselves to food. This sounds like an academic version of something that has become more or less common understanding: the world as whole produces enough calories to feed everyone, those calories are just not evenly available for consumption. While this is true, it obscures the power of Sen's point: people often starve to death not just in the context of a world that produces enough

calories, but often in countries or smaller geographic areas that produce enough calories to feed everyone.

For example, in 2001 and 2002 farmers in the south of Ethiopia produced a surplus crop that could have met normal demand for two years. There was so much maize that the price dropped by 80% and farmers left their maize to rot in the field because harvesting it was not worth it. In July 2002 a food crisis was announced in the north of Ethiopia with 14 million people at risk of starvation.

The failure can be attributed to market opacity, preventing price and crop information being shared throughout the country, and a lack of trusted rules of exchange that would lend enough confidence to farmers for them to trade over vast geographic distances. This is a market problem, but solving it also helps to improve food production from a technological perspective. Making farming more lucrative helps existing farmers adopt new technologies and grow their farms, but it also attracts businesses to the sector, ideally boosting food production and improving food security.

The fall after working with Eleni, I transferred to Brown and began studying Computer Science, which opened up new opportunities. I spent the summer after my sophomore year working for FarmLogs, whose web and mobile app is now used by one in three farmers in the US, and I spent this past summer working in logistics for a large agricultural company.

THESIS

This past summer I began talking with Rodrigo Fonseca, who is my thesis advisor, about potentially doing something related to commodity exchanges and logistics. I still had this idea in the back of my mind of building a less formal commodity exchange, one that would simply let people trade on cellphones, and, luckily, Rodrigo was up for working on it.

I spent the fall building the platform, getting help from many people in different departments, and then went to Ghana for nineteen days in January to pilot it. I decided to focus on maize, as it is the most important cereal crop in the country (55% of all output) and was reported in a 2012 study to have marketing margins (middlemen's margins) of 50-70%.

I divided my testing into two phases. First, I went to villages and told farmers that I would buy from them right then, paying above market, if they used the platform to facilitate the transaction. I then went to buyers in the urban markets and told them that I would sell to them right then, below market, if they used the platform to facilitate the transaction. I traded six bags (100kg each) like this.

I then told everyone that I bought from or sold to that I would be back in a few days and that if they wanted to buy or sell anything before then, they could text in to the platform to do so. This was sort of the holy grail of the trip: an open trading period. It was the holy grail, first, because I wanted to see if I could actually create a new market and have it function. But, second, because I wanted to know if people would use the platform when I wasn't standing right in front of them.

I've done enough surveys in the developing world to know that what someone tells you when you're a foreigner just meeting them for the first time can mean very little. I was giving both buyers and sellers a good price. There was high incentive to just nod and tell me what they thought I wanted to hear and get to the part where I gave them the money/maize.

A person's time is very valuable to them anywhere in the world, but especially in the developing world and especially in this context. Almost everyone we traded with was illiterate and so either had to have their children or a friend text the platform for them. And though I tried to make entering an order as simple as possible, there was certainly a learning curve. All of that to say, if they continued to use it after I left, that would be pretty high praise. Amazingly, they did!

Above // Riding with Wisdom, a local university student

Right // Rebecca, our first buyer!

In just a couple of days, buyers posted bids for 5400kg and sellers posted offers for 1000kg, and we delivered 500kg.

WHAT'S NEXT AND THOUGHTS ON OPPORTUNITIES IN CS

I'll now spend the spring semester further honing the platform, and I hope to go back to Ghana in July to work on this full time. I am incredibly grateful for the role of the university in the development of this project: for equipping me with the skills to build this and also for providing an excellent environment to explore the idea in an interdisciplinary way.

The concept seems so simple and obvious that many think there must be some key reason others haven't already done it. It is a good question to ask when exploring any idea, but after working for companies and aid organizations that have prestigious reputations, I've found that there is not often a good answer to it, even from the people who you'd expect to have it. In fact, especially in the developing world, often the answer is that the reason nothing has been done is that the people who experience the problem are unfortunately not the ones who have the means to create the proposed solution. In this case, rural farmers in Ghana do not know computer science.

I feel really lucky because I have the chance to work on a problem that I really care about and that I think is really important for the future of the world. I feel lucky because I think I was lucky. I don't think I would be working on this if I hadn't spent time living in India, and I was really lucky to have parents that encouraged me in taking a year off and supported me financially. To be clear, working (volunteering) in India for five months is a lot cheaper than a semester at almost any college in the US: I spent less than $6k in total, including airfare. But such a price is obviously still out of reach for many.

My career trajectory is unique, at least among the CS undergrads that I know, but I wish that it wasn't. I wonder why someone has to be really lucky to care enough about a problem affecting a majority of the world's population to forgo the normal CS career paths. Maybe the answer to this is obvious: someone has to be really lucky because of what I said in my above paragraph. It takes going to these places and seeing these problems first hand to care enough, and most don't have that opportunity.

But I don't think this is a satisfactory answer, at least not for my community. Though it's true most can't have such an opportunity, many at Brown could. First, many CS undergrads make much more than 6k in a summer internship. And, second, a disproportionate amount of our student body come from families that could comfortably afford such an experience.

If the opportunity is available and not taken or not encouraged, if many of the most well educated and privileged people in the world have zero experiential familiarity with the problems experienced by the majority of the people in the world, then I think we should question whether there is not something wrong with the way we are educating ourselves.

We bought from the woman above and her husband individually: she was emphatic that she wanted to text in with her own phone to ensure that she was the one getting paid

But aside from the fact that many in our community could have such an experience, should we really accept that the only way someone can "care enough" about one of these problems is to experience it first hand?

I wonder if our not pursuing these majority-world problems, if I can call them that, has a lot to do with our expectations of each other, or worse, a sheer lack of thought and imagination. CS students have been equipped with a set of tools in their undergrad career. These tools are being applied to problems in every sector across the world with dramatic results, and I think there remain swaths of serious problems that are still untouched. I think we should feel a sense of obligation and more regularly ask each other, "What are you going to do to make the world a better place?" There are obviously a lot of good answers to this question, and I am not saying that a "normal" CS postgrad career is not one of them, but at this point I don't think we are even having a conversation about it. There are many people I talk to who seem moderately excited about what they are doing after graduation, but mainly see it as a stepping-stone to some unknown better thing that will come after. There is a sense that people are being forced into these careers: offered a lot of money and a quickly approaching deadline, the offer is taken, as if between classes and our social lives we really don't have time to thoroughly think through what we want to do with our lives or what we think is worth doing at all.

"I AM **INCREDIBLY GRATEFUL** FOR THE ROLE OF THE UNIVERSITY IN THE DEVELOPMENT OF THIS PROJECT:"

I think we are crazy if we think anyone else is ever going to slow down our life for us, make time for us to seriously think through the important questions. We can only do that for ourselves, and I worry that we aren't.

PETER WEGNER: A Life Remarkable

Students who have recently arrived at the CIT may only know Peter Wegner, now retired for more than a decade, through his vigorous participation in colloquia. Others will be familiar with his research in programming languages. But as he hears a half-century at Brown, not just the course of his career but the story of his life bears telling. We open a three-part series by inviting Peter to tell the story of his early years in his own words.

Peter was born on August 20, 1932 in Saint Petersburg (then known as Leningrad) to Austrian parents, Leo Weiden and Hermine Wegner. They had met as members of a Viennese Jewish left-wing group that included Bruno Kreisky, who became Chancellor of Austria in the 1970s. After their marriage in 1931, Peter's parents emigrated to Russia, because they believed that the Russian Communist system would provide a better life for Jews than the increasingly anti-Semitic Austrian society. They named their son "Putilo" after the famous Putilov steelworks where Peter's father was first sent to work when he arrived in Leningrad.

In 1937 Peter's father, who had become a senior librarian at the National Library of Science, was arrested during a Stalinist purge by the Russian secret police (NKVD). Peter remembers two men coming to the apartment one evening and taking his father away. He never saw his father again. He now believes that his mother discovered in November 1937 that her husband had been executed after a trumped-up trial that charged him (along with many other visiting intellectuals) with sedition. But Peter's mother could not bear to admit that the system they so admired had executed her husband. Instead, she later claimed to friends and relatives that her husband had gone to Spain to fight against Franco and had been killed in the Spanish Civil War.

Peter's return to Vienna with his mother in January 1938 proved to be a case of "out of the frying pan into the fire." Two months later, on March 12, 1938, Hitler marched into Austria and annexed it in a move called the *Anschluss*, and Peter's mother had to escape at once, because everyone knew that the first people Hitler would hunt down were the Austrian Jewish Communists—whose names were readily available at the communist headquarters in Vienna. Years later, Peter learned that his mother, along with several friends, had escaped by skiing cross-country into Switzerland. She could not take five-year-old Peter with her on such a trip, so she left him with her mother, Jetta Wegner. From Switzerland she was able to obtain a visa for England by agreeing to work as a live-in maid in someone's house.

The day after the *Anschluss*, watching from a window in his grandmother's apartment overlooking the *Praterstrasse* (a main street in Vienna), Peter remembers seeing a parade driving up the street, with cheers of "*Heil Hitler!*" from Austrians lining the street. The parade included a car in which Hitler was standing. At age five, Peter did not understand what was going on, but he remembers his grandmother telling him: "That is the Chancellor of Germany, who is a very bad man."

Soon after the *Anschluss*, many Austrian Jews, including some of Peter's relatives, were arrested and sent to the Dachau concentration camp near Munich. When they returned six months later, they had lost a lot of weight, and Peter heard them tell his grandmother how badly they had been treated. During a walk in the park with his grandmother, Peter recalls being chased by some children who had found out he was Jewish and wanted to beat him up. He managed to escape, but from then on he learned to be careful when out on the public streets.

The Nazi authorities soon deprived Austrian Jews of their nationality and expelled Jewish children from the state schools, so Peter (whose family were secular, non-observant Jews) had no choice but to attend a Jewish school at the local synagogue.

On November 9-10, 1938, there occurred *Kristallnacht* ("the Night of Broken Glass"), when Nazis in Germany, Austria, and Czechoslovakia (occupied by Hitler before the outbreak of World War Two) broke the windows of Jewish shops and burned down many synagogues in those countries. The next day, Peter remembers being taken to school, and

finding that the local synagogue had burned to the ground. He also recalls going for a walk in the park with his uncle and seeing Jewish women being molested by Nazi soldiers and Austrian police.

Though *Kristallnacht* informed the world of the brutality of Nazi persecution of Jews, the only country to offer assistance was England. In December 1938, the British Parliament passed a resolution offering political asylum to 10,000 Jewish children from Germany, Austria, and Czechoslovakia by letting their parents send them to England on special trains called *Kindertransports*, provided that their financial support was guaranteed by a payment of £50. Peter's mother found a financial sponsor for him (Marks & Spencer, the well-known Anglo-Jewish retail chain) who put up the money.

Peter at six years, when sent to London

On April 25, 1939, Peter's grandmother took him to Vienna's *Westbahnhof* where, along with about 300 other children aged between three and 17 years, he boarded a train bound for London, organized by the Quakers and the international Red Cross. The *Kinder* were supervised by young Jewish adults, who were required to return after accompanying them to the Dutch border, or no further transports would be permitted. Peter's train left Vienna at 5 pm, passed through Munich at midnight, and reached the Dutch border at 9 am the next morning. The train had been locked throughout the ride through Germany but the doors were opened once they crossed the Dutch border, and Peter remembers being allowed to get out and walk on the platform, where he received a mug of hot cocoa dispensed by kind Dutch ladies showing their concern for these Jewish children who had been persecuted by the Nazis.

The *Kinder* continued their journey to the Hook of Holland, crossed the English Channel during the night, and arrived at Liverpool Street Station in London in the early afternoon of April 27, two days after leaving Vienna. (Much later, in 2006, a realistic bronze sculpture by Frank Meisler, a former *Kind*, which depicted the *Kinder* arriving with their suitcases, was erected by the Association for Jewish Refugees outside the station entrance as a memorial.)

Peter was luckier than most of the children, as his mother came to meet him at the station and brought him to a house in the London suburb of Willesden, where she was working as a live-in maid. Her employers allowed Peter to stay there for four months, during which he attended a local primary school and learned some English. But after that he was sent away to boarding school, as his mother could not keep him indefinitely at the house.

BUNCE COURT SCHOOL

The Jewish Refugee Committee in London agreed to pay Peter's tuition, room, and board at a primarily Jewish co-educational boarding school in Kent, about 50 miles from London. Its headmistress, Anna Essinger, had transferred her school from Herrlingen (near Ulm in Germany) to England in 1933, realizing that a mainly Jewish school could not survive in Hitler's Germany. She was ethnically Jewish but had adopted Quaker principles during World War One while a graduate student in German at the University of Madison, Wisconsin. Once the school was set up in England she was able to accept *Kinder* and other Jewish refugee children and to hire some excellent teachers who had fled from Germany to England. The headmistress insisted that all classes be taught in English, and encouraged the children to speak English rather than German among themselves. (She had been born in Ulm in 1879, the same city and year where Einstein was born, and would die in England in 1960, outliving Einstein by five years.)

Peter entered Bunce Court School, near Faversham, Kent, in September, 1939, and stayed there for nine years until the summer of 1948. When war broke out, the government required the school to move elsewhere, because it was thought that if the Germans invaded England, they would probably land in Kent. So in early 1940, Bunce Court School was moved to Trench Hall near Wem, Shropshire, close to the Welsh border, and allowed to return to Kent at the end of the War in 1945. The dining room at Trench Hall had a nice view of the Welsh mountains, and Peter enjoyed hiking through the beautiful Shropshire countryside.

Peter visited his mother occasionally in London during school vacations, and he remembers spending one night in a London Tube station to avoid the bombs during an air raid. His mother sometimes visited him at school, and his schoolmates remembered her because their own parents remained in Europe and could not visit their children.

As a long-term student, Peter was often consulted by his teachers about administrative actions, to gauge the reactions of students. He was known to be good at mathematics and science, and passed his School Certificate examination with high marks at age 15 in 1947, before leaving the school when it closed down in 1948. After leaving school, he continued to interact with his former classmates, and visited Anna Essinger in her old age to comfort her and read to her.

Meanwhile, Peter's mother had earned some money running a clothing store, and was able to purchase a small house near Ascot, about 25 miles from London, and Peter moved there in 1948, travelling up to London every day to study at the Regent Street Polytechnic, near Oxford Circus. During the summer of 1949, Peter went hitchhiking through Europe, climbing Mont Blanc in July and visiting relatives in Vienna in August. But on his arrival in Vienna, he was informed that his mother was seriously ill in hospital and that he must return by train to London immediately. On arrival, he went straight to the Middlesex Hospital to see his mother, but found her to be comatose. The following morning he received a phone call informing him that his mother had died during the night, and the hospital asked him to arrange for her burial.

Peter's story will continue in Part 2, to be featured in the next issue of Conduit.

Immersive Visualization To Support Scientific Insight

BY BRUCE DONALD CAMPBELL
(RISD, BROWN CS VISITING SCIENTIST)

What do you get when you bring together a seasoned popular art school illustrator, an accomplished computer graphics scientist with a new virtual reality surround theater, a respected scientist with a compelling field of study, and twenty-one eager participants in a semester-long course focused on how to support science through design in virtual reality? You get plenty of material for a *Conduit* column at a minimum. And, if the you is a year-long visiting scientist to Brown CS, you aren't about to pass up the opportunity to observe as a participant.

A VR DESIGN FOR SCIENCE COURSE

I had heard of the cross-listed course known as *VR Design for Science* from a "Visualization Viewpoints" article Dan Keefe wrote for CG&A in 2005.[1] In that article, Dan hypothesized:

> "[Virtual Reality is] one of the technologies that can most benefit from artistic insight, since guidelines for good visual depiction are far less developed in the unconventional visual space of VR than in more traditional media".

Having watched artists make huge contributions in many research projects at the University of Washington's Human Interface Technology Lab in the 1990s, that hypothesis is certainly one I continue to feel confident is evident through observation, at least when providing narrative or illustration is useful. Based on Dan's article and discussion with the course's long-term providers from Brown CS and the Rhode Island School of Design's (RISD's) illustration department, I expected the class to facilitate further investigation.

David Laidlaw and Fritz Drury co-taught the 2005 class around the science of arterial blood flow and a winged mammal's ability to fly. For 2015's class, they enticed paleontologist Stephen Gatesy to share his work on foot movements in Triassic theropod dinosaurs. The leap upgrade of the class' VR reality theater, called the Yurt for the shape of the physical projection space, added to the excitement of a new science focus.

Needless to say, the course reading list provides a large bibliography of relevant material to exploring the science of bird and dinosaur tracks in VR. I found the most useful readings to include historical perspectives on visualization from art,[2,3] science,[4,5] and computer science.[6] To weave those themes together, useful reads included publications on human perception,[7] human-computer interaction,[8] and the hardware/software affordances of the new Yurt. In front of that background, science publications presented perspectives relevant to dinosaur science and related bird science.[9,10]

Right, Figure 1 // Students created weather maps to explore multivariate data representation. In this example, Timothy Blaine-Kuklo represents temperature, wind, cloud cover, pressure, and front delineation on a single map. The Yurt allowed the artist to explore intermediate steps and stack them to try out encoding combinations while also then discussing them interactively in critique. All article photos courtesy of Johannes Novotny.

By having experts and emerging experts in these subjects co-present in a shared studio (with a minimum of two, two-hour required sessions a week), the course also provided an ideal configuration for distributed cognition: tangible objects for interaction (casts of dinosaur tracks and bird anatomy), artifacts of embedded cognition (diagrams and multimedia exhibits), and rich social processes for creating new objects and artifacts to support a scientific pursuit (as a well-designed studio class readily provides).

Through Steve's introductory presentation, I quickly confirmed my expectation that dinosaur science is hard. While many uncover and describe new fossil skeletons, his work involves the tracks dinosaurs of varying anatomy made in varying substrates using varying kinematics. Steve studies modern day birds like turkeys and guineafowl and sets up experimental environments where they walk through designed substrates—all done with an eye on optimal x-ray recording of foot movement underground. Track simulations that can be generated from reconstructed foot motion involve a large number of variables. Visualizing the data involves choosing among the variables and encoding them to co-exist in a meaningful perceptive presentation.

To help develop their own personal process for making visual trade-offs, participants first grappled with a popular visualization opportunity: weather map generation (see Figure 1). Discussions about scope and relevance shed light on choosing which weather characteristics to consider including in a visualization. Discussions about human perception and cultural implications shed light on how to represent those characteristics visually and texturally.

Students were asked to bring forward the lessons learned from weather visualization to create hypothetical dinosaur foot movement visualizations aimed at providing useful science support—helping scientists rethink their hypotheses, clarify their experimental goals, and even alter the way they collect data. The potential of possible support became evident through the diversity of solutions participants suggested, sketched, created in the Yurt, and then critiqued over many weeks. Participants critiqued visual reasonableness based on watching bird

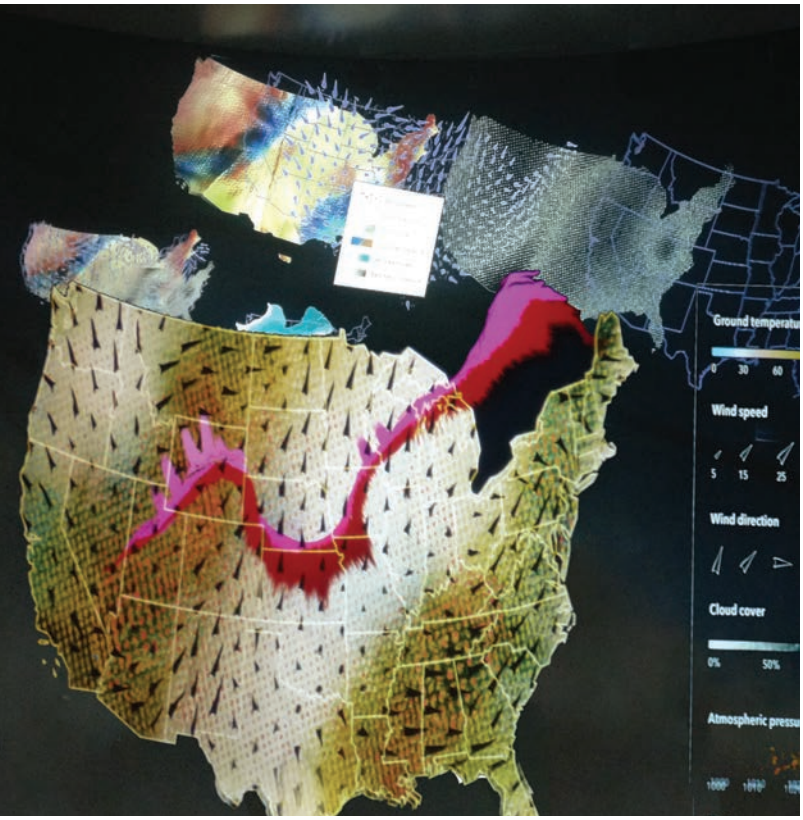

Below, Figure 2 // Students created visualizations that honed in on an aspect of foot movement in order to verify their understanding of what active bird footprint generation data should provide for visualization. In this example, Johannes Novotny traced points of foot anatomy through a walk cycle and used color to represent displacement of the substrate. Orange represents displacement of the substrate at one point in the down motion and yellow represents displacement slightly later. Red and purple were added later to illustrate upward motion. As in many cases, the resulting visualization appeared as an attractive sculpture floating in the Yurt's 3-D projection space.

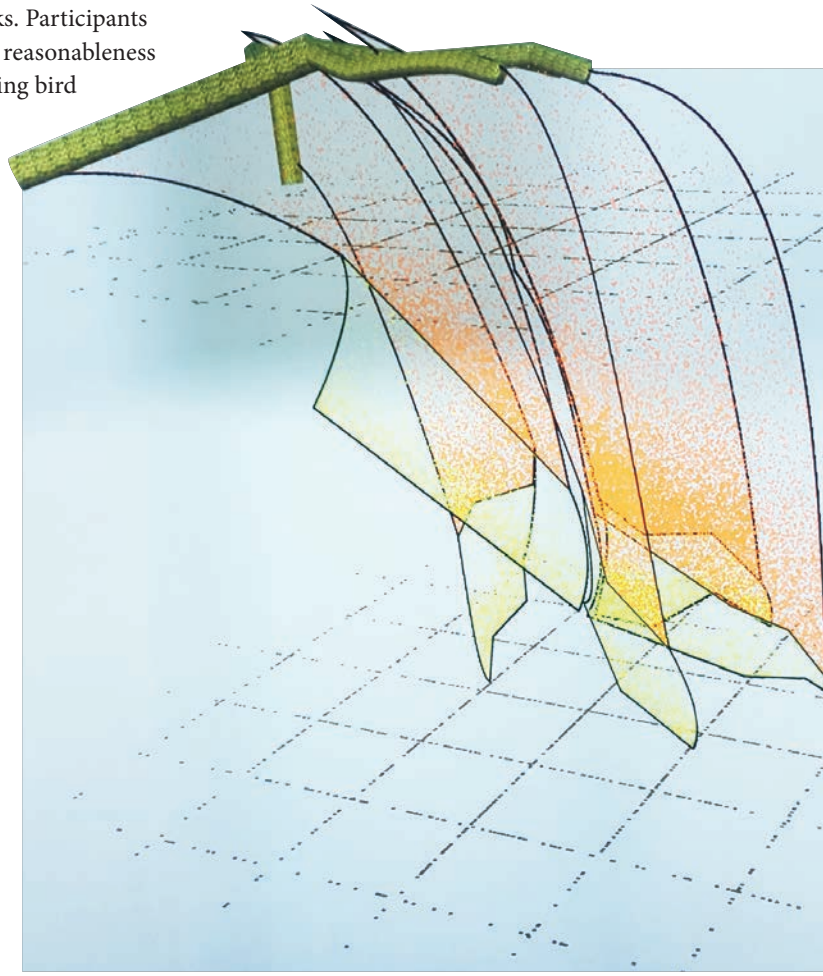

walking movies and data simulation videos that demonstrated foot movement and substrate deformation. Incorporating those hints, some students created visualizations that focused on structure, while some focused on motion, and some focused on representing the forces involved in the interaction of both (see an example in Figure 2).

During most studio sessions, class participants filled the Yurt to participate in critique sessions. Heightened by the retinal visual fidelity, bold color representation, and a sense of being immersed inside data and narrative representations, I appreciated that some students had tremendous eye-hand coordination in creating detailed immersive visualizations. I noticed others compensate through a control of descriptive language they used to communicate about missing or deformed features to the mind's eye. I found some had an inherent ability to sketch ideas out quickly, describe fuller implementations verbally, and then readjust their sketches to consider feedback—even catching themselves mid-statement to adapt to a new train of thought that came from describing an idea to us. Over time, as students improved their ability to dynamically manipulate the point of view, scale, and orientation of their creations, we noticed their ability to promote ideas that required manipulation with VR peripherals (wand, tracker, button-based devices) to fully integrate.

The final project assignment required students to expand their visualization ideas into science-support tool prototypes. Students created storyboards to explain the tool user's experience, identified menus and virtual objects that would allow the user to make meaningful choices, and then built a visual representation of a practical state of the tool in action in the Yurt. Figure 3 shows one example of a final project. Before reaching this state in the tool, the user has selected from different bird anatomy options, created a substrate shape and consistency, chosen a locomotion speed, and suggested a tile density for data collection purposes.

I noticed each student's evolution in developing a unique confident personal style as they iterated upon the role of lead commentator when considering an idea in critique. Confidence seemed facilitated by mastering all the aspects of the course suggested by lectures and reading list and incorporating the more effective aspects of each participant's emerging abilities. Confidence became evident in the growing comfort and ease demonstrated when suggesting a relevant contribution to critique. A social comfort came with familiarity of each other and the process. While some strengths of each participant could be anticipated from their previous experiences and areas of expertise, the willingness of students to offer critique outside of that expertise while among those who were their senior gave me the ultimate satisfaction as an observer. When one student struggled with a personal weakness, other participants assisted to promote their perspective, often by temporarily taking over the presentation. The supportive demeanor of the career accomplished scientists, artists, and computer scientists created a positive environment in which the students could evolve their confident style.

Below, Figure 3 // A scientific support tool prototype provides a data visualization for a hypothetical running guineafowl's footprint sequence. Three depths of substrate are disrupted by foot interaction and uniquely colored tiles are displaced to represent their final resting spot. Visual filaments show the paths taken by each tile and tile orientation identifies the direction of last movement. The initial concept offered for critique is inset in the upper-left.

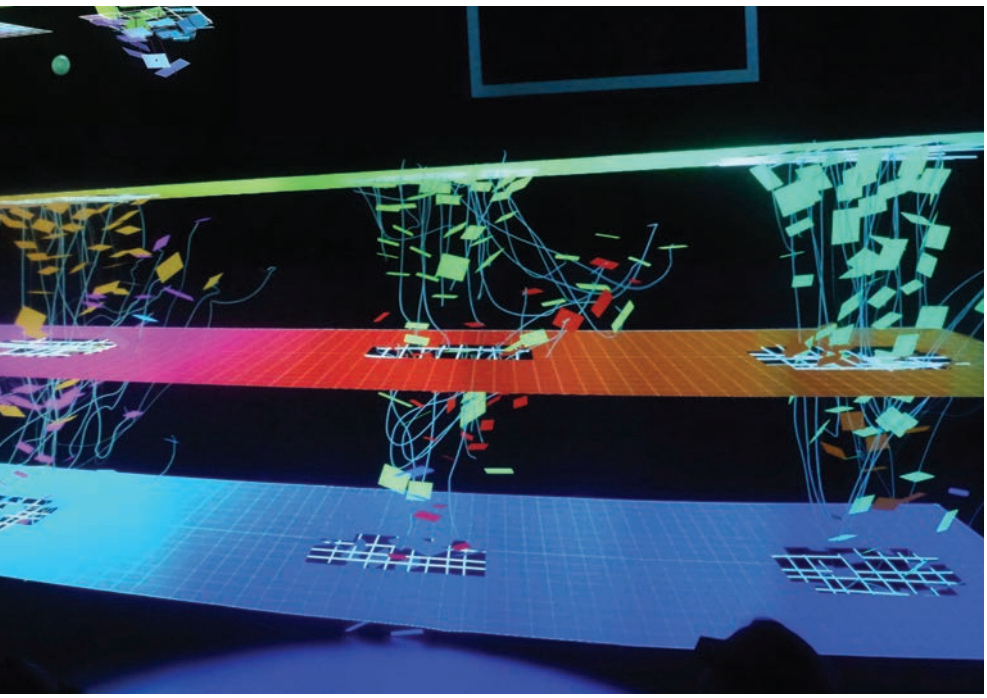

NOTES ON THE TECHNOLOGY

Fritz consistently provided a perspective that the materials involved in illustration offer strengths

Images from student projects

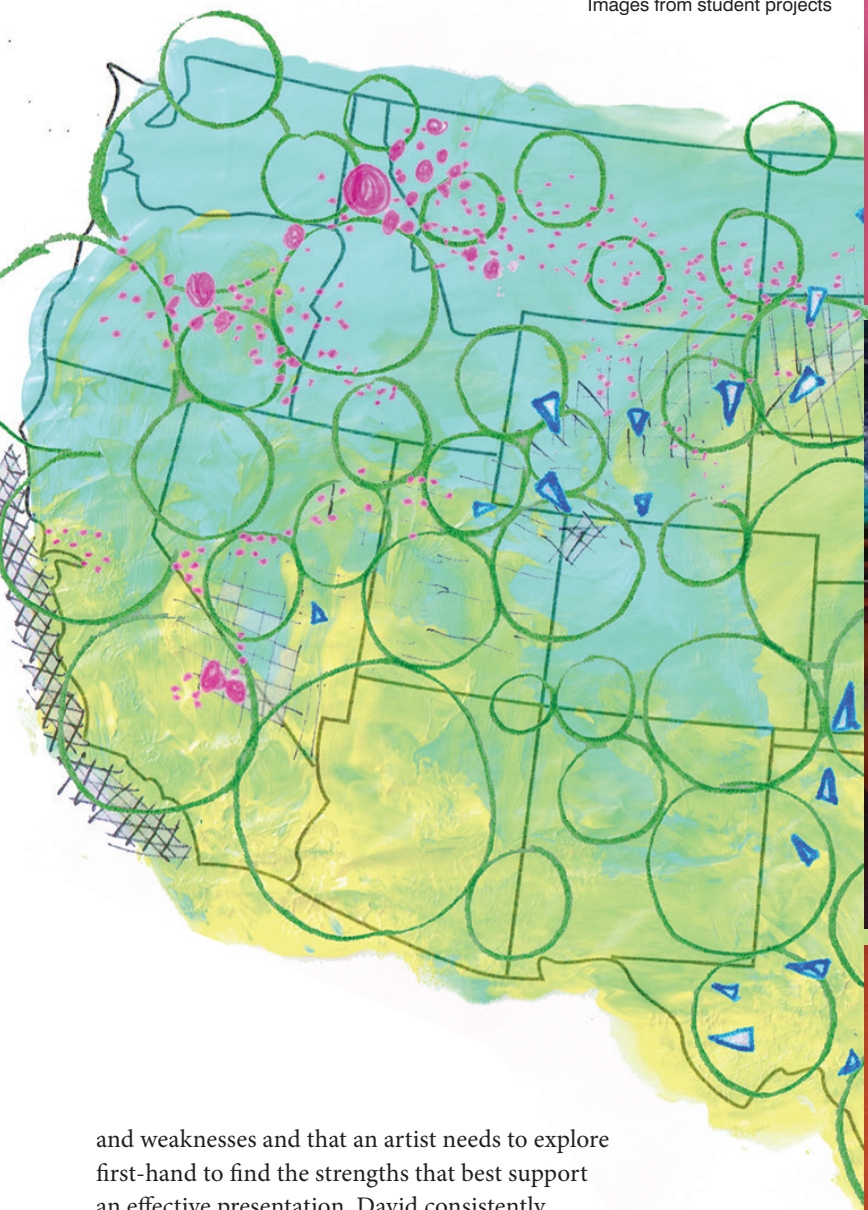

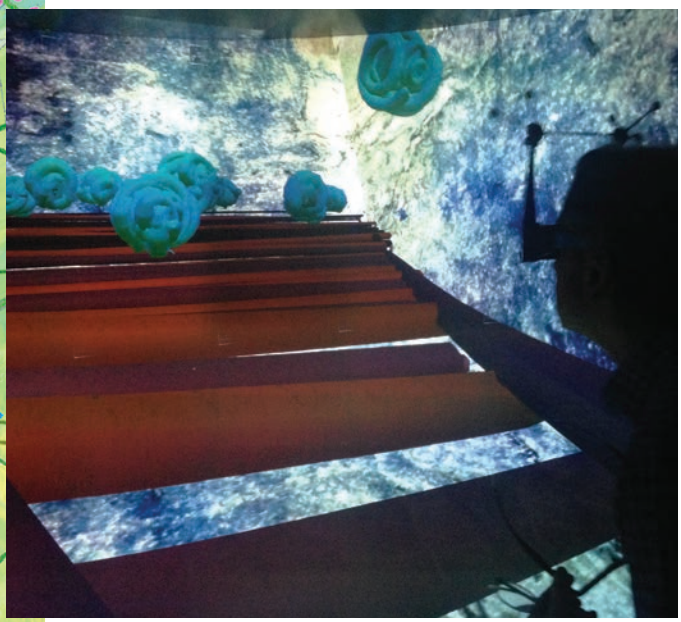

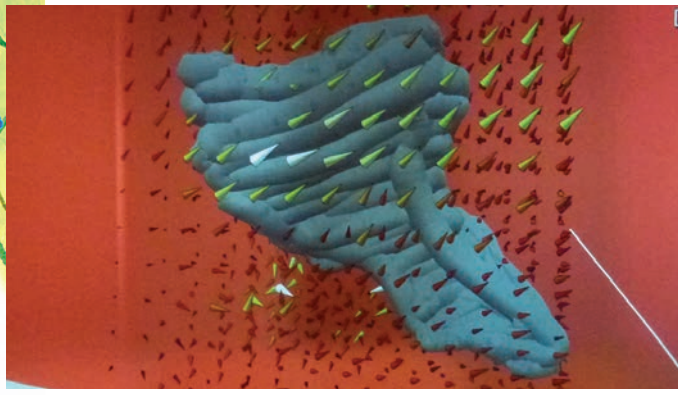

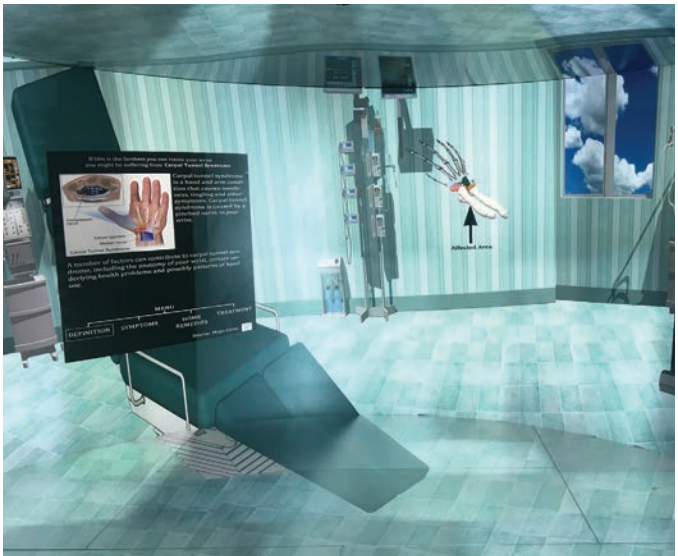

and weaknesses and that an artist needs to explore first-hand to find the strengths that best support an effective presentation. David consistently reminded participants that technologies provide affordances that can only be taken advantage of when perceptible and cognitively accessible. The Yurt's highly responsive 360-degree retinal resolution provided by 69 high resolution, high luminosity projectors, and the CavePainting software that runs within it,[11] provided a unique technology not readily found in any close configuration outside of the class—the most extensive use of the Yurt since it completed construction in September 2015. I watched as students took ideas they sketched out with traditional artist materials (often paper-based) and mature software tools on common computing workstations (Photoshop and 3-D modeling tools), and mapped those sketches to useful

Above, Figure 4 // Emily Schnall took advantage of the Yurt's high fidelity to immerse us in a photorealistic dinosaur habitat. Jurassic Park and other dinosaur films had provided us previous experience with dinosaur environments but the freedom to choose our own visual experience through such an environment feels different and evokes a different type of emotional engagement—one perhaps hard to put into words and perhaps more valuable in the artist's domain.

representations in the Yurt. Students found constraints in the CavePainting software a relief at times given the enormity of choices and their personal time constraints. Students found the fidelity of the experience expressive in describing what 'could be done' with different software available. The Yurt's ability to simulate the look (and even suggest the feel) of a wide variety of materials seemed critical for the class' success. Through tackling opportunities and constraints, participants provided valuable feedback for considering the evolution of the Yurt going forward.

THOUGHTS FROM THE EXPERTS

Fritz reminded that the course always takes on the problem of working with a kind of immersive experience that is really new and fundamentally different from the previous territory of traditional design fields, but also from the quotidian physical world. That's our challenge and our inspiration. When asked about what made the 2015 class memorable, Fritz called out the photorealistic background objects students included to help immerse our focus into the human relevance of the data being explored (see Figure 4). He reports that the new Yurt significantly improves the illusion of immersion provided by a color range far beyond what he had to work with in the older Cave, and that the Yurt is big enough for the everybody to comfortably get inside for critiques, an essential component of any collaborative design process. The improvements provide an increased sense of engagement and suggest added value for supporting context, metaphor, and narrative.

Fritz and David acknowledged that they were fortunate to have some graduate students in computer science who approached the technicalities of scientific visualization in a substantive way, but also some illustration students who pushed the boundaries of using the virtual space of the Yurt to create elaborate immersive environments. Another group of particularly cogent thinkers were able to construct "clean" designs that facilitated usability and visual clarity. This variety reflects positively a number of potential applications of the skill set that participants can acquire in class: visualization for scientific research; creation of immersive displays for public use; interactive interface design, and pure entertainment. Although they stress the importance of designing for the scientist user, the course embraced all these corollary skills and interests.

David echoed the success of many of these objectives he had when creating the Yurt and was still integrating the feedback provided by the class experience. In its first three months, the Yurt has been most successful when using a megapixel image viewer or high-resolution video player or building new applications upon the legacy process—the one used to evolve the previous Cave's technology. As the talented support staff looks to support new and improved methods for getting applications running in the Yurt, David wondered about the metrics to use in pursuing a promising best solution compared to an ecosystem of marginally improved ones (there are many candidate approaches the support staff has been considering). The pursuit is complicated as many applications benefit from specialized trade-offs in

data size and format considerations, messaging strategies, performance considerations, and interaction facilities.

Steve reminded me that currently a track simulation involves over 3 million particles. He was fascinated to see the object of his research through the eyes of bright, creative students. Sometimes a key point that he thought was obvious from his papers and presentation was not communicated well enough. In other cases, an extremely subtle aspect that he never expected anyone to understand was picked up perfectly. Although it's always instructive to get feedback from non-scientific audiences, that feedback is typically verbal. Artists creating images and material for the Yurt provided visual feedback that was extremely powerful.

At an intermediate level, Steve was repeatedly struck by how art students were comfortable working with approximations and caricatures of real data. Sketches, cartoons, and inexact representations usually conveyed relationships just as well, and often even better, than perfectly scaled and anatomically accurate objects. The freedom to exaggerate or ignore certain elements made him see his data differently. He really appreciated the class' attention to color, line and perceptual cues—and yet they varied their approaches far more than expected.

Steve added that

> "Seeing our project in the Yurt was invigorating. 2-D sketches and concept drawings that looked promising did not always pan out in 3-D. Yet other examples that I thought might go nowhere shocked me in the Yurt. Sometimes the impact was simply a novel perspective (standing above a lifesize trackway or plunging down and into the depths of a magnified track). The stereo and textures often combined to offer visual overlap that I've never had before, seeing through multiple layers without losing spatial context."

CONCLUSION

In addition to the experts, student commentary both implicitly, and explicitly via Tim, overwhelmingly suggested that "the class gave me a lot to think about". Upon watching a trajectory of confidence and effectiveness emerge by the predominantly college age participants, and experiencing their growth in being able to work together face-to-face within a technology (instead of isolated with personal gadgets as buffers), I find myself wishing that their careers include many art-tech-science collaborations where they can innovate in new ways with their unique trajectory as an age group. Or, if not, that they will continue to weave technical information from visualization research with multiple technical fields of knowledge to find innovative directions for their work.

REFERENCES

[1] D. F Keefe, D. B. Karelitz, E. L. Vote, D. H. Laidlaw, *Artistic collaboration in designing VR visualizations*, Computer Graphics and Applications, IEEE 25 (2), 18-23

[2] R. M. Kirby, H. Marmanis, and D. H. Laidlaw. *"Visualizing Multivalued Data from 2D Incompressible Flows Using Concepts from Painting." In VIS '99*: Proceedings of the Conference on Visualization '99, 20, pp. 333-340, Los Alamitos, CA: IEEE Computer Society Press, 1999.

[3] D. H. Laidlaw. Loose, artistic "textures" for visualization. *IEEE Computer Graphics and Applications*, 21(2):6–9, March/April 2001

[4] P. S. Stevens, *Patterns in Nature*, Little, Brown and Company; 1st edition (1974).

[5] N. C. Hale, *Abstraction in Art and Nature*, Chapter 3, pp. 57-85, Dover Publications; Reprint edition (1993).

[6] D. H. Laidlaw, R. M. Kirby, C. D. Jackson, J. S. Davidson, T. S. Miller, M. da Silva, W. H. Warren, and M J. Tarr, Comparing 2D Vector Field Visualization Methods: A User Study, *IEEE Transactions on Visualization and Computer Graphics*, Vol 11, No 1, Jan-Feb 2005.

[7] B Martinez and J. Block, *Perception, Design, and Practice*, Chapters 1-5, Prentice Hall College Div (1985).

[8] J.J. Gibson, *The Ecological Approach to Visual Perception*. Chapter 8, Boston: Houghton Mifflin, (1986).

[9] P. L. Falkingham, S. M. Gatesy (2014), The birth of a dinosaur footprint: Subsurface 3D motion reconstruction and discrete element simulation reveal track ontogeny, *Proceedings of the National Academy of Sciences*

[10] *S. M. Gatesy, K. M. Middleton*, F, A. Jenkins Jr & N, H. Shubin, Three-dimensional preservation of foot movements in Triassic theropod dinosaurs, *Macmillan Magazines Ltd letters to nature, Nature*, Vol 399, May 13, 1999

[11] D. F. Keefe, D. Acevedo, T. Moscovich, D. H. Laidlaw, and J. LaViola. Cavepainting: A fully immersive 3D artistic medium and interactive experience. In *Proceedings of ACM Symposium on Interactive 3D Graphics 2001*, pp 85–93, March 2001

The Thinking Behind Startup@Brown

BY VALENTIN PEREZ '18

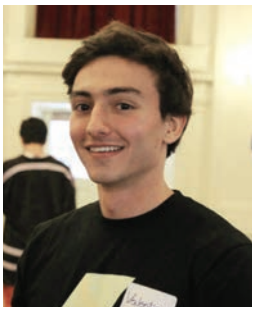

When I got to Brown as a freshman, I knew I wanted to dive right into the startup scene. Since I can remember, I've always loved creating things and working on projects with teams. I had come to the first Hack@Brown as a senior in high school and it was a no-brainer to join the organizing team once I got to campus. I was also part of the first iteration of the Innovation Dojo, a semester-long workshop series for underclassmen by the Brown Entrepreneurship Program, and I immediately joined Brown EP as well.

Most of these initiatives were new, and I felt the entrepreneurship scene at Brown was catching sparks. I went to the fall career fair and there were a lot of big companies, but I felt it wasn't as welcoming for freshmen, and I also didn't see a clear way to apply to startups. I really wanted to feel the rush of working at a VC-backed startup, so through the hackNY Fellowship, I was fortunate to end up interning at a startup in NYC for the summer.

I felt Startup@Brown was a great way to expose more students to this thrilling world of startups and kick off the year by trying to start a fire.

Startup@Brown relies on a simple premise: connect students and startups. It's a single idea, but the intention is threefold. First, demystify the process of founding a startup. Second, enable a great recruiting starting-point for both startups and students. And third, but not least, inspire, grow, and bring together the startup community on College Hill. We structured the conference around these goals.

Startup@Brown was a weekend-long conference at Brown that brought together students and startups.

DEMYSTIFY THE PROCESS OF FOUNDING A STARTUP

There's been an explosion of interest on the startup world. We only had spots for 250 students, but when we opened registration for a bit over 2 weeks, more than 500 students from Brown and RISD responded. In many of the responses and conversations, students said they had startup ideas but that they would like to learn more about the process of bringing them to reality.

We didn't want Startup@Brown to be another 300-person lecture (apart from the keynote), so we took a more dynamic yet intimate approach, mirroring the idea of closeness and face-to-face contact in startups. We had workshops in which 20-30 students learned from startup leaders, from founders and CTOs of startups such as Teespring, Casper, and Figma, to venture capitalists such as Y Combinator, Techstars, and Andreessen Horowitz. In total, we had 17 workshops ranging from "Building Your First Physical Product" to "Funding 101" to "User Experience". Kevin Hale, partner at Y Combinator, one of the top startup incubators in the world, also held office hours for student startups.

ENABLE A GREAT RECRUITING STARTING POINT

It's hard for startups to access top talent the way big companies do. Startups have way less time and money to spend on recruiting than corporations. Big companies have full-time recruiters and startups have founders focusing on their product.

And for students, the application process to work at small startups is shrouded in mystery, if not nonexistent. In contrast to larger organizations, startups have no information about intern programs on their websites and only a few have visible job postings. Career fairs are usually filled with large companies and little to no startups.

Therefore, right after the opening keynote we had the Startup Fair, a startup-only career fair. Students had a chance to walk around, meet with startups and learn about their work, challenges, and career opportunities. Similar to a hackathon, students didn't have to bring their resumes to each table—a student could choose to have her resume shared with all the companies before the event.

Above Left // Eveline Buchatskiy, Director of Techstars Boston, gave the opening keynote on the startup journey —a playbook on creating a startup

Above Right // Allan Tear, co-founder of Betaspring, giving a workshop of the six most common entrepreneurial mistakes and how to avoid them

Above Left // The Startup Fair, with Jerry Talton, CEO of Apropose, a data+design company, talking to students

Above Right // A small-group conversation with Sam Pierce Lolla, Farmlogs Lead Designer

Apart from the Startup Fair and workshops on Saturday, on Sunday there were small-group conversations; startup representatives (engineers, designers, founders) and six to ten students sat in a circle and the students could ask questions and hear stories from the leader for 30 minutes, getting a better sense of the different company cultures.

INSPIRE, GROW, AND BRING TOGETHER THE STARTUP COMMUNITY ON COLLEGE HILL

The reason why some people go to Silicon Valley to start companies is because of the community. College Hill and Providence have an enormous potential for creative endeavours; it's a very fertile ground that has cultivated many great entrepreneurs, startups, and social change, but we believe there's always room for growth.

Many students said one of their favorite moments was meeting and interacting with like-minded people. Bringing close to 300 students and startup leaders together for a weekend in September made for a great start of the year.

Naturally, many of the founders, engineers, and designers were alums, and they brought plentiful experience and enthusiasm back to campus. The help of alums was one of the main catalysts for Startup@Brown: we were thrilled to see over 20 alums from the different startups back on campus.

We sent a survey after the event where we got some feedback from attendees. 79% of the students are now more inclined to start their own startup, including six students saying they were starting one after Startup@Brown, and 97% would consider working with at least one of the startups that came. Feedback from the startups included that it was their "favorite career-fair" and that "the focus on startups was a huge win because every student was very engaging".

Out of another survey we sent in January, 59% of respondents interviewed with at least one Startup@Brown startup, and 37% got offers for a summer internship, full-time job, or non-summer internship.

The close interactions created new friendships and now a stronger startup community vibrates through College Hill. There has been a snowball effect: out of conversations with Y Combinator, a new Brown CS startup class, *csciStartup*, is being offered in Spring 2016, taught by Prof. John Jannotti, supported by Y Combinator, and with guest lectures from some of Startup@Brown's speakers.

Startup@Brown was hosted by Brown University's Department of Computer Science and organized in collaboration with the student-clubs Hack@Brown and Brown Entrepreneurship Program. We owe special thanks to alums: Startup@Brown wouldn't have been possible without you.

If you have suggestions for next year, want to support, or just want to say hi, email us at hello@startupatbrown.org. We learned from the first edition and hope to make next year better. Stay updated for the next iteration at startupatbrown.org.

Left // Brown and RISD students together for the opening keynote, with Valentin Perez ‘18 on stage

Below Center // Kevin Hale, partner at Y Combinator, giving the closing keynote on Sunday

Below // Part of the Startup@Brown organizing team: Katie Hsia, Sharon Lo, Valentin Perez, Athyuttam Eleti, and Krishna Chaitanyu Aluru

Faculty Notes

Jennifer Wortman Vaughan, Amy Greenwald, and Hanna Wallach, who delivered the opening remarks at Women in Machine Learning (WiML) Workshop 2015

AMY GREENWALD

SORIN ISTRAIL

In 2015, Sorin gave two keynote lectures to the 26th International Conference on Combinatorial Pattern Matching, Ischia Island, Italy and the 8th Congress of Romanian Mathematicians, University of Iasi, Romania, and an invited lecture at the University of Iasi, Summer School in Computer Science, ECODAM.

Sorin has been advising three very strong undergraduates doing honor theses on track to graduate in May 2016: Samuel Crisanto working on long-range haplotype phasing, Youn Kim working on spectral graph theory methods with applications to computational genomics, and Daniel Seidman working on haplotype tracts inference on unphased haplotype data. Doug McErlean, now at Google, continued working on the research project of his honor thesis done under Sorin's supervision. He obtained a breakthrough result about computational methods for maximum likelihood haplotype phasing by speeding up computation using symmetries of the likelihood function, which he presented at the "Probabilistic Modeling in Genomics" Conference, in October 2015 at Cold Spring Harbor Laboratory. As John von Neumann commented, in order to have a new theory you need first new theorems. The paper "CorePhase: Reducing maximum likelihood phasing problems of 2 ^100 variables or more to EM-practical sizes without loss of optimality via graph-theoretic symmetries of the likelihood function" presents such a theorem for Expectation-Maximization meta-algorithms for big data. Sorin also continued work with Hammurabi Mendes on extending methods from mathematical economics and voting theory to the protein folding energy function inference problem. A collaboration with statistician Wendy Wong from INOVA Translational Medical Institute in Northern Virginia was established focusing on extending the HapCompass algorithms (Derek Aguiar & S.I.) to de novo mutations inference and with applications to parents-of-origin alleles inference and genetic mosaicism.

In teamwork with two university presidents, the department of computer science and colleagues from physics, applied mathematics and economics departments, over almost two years, Sorin designed and chaired the "Brown University 250th Anniversary Symposium: The Next 250 Years," held in May 2015. The Symposium was co-organized with the Office of the President, the Department of Computer Science and colleagues professors Leon Cooper (Physics), Stu Geman (Applied Mathematics) and Roberto Serrano (Economics). The Symposium featured fourteen John von Neumann lectures on computer science, economics, physics and neuroscience, unified by John von Neumann's vision of "Computation as a Scientific Lens." Each lecture was followed by a tough Q&A session with the speaker, a so called Sweatbox session, concept inspired by professor Eric Davidson of California Institute of Technology. The distinguished lecturers included three Nobel Laureates and one Turing Award winner: Ken Arrow (Stanford), Leon Cooper (Brown), Frank Wiczek (MIT), Leslie Valiant (Harvard), Nima Arkani-Hamed (Institute for Advanced Studies), David Berson (Brown), Patricia Churchland (USCD), Vincent Crawford (Oxford), Freeman Dyson (Institute for Advanced Studies), Michael Jordan (Berkeley), Tom Leighton (MIT), Christos Papadimitriou (Berkeley), Mark Satterthwaite (Northwestern), and Susanne Schennach (Brown) were each questioned in the Davidson–Sweatbox style. Professor Eric Davidson delivered the John von Neumann Lecture on Biology at Brown University's 2010 John von Neumann Symposium.

Eric Davidson (Caltech) and Sorin, Caltech 2005

Franco Preparata, Alberto Apostolico, and Sorin, Brown University 2014

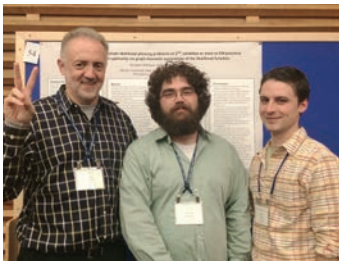

Sorin and his former students, Doug McErlean (Google) and Derek Aguiar (Princeton), Cold Spring Harbor Laboratory, 2015

In the past year, Sorin was devastated by the loss of three individuals who had tremendous impact on his scientific life: Professor Eric Davidson of California Institute of Technology, his closest friend, most important scientific collaborator and mentor; Professor Alberto Apostolico of Georgia Institute of Technology, his close friend and collaborator; and Julie Nguyen Brown, the donor of Sorin's Brown University endowed chair professorship.

SHRIRAM KRISHNAMURTHI

Photos from Shriram's hikes in the Berner Oberland, Switzerland

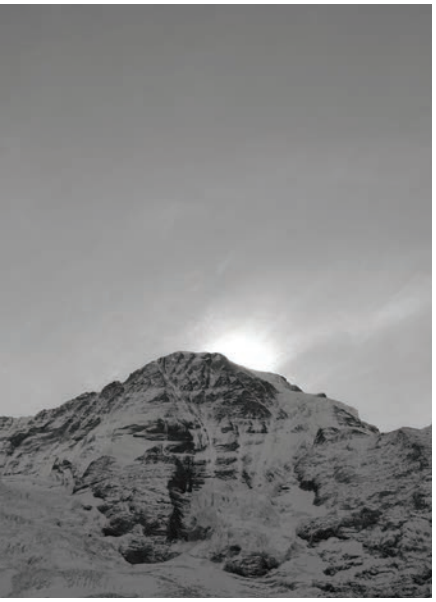

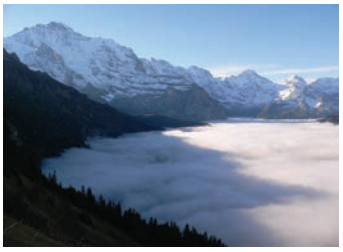

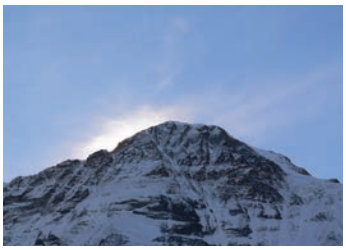

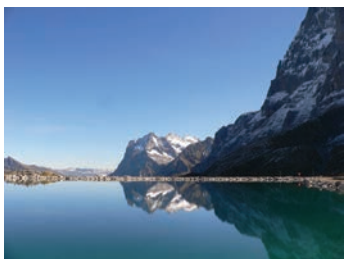

BARBARA MEIER

A highlight this year was bicycling in the Netherlands last summer. In addition to getting to know Andy van Dam's homeland, I was able to recreate the scene from one of my favorite animated films, *Father and Daughter* by Michael Dudok de Wit, in which a woman's life is portrayed in a series of bike rides along a dike. Indeed, the bicycle culture, infrastructure, and ease of getting everywhere by bike in the Netherlands is to be envied. David Laidlaw and I regularly commute the ten miles from Barrington to Providence by bicycle, and we are thrilled that the new Washington Bridge Linear Park finally opened after three years of bridge closure (to bikes). Next time you are in Providence, head down to India Point Park to check it out and don't forget to support cycling culture in your own community.

JOHN SAVAGE

John Savage continues to be active in cybersecurity policy and technology. He continues to be a Professorial Fellow at the EastWest Institute where he participates in the EWI Breakthrough Groups on Governing and Managing the Internet and Promoting Measures of Restraint in Cyber Armaments.

Since last spring he was invited to participate in many meetings on Internet governance. They include 1) the Russian International Information Security Research Consortium in Germany in April, 2) the IEEE Experts in Technology and Policy Forum (ETAP) in San Jose in May, 3) the EWI Cyberspace Cooperation Summit in New York in September, 4) the Munich Security Conference Roundtable in New York in September, 5) the India Conference on Cyber Security and Internet Governance in New Delhi in October, 6) a meeting run by the Department of State and the National Intelligence Council entitled "Cyber Deterrence: What is Really Required?" in DC in November, and 7) the Chinese 2nd World Internet Conference (Wuzhen Summit) in Wuzhen in December.

In the spring John taught his writing course CSCI 1800 Cybersecurity and International Relations to 77 students and co-taught CSCI 1951E Computer Systems Security: Principles and Practice with help from six colleagues to 62 students in which he gave more than half the lectures.

During the last year John has served the Faculty as Secretary of the Faculty and Secretary of the Faculty Forum.

Jared Bielby, John Savage, and Rod Beckstrom at the World Internet Conference (Wuzhen Summit)

Department Awards and Honors

Each item below is just the headline of an entire story for you to read, complete with photos and links and additional resources. For full versions, please visit **www.cs.brown.edu**, where they're featured at either CS News or CS Blog:

Bootstrap Plays A Role In White House CS Education Initiative

Cybersecurity: Brown Undergraduates Win It All

Brown CS Students Win Three Awards At HackMIT

Multiple Brown CS Collaborators Win Two Best Paper Award Runner Up Honors At HCOMP 2015

Two Brown CS Teams Win CyberSEED Prizes

Brown CS Earns Landmark Recognition In Computer Operating Systems At SOSP 2015

Four Brown CS Students Recognized As 2015 Google Scholars

Michael Chang Wins KPCB Fellowship

Crotty, Galakatos, Zgraggen, Binnig, And Kraska Win Best Demo At VLDB 2015

Esha Ghosh, Olya Ohrimenko, And Roberto Tamassia Win The ACNS 2015 Best Student Paper Award

2nd Place Award For Software At HackPrinceton Goes To Aaron Gokaslan '18 And Laura Shea '18

Amy Greenwald Delivers Opening Remarks At Women In Machine Learning Workshop's Tenth Anniversary

Maurice Herlihy Joins Dijkstra, Needham, And Others In The SIGOPS Hall Of Fame

Kraska, Binnig, Cetintemel, And Upfal Contribute To Northeast Big Data Innovation Hub

Tim Kraska Wins NSF CAREER And AFOSR Young Investigator Awards

Tim Kraska, Andy van Dam, And Carsten Binnig Win A Google Faculty Research Award

Michael Littman Named To DARPA ISAT Study Group

Michael Littman Wins IFAAMAS Influential Paper Award

Anna Lysyanskaya Reelected An International Association For Cryptologic Research Director

Mace, Roelke, And Fonseca Win A Best Paper Award At SOSP 2015

Jonathan Mace Receives A Facebook Graduate Fellowship

Christian Mathiesen And Teammates Take First Place At LinkedIn's Intern Hackday

Research By Undergraduate Sarah Sachs Gets Attention From Wired, The Today Show, And Others

John Savage Is Awarded A New Patent And Travels To The Munich Security Conference

John Savage's Participation At The World Internet Conference's Wuzhen Summit Airs On Chinese Television

John Savage Continues His Participation In International Cyber Conferences

Erik Sudderth And Collaborators Advance Seismic Monitoring And Nuclear Non-Proliferation, Earning A Top Prize In Bayesian Analysis

Stefanie Tellex Named One Of Wired UK's Women Who Changed Science In 2015

Stefanie Tellex And John Oberlin's Award-Winning Video Earns Brown CS A New Baxter Robot

Stefanie Tellex Wins A DARPA Young Faculty Award For Research In Human-Robot Communication

Eli Upfal Wins A Brown Institute For Brain Science Innovation Award

"A New Form Of Computer-Based 'Manuscript', Called A Hypertext": Time Magazine Cites AVD's Pioneering 1974 Work

Stan Zdonik And Co-Authors Win VLDB's 10-Year Best Paper Award

The 2015 Paris C. Kanellakis Memorial Lecture

On December 16, 2015, Brown CS and Assistant Professor Paul Valiant hosted the Fifteenth Annual Paris C. Kanellakis Memorial Lecture. Shafi Goldwasser of MIT and the Weizmann Institute delivered a talk ("The Cryptographic Lens") on how cryptography has changed the way we think of proofs, reductions, randomness, secrets, and information, as well as recent developments in cryptography that may allow the next successful shift from local to global computation. A reception followed the well-attended lecture.

Above // Audience for "The Cryptographic Lens"

Left // Kanellakis fellow Ahmad Mahmoody, Shafi Goldwasser, and Paul Valiant

Bottom Left // Bruno Harris, Professor of Mathematics (Emeritus)

Bottom Right // Brown CS student Gabriel Bankman-Fried at the reception following the lecture

Recent PhDs

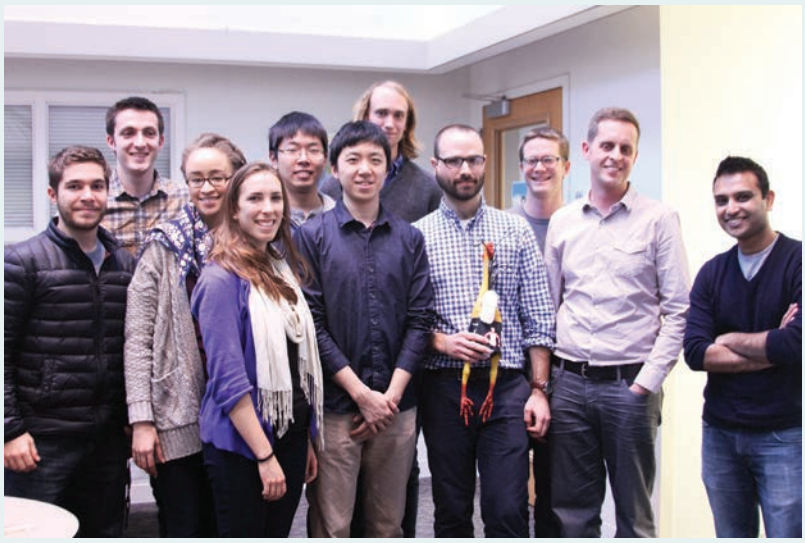

THIS PAGE CLOCKWISE FROM TOP LEFT

Jason Pacheco

Irina Calciu

Mark Buller

Hammurabi Mendes

Layla Oesper

THIS PAGE CLOCKWISE FROM TOP LEFT

Soumya Ghosh

Justin deBrabant

Steven Gomez

Rebecca Mason

PIVOT TRACING: Dynamic Causal Monitoring For Distributed Systems

BY JONATHAN MACE, RYAN ROELKE, AND RODRIGO FONSECA

Pivot Tracing is a monitoring framework for distributed systems that can seamlessly correlate statistics across applications, components, and machines at runtime, without needing to change or redeploy system code. Users can define and install monitoring queries on-the-fly, to collect arbitrary statistics from one point in the system while being able to select, filter, and group by events meaningful at other points in the system. Pivot Tracing does not correlate cross-component events using expensive global aggregations, nor does it perform offline analysis. Instead, Pivot Tracing directly correlates events as they happen, by piggybacking metadata alongside requests as they execute—even across component and machine boundaries. This gives Pivot Tracing a very low runtime overhead—less than 1% for many cross-component monitoring queries.

MONITORING AND TROUBLESHOOTING DISTRIBUTED SYSTEMS

Problems in distributed systems are many and varied: component failures due to hardware errors, software bugs, and misconfiguration; unexpected overload behavior due to hot spots and aggressive tenants; or simply unrealistic user expectations. Due to designs such as fault-tolerance and load balancing, the root cause of an issue may not be immediately apparent from its symptoms. However, while troubleshooting distributed systems is inherently challenging, many of the monitoring and diagnosis tools used today share two fundamental limitations that further exacerbate the challenge.

One size does not fit all

First, many tools only record information that is selected a priori at development or deployment time. Even though there has been great progress in using machine learning techniques and static analysis to improve the quality of logs, they still carry an inherent trade-off between recall and overhead. The choice of what to record must be made a priori, so inevitably the information needed to diagnose an issue might not be reported by the system. Even if a relevant event is captured in a log message, it can still contain too little information; similarly, performance counters may be too coarse grained, or lack the desired filters or groupings.

On the other hand, if a system does expose information relevant to a problem, it is often buried under a mountain of other irrelevant information, presenting a "needle in a haystack" problem to users. Any time a user or developer patches a system to add more instrumentation, they contribute to this information overload. They also potentially add performance overheads for any monitoring that is enabled by default. Unsurprisingly, developers are resistant to adding additional metrics or groupings, as can be observed in a plethora of unresolved and rejected issues on Apache's issue trackers.

Crossing Boundaries

Second, many tools record information in a component- or machine-centric way, making it difficult to correlate events across these boundaries. Since today's datacenters typically host a wide variety of inter-operating components and systems, the root cause and symptoms of an issue often appear in different processes, machines, and application tiers. A user of one application may need to relate information from some other dependent application in order to diagnose problems that span multiple systems. To do this manually is cumbersome and in many cases impossible, because it depends on sufficient execution context having been propagated across software component and machine boundaries.

DYNAMIC INSTRUMENTATION AND CAUSAL TRACING

Pivot Tracing overcomes these challenges by combining two key techniques: dynamic instrumentation and causal tracing. Dynamic instrumentation systems, such as DTrace,[1] Fay,[2] and SystemTap,[6] let users defer until runtime their selection of information reported by the system. They allow almost arbitrary instrumentation to be added dynamically at runtime as needed, and have proven extremely useful in diagnosing complex and unanticipated system problems. Pivot Tracing also uses dynamic instrumentation, enabling users to specify new monitoring queries at runtime. Pivot Tracing queries are dynamically installed without the need to change or redeploy code.

Dynamic instrumentation alone does not address the challenge of correlating events from multiple components. To address this challenge, Pivot Tracing adapts techniques presented in the causal tracing literature by systems such as X-Trace[3] and Dapper.[7] These systems maintain a notion of context that follows an execution through events, queues, thread pools, files, caches, and messages between distributed system components. Likewise, Pivot Tracing propagates a tracing context alongside requests. Unlike end-to-end tracing, Pivot Tracing does not record or reconstruct traces of executions for offline analysis. Instead, its tracing context is a means for propagating a small amount of state directly along the execution path of requests, including when they cross component and machine boundaries.

PIVOT TRACING

Pivot Tracing exposes these two features by modelling system events as the tuples of a streaming, distributed data set. Users can write relational queries about system events using Pivot Tracing's LINQ-like query language. Pivot Tracing compiles queries into instrumentation code and dynamically installs the code at the sources of events specified in the query. Each time one of the events occurs, the instrumentation code is also invoked.

Happened-Before Join

In order to reason about causality between events, Pivot Tracing introduces a new relational operator, the "happened-before join," ⋈, for joining tuples based on Lamport's happened-before relation.[4] For events a and b occurring anywhere in the system, we say that a happened before b and write a→b if the occurrence of event a causally preceded the occurrence of event b and they occurred as part of the execution of the same request. Using the happened-before join, users can write queries that group and filter events based on properties of events that causally precede them in an execution. Pivot Tracing evaluates happened-before join by putting partial query state into the tracing contexts propagated alongside requests. This is an efficient way to evaluate happened-before join, because it explicitly follows the happened- before relation. It drastically mitigates the overhead and scalability issues that would otherwise be required for correlating events globally.

Figure 1

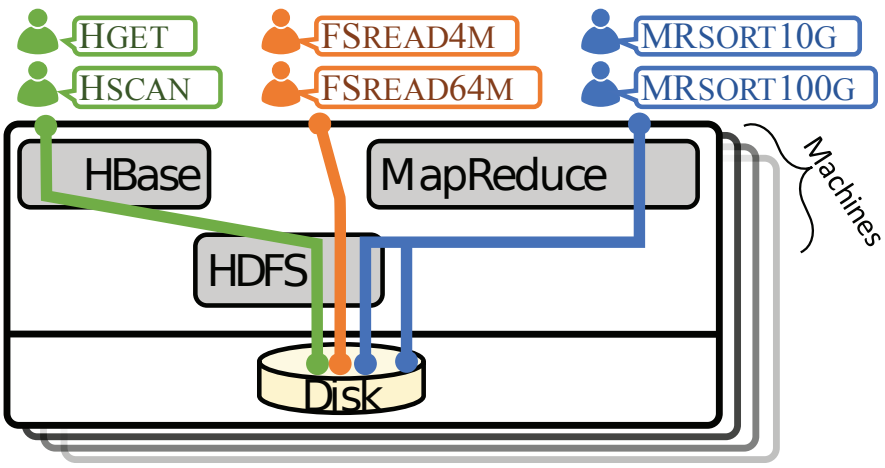

Figure 1 // Six client workloads access the disks on eight cluster machines indirectly via HBase, a distributed database; HDFS, a distributed file system; and MapReduce, a data processing framework.

HGet	10kB row lookups in a large HBase table
HScan	4MB table scans of a large HBase table
FSread4M	Random closed-loop 4MB HDFS reads
FSread64M	Random closed-loop 64MB HDFS reads
MRsort10G	MapReduce sort job on 10GB of input data
MRsort100G	MapReduce sort job on 100GB of input data

PIVOT TRACING *continued*

PIVOT TRACING IN ACTION

To motivate Pivot Tracing's design and implementation, we present a brief example of Pivot Tracing with a monitoring task in the Hadoop stack. Suppose we are managing a cluster of eight machines, and we want to know how disk bandwidth is being used across the cluster. On these machines, we are simultaneously running several clients with workloads in HBase, MapReduce, and HDFS. It suffices to know that HBase is a distributed database that accesses data through HDFS, a distributed file system. MapReduce, in addition to accessing data through HDFS, also accesses the disk directly to perform external sorts and to shuffle data between tasks. Figure 1 depicts this scenario.

By default, our distributed file system HDFS already tracks some disk consumption metrics, including disk read throughput aggregated on each of its DataNodes. To reproduce this metric with Pivot Tracing, we can define a tracepoint for the method incrBytesRead(int delta) in the DataNodeMetrics class in HDFS. A tracepoint is a location in the application source code where instrumentation can run. We then run the following query in Pivot Tracing's LINQ-like query language:

```
Q1:  From incr In DataNodeMetrics.incrBytesRead
     GroupBy incr.host
     Select incr.host, SUM(incr.delta)
```

This query causes each machine to aggregate the delta argument each time incrBytesRead is invoked, grouping by the host name. Each machine reports its local aggregate every second, from which we produce the time series in figure 2a.

Things get more interesting if we wish to measure the HDFS usage of each of our client applications. HDFS only has visibility of its direct clients, and thus it only has an aggregate view of all HBase and all MapReduce clients. At best, applications must estimate throughput client side. With Pivot Tracing, we define tracepoints for the client protocols of HDFS (DataTransferProtocol), HBase (ClientService), and MapReduce (ApplicationClientProtocol), and use the name of the client process as the group-by key for the query. Figure 2b shows the global HDFS read throughput of each client application, produced by the following query:

```
Q2:  From incr In DataNodeMetrics.incrBytesRead
     Join cl In First(ClientProtocols) On cl -> incr
     GroupBy cl.procName
     Select cl.procName, SUM(incr.delta)
```

The -> symbol indicates a happened-before join. Pivot Tracing's implementation will record the process name the first time the request passes through any client protocol method and propagate it along the execution. Then, whenever the execution reaches incrBytesRead on a DataNode, Pivot Tracing will emit the bytes read or written, grouped by the recorded name. This query exposes information about client disk throughput that cannot currently be exposed by HDFS.

Figure 2a

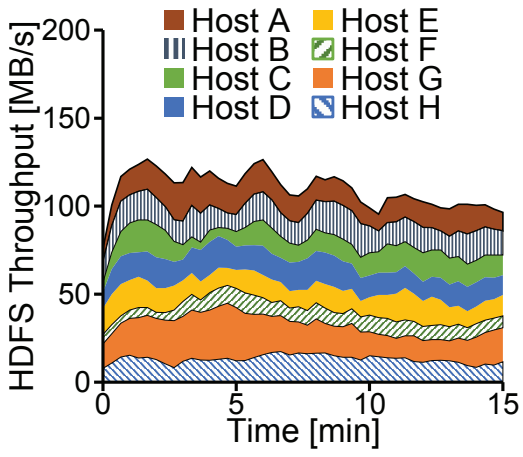

Figure 2b

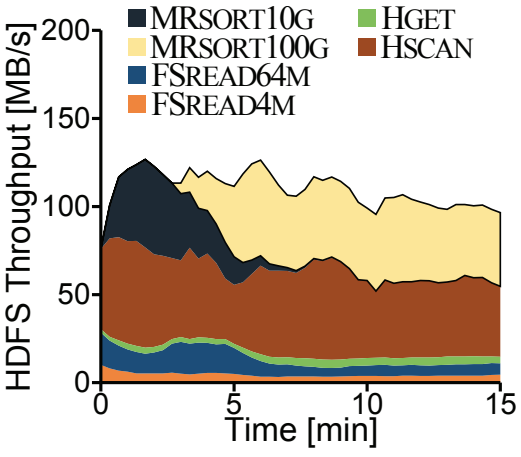

Figure 2 // In this example, Pivot Tracing dynamically instruments HDFS to expose read throughput grouped by client identifiers from other applications.

(a) HDFS DataNode throughput per machine

(b) HDFS DataNode throughput per high-level client application.

PIVOT TRACING DESIGN AND IMPLEMENTATION

We opted to implement our Pivot Tracing prototype in Java in order to easily instrument the aforementioned open-source distributed systems. However, the components of Pivot Tracing generalize and are not restricted to Java—a query can even span multiple systems written in different programming languages. Full support for Pivot Tracing in a system requires two basic mechanisms: dynamic code injection and causal metadata propagation. For full details of Pivot Tracing's design and implementation, we refer the reader to the full paper[5] and project website (http://pivottracing.io)

Figure 3 presents a high-level overview of how Pivot Tracing enables queries such as Q2. We will refer to the numbers in the figure (e.g., ①) in our description.

Writing Queries

Queries in Pivot Tracing refer to variables exposed by one or more tracepoints (①)—places in the system where Pivot Tracing can insert instrumentation. Tracepoints export named variables that can be accessed by instrumentation. However, the definitions of tracepoints are not part of the system code, but rather instructions on where and how Pivot Tracing can add instrumentation. Tracepoints in Pivot Tracing are similar to pointcuts from aspect-oriented programming, and can refer to arbitrary interface/method signature combinations. Pivot Tracing's LINQ-like query language supports several typical operations including projection, selection, grouping, aggregation, and happened-before join.

Compiling Queries

Users submit queries to the Pivot Tracing front-end (②), which is responsible for optimizing queries using some simple static rewriting rules, pushing projection, selection, and aggregation as close as possible to the source tracepoints. The front-end then compiles queries into advice, an intermediate representation of the system-level instrumentation needed to evaluate the query. Advice specifies the operations to perform at each tracepoint used in a query.

Installing Queries

The Pivot Tracing front-end distributes advice to local Pivot Tracing agents running in each process (③). Pivot Tracing agents are responsible for

Figure 3

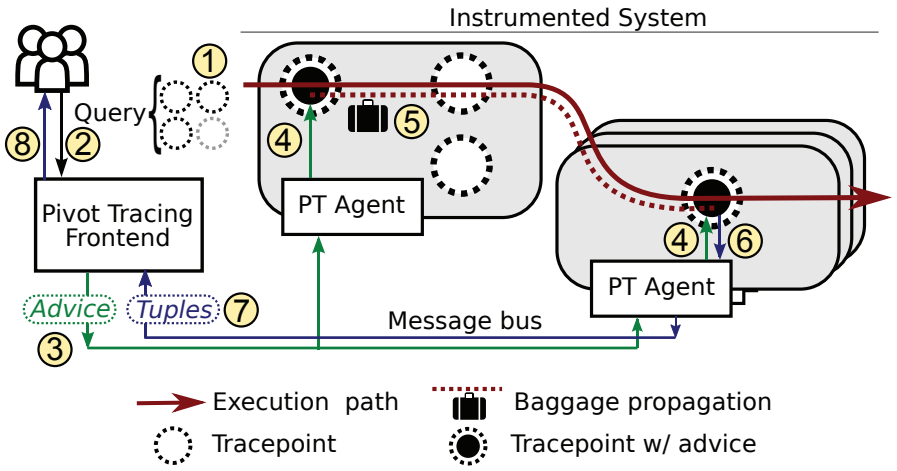

Figure 3 // Pivot Tracing overview

dynamically instrumenting the running system so that advice is invoked at tracepoints. The agents weave advice into tracepoints ((④)) by: 1) generating code that implements the advice operations; 2) configuring the tracepoint to execute that code and pass its exported variables; 3) activating the necessary tracepoint at all locations in the system. Later, requests executing in the system will invoke the installed advice every time their execution reaches the tracepoint.

Crossing Boundaries

In order to implement the happened-before join, advice invoked at one tracepoint needs to make information available to advice invoked at other tracepoints later in a request's execution. For example, in Q2, advice at the ClientProtocols tracepoint needs to make its procName available to later advice invoked at the DataNodeMetrics tracepoint. This is done through Pivot Tracing's baggage abstraction, which uses causal metadata propagation ((⑤)). Baggage is a per-request container for tuples that is propagated alongside a request as it traverses thread, application, and machine boundaries. At any point in time, advice can put tuples in the baggage of the current request, and retrieve tuples that were previously placed in the baggage by other advice.

Evaluating Queries

Advice uses a small instruction set to evaluate queries and maps directly to the code that local Pivot Tracing agents generate. Advice operations are as follows: advice can create a tuple from tracepoint-exported variables (Observe); filter tuples by a predicate (Filter); and output tuples

PIVOT TRACING *continued*

for global aggregation (Emit). Advice can put tuples in the baggage (Pack) and retrieve tuples from the baggage (Unpack). Unpacked tuples are joined to the observed tuples (i.e., if to is observed and tu1 and tu2 are unpacked, then the resulting tuples are totu1 and totu2). Both Pack and Emit can group tuples based on matching fields and perform simple aggregations such as SUM and COUNT.

Query Results

Advice can emit tuples as output of a query using the Emit instruction (⑥). Pivot Tracing first aggregates emitted tuples locally within each process, then reports results globally at a regular interval, e.g., once per second (⑦). The Pivot Tracing front-end collects and forwards query results to the user (⑧). Process-level aggregation substantially reduces traffic for emitted tuples; Q2 is reduced from approximately 600 tuples per second to 6 tuples per second from host.

Pivot Tracing Example

Recall query Q2 from our earlier Hadoop example:

```
Q2:  From incr In DataNodeMetrics.incrBytesRead
     Join cl In First(ClientProtocols) On cl -> incr
     GroupBy cl.procName
     Select cl.procName, SUM(incr.delta)
```

Q2 compiles to two advice specifications, A1 and A2, to be invoked at the ClientProtocols and DataNodeMetrics tracepoints respectively:

```
A1:  OBSERVE procName
     PACK procName

A2:  UNPACK procName
     OBSERVE delta
     EMIT procName, SUM(delta)
```

When a request invokes any of the ClientProtocols methods, the instrumented code will invoke advice A1. The advice will observe the value of the procName variable and pack a tuple into the request's baggage, e.g. <procName="HGet">. The request will continue execution, carrying this tuple in its baggage. If the request subsequently invokes the DataNodeMetrics.incrBytesRead method, the instrumented code will invoke advice A2. The advice will unpack the previously packed procName and observe the local value of the delta variable, e.g. <delta=10>. The advice will then join the unpacked procName with the observed delta and emit the result as output, e.g. <procName="HGet", delta=10>. The output tuple will be aggregated with other tuples in the process's Pivot Tracing agent, and included in the next interval's query results.

Figure 4 gives a final demonstration of how Pivot Tracing can group metrics along arbitrary dimensions. It is generated by two queries similar to Q2 which instrument Java's FileInputStream and FileOutputStream, still joining with the client process name. We show the per-machine,

Figure 4

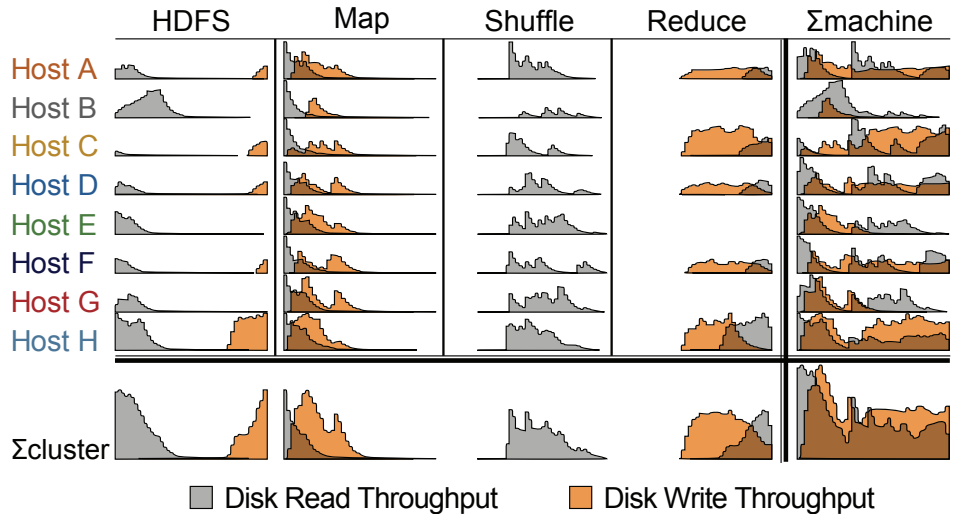

Figure 4 // Pivot table showing disk read and write sparklines for MRsort10g. Rows group by host machine; columns group by source process. Bottom row and right column show totals, and bottom-right corner shows grand total.

per-application disk read and write throughput of MRsort10g from the same experiment. This figure resembles a pivot table, where summing across rows yields per-machine totals, summing across columns yields per-system totals, and the bottom right corner shows the global totals. In this example, the client application presents a further dimension along which we could present statistics.

PIVOT TRACING SUMMARY

In this article we gave an overview of how Pivot Tracing can evaluate cross-component monitoring queries dynamically at runtime, using a combination of dynamic instrumentation and causal tracing. For full details of Pivot Tracing's design and implementation, we refer the reader to the full paper[5] and project website (http://pivottracing.io) . In our full evaluation, we present several case studies where we used Pivot Tracing to successfully diagnose root causes, including real-world issues we encountered in our cluster. We also evaluate the overheads imposed by Pivot Tracing, including the additional costs of invoking advice and the overheads of propagating tuples alongside requests at runtime. Of the examples presented in this article, Q2 only required the propagation of a single tuple per request, and imposed less than 1% overhead in terms of end-to-end latency on several application-level HDFS benchmarks.

Pivot Tracing is the first monitoring system to combine dynamic instrumentation with causal tracing. Its novel happened-before join operator fundamentally increases the expressive power of dynamic instrumentation and the applicability of causal tracing. Pivot Tracing enables cross-tier analysis between any inter- operating applications, and the overheads of evaluating the happened-before join are sufficiently low that we believe Pivot Tracing is suitable for production systems, both for high-level standing queries and for digging deeper when necessary. Ultimately, its power lies in the uniform and ubiquitous way in which it integrates monitoring of a heterogeneous distributed system.

REFERENCES

[1] Bryan Cantrill, Michael W Shapiro, and Adam H Leventhal. Dynamic Instrumentation of Production Systems. In USENIX Annual Technical Conference (ATC), 2004.

[2] Ulfar Erlingsson, Marcus Peinado, Simon Peter, Mihai Budiu, and Gloria Mainar-Ruiz. Fay: Extensible Distributed Tracing from Kernels to Clusters. ACM Transactions on Computer Systems (TOCS), 2012.

[3] Rodrigo Fonseca, George Porter, Randy H. Katz, Scott Shenker, and Ion Stoica. X-Trace: A Pervasive Network Tracing Framework. In 4th USENIX Symposium on Networked Systems Design and Implementation (NSDI), 2007.

[4] Leslie Lamport. Time, clocks, and the ordering of events in a distributed system. Communications of the ACM, 21(7):558–565, 1978.

[5] Jonathan Mace, Ryan Roelke, and Rodrigo Fonseca. Pivot Tracing: Dynamic Causal Monitoring for Distributed Systems. In 25th ACM Symposium on Operating Systems Principles (SOSP), 2015.

[6] V. Prasad, W. Cohen, F. C. Eigler, M. Hunt, J. Keniston, and B. Chen. Locating system problems using dynamic instrumentation. In Ottawa Linux Symposium (OLS), 2005.

[7] Benjamin H Sigelman, Luiz Andre Barroso, Mike Burrows, Pat Stephenson, Manoj Plakal, Donald Beaver, Saul Jaspan, and Chandan Shanbhag. Dapper, a large-scale distributed systems tracing infrastructure. Google research, 2010.

Celebrate With Andy:
50 YEARS OF CS AT BROWN

On Friday, May 22, 2015, the Brown CS family gathered in Pizzitola Sports Center to celebrate three golden anniversaries: the Undergraduate Teaching Assistant program, undergraduate participation in research, and Andy van Dam at Brown. The tributes and reminiscences were insightful, heartfelt, and often extremely funny. You can watch a full video of the formal program at http://tinyurl.com/cwavideo.

Right // Tony Fross, Katrin van Dam, Debbie van Dam, and Andy van Dam

Tom Banchoff and Tom Doeppner

Brown CS Alum Mary Fernandez shares a memory of AVD

Ed Lazowska and Ingrid Carlbom

Brown CS Alum Janete Perez

Welcoming guests to *Celebrate With Andy*

Carolyn Eckel, James Michener, and Rich Kogut

Around The Department

DJ Hoffman, Indy Prentice, Jake Ellis, and Aisha Ferrazares

Emma Herold and a friend

students at the 2015 Tech Fair

Nediyana Daskalova, Paige Selby, and Hannah Quay-de la Vallee

Left // Sorin Istrail and Tom Leighton, (Professor, MIT, and CEO, Akamai Technologies) at the *Brown University 250th Anniversary Symposium: The Next 250 Years*

Below // students at the 2015 Undergraduate Research Symposium

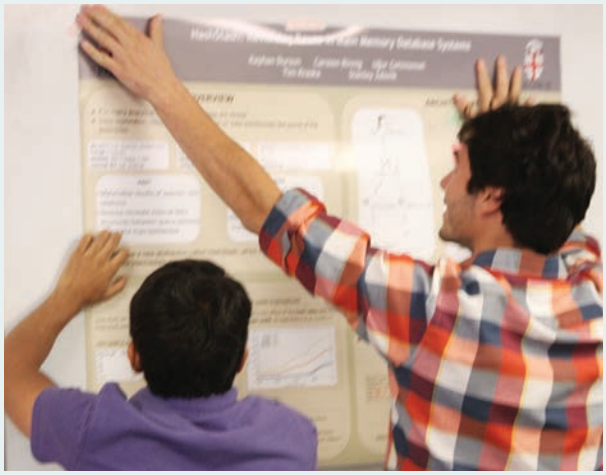

Joe Gibbs Politz, receiving the Sigma Xi award for outstanding graduate CS research

high school students watching a drone demonstration

Introducing Mosaic+

BY NIFEMI MADARIKAN

> We strive to foster a sense of community for marginalized and underrepresented racial minority (URM) students within the computer science department. We will welcome more racial and ethnic minority students to computer science at Brown through a community they can easily relate to and learn from. We will launch Mosaic+ as a platform from which to motivate administration in Brown University's computer science department to create a receptive and equitable space for minority students and faculty. We hope our efforts will draw attention to the gross lack of diversity in computing at Brown and beyond. Our ultimate objective is for the discourse around diversity in computing at Brown University to be acknowledged and duly acted upon.
>
> MOSAIC+ CONSTITUTION, FALL 2015

Founded in the spring semester of 2015, Mosaic+ is a student-led diversity initiative to create an inclusive space for racially and ethnically underrepresented minority (URM) students. We believe collaboration between students, faculty, staff, and partners in academia and industry is necessary to tackle the persistent prevalence of under-representation in technology today.

Here at Brown, Mosaic+ adopts a multifaceted approach to developing diversity and inclusion in our computer science department - one made possible by our versatile structure. Much of the work we do falls into two main programs which have received the encouragement, endorsement, and active help of Brown CS: *the Transition Program and the Mentorship Program*. The first is a summer pre-orientation program targeted at newly admitted, URM first-year students curious about computer science and technology at Brown University. Over the course of this pre-orientation program, students will be exposed to meaningful opportunities and exciting challenges as they engage with both technical and non-technical

activities and challenges. Our Transition Program launches this summer of 2016 in collaboration with Brown's Catalyst summer program. Launching this spring semester, our Mentorship Program provides focused guidance and attentive support to URM students through one-on-one mentorship, class-specific group study sessions, and large interactive workshops. Over the course of each semester, students will have the opportunity to gain a firmer grasp of concepts they learn in class as well as exposure to new and exciting concepts to expand their experience of computer science and technology.

Crucial to our success is the collaboration between students, faculty and staff; in particular we would like to appreciate the efforts of: Chadwicke Jenkins (former Brown CS faculty member), Department Chair Ugur Cetintemel, Associate Professor (Research) and Vice Chair Tom Doeppner, Laura Dobler, the Brown CS Diversity Committee, Nifemi Madarikan '17, Chelse-Amoy Steele '18, Kevin Anderson '18, Joanna Simwinga '18, Peter Clarke '18, Jonathan Powell '18, Kendal Cockrel '17, Jide Omekam '18, Asha Owens '16, Hassan Sufi '17, Austin Peete '18, Charles Njorge '18, Victoria Chavez '18, and all of the newly appointed Mosaic+ student mentors and facilitators.

Alum Update

DAVE JACKSON '98

As you can see, my daughter is already learning about networking!

Dave Jackson '98 has been busy: Three years ago, he began freelancing, and at the beginning of 2015, he joined forces with his old friend, Mik, to start Green Mars Consulting, a web-app+ software consulting business in the San Francisco Bay Area. He's also been busy with his now 13-month-old daughter, Elizara Jackson-Hall, and his almost-completed science fiction novel about hipster bicycle gangs in outer space.

After getting his integrated bachelors and Masters at Brown, he's worked for Oracle, Yahoo, NASA Ames Research Center, and a couple of start-ups. He finally decided he wanted the freedom and flexibility to work on projects he chose, and that led to Green Mars. In addition to consulting, Dave started a company with a former colleague to provide an automated product safety service.

Dave can be reached to reconnect at: dave@jacksonofalltrades.net or to talk business at:dej@greenmars.consulting.

He could never quite give up his old Brown CS username. :)

Department of Computer Science
Brown University
Box 1910
Providence, RI 02912
USA

Staying In Touch

- Send a blank email to alumsubscribe@cs.brown.edu to join the Brown CS Alum Email Group, which we'll use to occasionally share updates. It's only for alums, and members can't e-mail each other, so you don't need to worry about a high volume of messages.
- Visit our web site (www.cs.brown.edu), which we update daily with news and stories
- Visit (and please be sure to like) our Facebook page (www.tinyurl.com/browncsfb)
- Follow us on Twitter (www.twitter.com/browncsdept)

Industry Partners Program

The IPP provides a formal mechanism for interactions between companies and students in the CS Department. Member companies benefit from superior visibility in the Department, exclusive access to event/interview space in the CIT Building and assistance with recruiting events; students benefit from specific information about opportunities for summer internships and permanent employment.

The department wishes to thank our Industrial Partners for their support:

Premier
Adobe

Affiliates
Citizens Bank
Cloudera
Delphix
Facebook
Google
Jane Street
Microsoft
MongoDB
Onshape
Oracle
RetailMeNot, Inc
SIG
TripAdvisor
Two Sigma
World Quant
Yelp

Advertising
IXL Learning

Start-Up/Small Company
FiveStars
Joyent
pMD
Vision Systems
VoltDB

Individuals
Paul Edelman,
Edelman & Associates

Brown Computer science
Industry Partners Program

For more information about the Industry Partners Program, contact:

Lauren Clarke
Program Manager
Tel: (401) 863-7655
lkc@cs.brown.edu

Ugur Çetintemel
Professor and IPP Director
Tel: (401) 863-7600
ugur@cs.brown.edu

To learn more about the IPP visit:
http://www.cs.brown.edu/industry

Ping!

Let us know what's happening in your life! Submit your news by e-mail to conduit@cs.brown.edu or by mail to:

Conduit
Department of Computer Science
Brown University
Box 1910
Providence, RI 02912

Designed and produced by Brown Graphic Services in partnership with Computer Science

BROWN