



## Industry Affiliates Meeting October 20, 2015

Agenda

Break-out Topics & Talks

Talks viewed here: <http://tinyurl.com/CSEtalks2015>

Lunchtime Keynote Address

Poster and Demo Session – **Be sure to vote for the People’s Choice in EACH LAB!** Vote for your favorite poster in EACH LAB THAT YOU VISIT!

Posters in each lab may be viewed here: <http://tinyurl.com/CSEposters2015>

Building Maps

Internet Access

Recruiting Policy

Meeting materials:



Posters:





# COMPUTER SCIENCE & ENGINEERING

UNIVERSITY of WASHINGTON

## AGENDA

Research presentations, posters and demos will occur on the second day of our annual meeting - October 20<sup>th</sup>. The meeting will be held in the [Paul G. Allen Center for Computer Science & Engineering](#).

---

### Tuesday, October 20 - Research Overviews, Talks, Demonstrations, and Posters

This day will begin with a state-of-the-department overview that includes initiatives and research begun over the past year. The bulk of the day will consist research talks (in parallel tracks), a working lunch, time for industry feedback sessions, and will wrap up with an open house reception and lab tours with poster presentations and demonstrations. The day will allow our affiliates to meet many of our students at the talks and demonstration/poster sessions.

10:00 - 10:30am	Coffee and Registration	<a href="#">Paul G. Allen Center</a> , Gates Commons, Room 691
10:30 - 11:10am	Welcome and Overview <a href="#">Ed Lazowska</a> and <a href="#">Hank Levy</a> + various faculty on research areas	<a href="#">Paul G. Allen Center</a> , Gates Commons, Room 691
11:15am - 12:20pm	<a href="#">Break-out topics: Session I</a>	<a href="#">Paul G. Allen Center</a> , Rooms 305, 403, and 691
12:25 - 1:30pm	Lunch + <a href="#">Keynote Talk</a>	<a href="#">Paul G. Allen Center</a> , Microsoft Atrium
1:30 - 2:35pm	<a href="#">Break-out topics: Session II</a>	<a href="#">Paul G. Allen Center</a> , Rooms 305, 403, and 691
2:40 - 3:45pm	<a href="#">Break-out topics: Session III</a>	<a href="#">Paul G. Allen Center</a> , Rooms 305, 403, and 691
3:50 - 4:55pm	<a href="#">Break-out topics: Session IV</a>	<a href="#">Paul G. Allen Center</a> , Rooms 305, 403, and 691
5:00 - 7:00pm	<a href="#">Open House: Reception and lab tours (with posters and demos)</a>	<a href="#">Paul G. Allen Center</a> , Microsoft Atrium and various labs throughout the building
7:15 - 7:45pm	Program	<a href="#">Paul G. Allen Center</a> , Microsoft Atrium



### Break-out Topics and Talks

Tuesday, October 20, 2015

<a href="#">Session I</a> 11:15am - 12:20pm	<a href="#">Data Science at UW</a> CSE 305	<a href="#">SANE 1</a> CSE 403	<a href="#">Natural Language Processing</a> CSE 691
1:00 - 1:30pm	<a href="#">Keynote Talk</a> Atrium		
<a href="#">Session II</a> 1:30 - 2:35pm	<a href="#">Big Data Management and Analytics</a> CSE 305	<a href="#">SANE 2</a> CSE 403	<a href="#">Innovations in Mobile Systems</a> CSE 691
<a href="#">Session III</a> 2:40 - 3:45pm	<a href="#">Human Computer Interaction and Self-Tracking</a> CSE 305	<a href="#">Next Generation Programming Tools: Performance</a> CSE 403	<a href="#">New Sensing and Interaction Technologies</a> CSE 691
<a href="#">Session IV</a> 3:50 - 4:55pm	<a href="#">Computer Graphics and Vision at UW</a> CSE 305	<a href="#">Next Generation Programming Tools: Correctness</a> CSE 403	<a href="#">Health Sensing</a> CSE 691
5:00 - 7:00pm	<a href="#">RECEPTION AND LAB TOURS (WITH POSTERS AND DEMOS)</a> <i>various labs and locations around the building</i>		
7:15 - 7:45pm	<b>Program: Madrona Prize, People's Choice Awards</b> <i>Microsoft Atrium</i>		

---

## Session I

- [Data Science at UW \(CSE 305\)](#)

- **11:15-11:20: Introduction and Overview, [Bill Howe](#)**

I'll give an overview of the UW eScience Institute, and how our activities are helping to advance data science research and education across UW. I'll highlight a new Master's degree program launching Fall 2016, an "incubator" program to advance promising internal and external data science projects, and a new smart cities initiative in partnership with the City of Seattle and a campus-wide consortium called [Urban@UW](#).



# COMPUTER SCIENCE & ENGINEERING

UNIVERSITY of WASHINGTON

- **11:20-11:40: *Data Science for Social Good*, [Ariel Rokem](#)**

During the Summer of 2015, the University of Washington eScience Institute ran an interdisciplinary summer internship program focused on urban informatics, civic engagement, and data-intensive social science. Building on our own previous consulting and "incubation" programs, we brought together teams of students, data scientists, project leads and stakeholders from the University of Washington and local NGOs to design, develop, and deploy new solutions to high-impact problems in the Seattle Metro Area. The DSSG cohort included 16 students accepted from a pool of 144 applicants, working on 4 projects chosen out of 11 submitted proposals: two addressing transportation access for people with limited mobility, one identifying factors affecting whether homeless families find permanent housing, and one deriving new metrics of community well-being from social media data and other relevant data sources.
- **11:40-12:00: *Inferring Dynamics from Data*, [Valentina Staneva](#)**

In this talk I will describe several scientific projects in which we aim to infer dynamical processes from spatiotemporal data. We face similar problems when working with data across different applications: modeling insect behavior, understanding neuronal activity, etc, and I will discuss methodology to address those problems.
- **12:00-12:20: *How Python is Changing Science*, [Jake Vanderplas](#)**

Research across scientific fields is becoming increasingly dependent on computation, whether that computation involves large-scale simulation, analysis of large datasets, or a combination thereof. Over the past decade, we have seen an accelerating shift within many fields from older domain-specific tools like IDL and MatLab to Python and its accompanying scientific ecosystem. In this talk, I'll give a brief summary of why Python has become such a tool of choice in the scientific community, and discuss the foundational emphasis on open science and reproducible research that have accompanied uptake of this tool.
- ***Systems, Architecture, and Networking (SANE) 1 (CSE 403)***
  - **11:15-11:20: Introduction and Overview, [Mark Oskin](#)**
  - **11:20-11:35: *Accelerating Data-Intensive Applications with Latency-Tolerant Distributed Shared Memory*, [Jacob Nelson](#)**

Conventional wisdom suggests that making large-scale distributed computations fast requires minimizing the latency of individual operations in the computation. Grappa is a modern take on software distributed shared memory that takes the opposite view: it tolerates latency by exploiting application parallelism to achieve overall higher throughput. In this talk I'll discuss Grappa's core ideas, performance results, and some future directions.
  - **11:35-11:50: *Building Consistent Transactions with Inconsistent Replication*, [Irene Zhang](#)**



Application programmers increasingly prefer distributed storage systems with distributed transactions and strong consistency (e.g., Google's Spanner) for their strong guarantees and ease of use. Unfortunately, existing transactional storage systems are expensive to use because they rely on expensive replication protocols like Paxos for fault-tolerance. In this talk, we take a new approach to make transactional storage systems more affordable; we eliminate consistency from the replication protocol, while still providing distributed transactions with strong consistency to applications. This paper presents TAPIR -- the Transaction Application Protocol for Inconsistent Replication -- the first transaction protocol to use a replication protocol, inconsistent replication, that provides fault-tolerance with no consistency. By enforcing strong consistency only in the transaction protocol, TAPIR is able to commit transactions in a single round-trip and schedule distributed transactions with no centralized coordination. We demonstrate the use of TAPIR in TAPIR-KV, a key-value store that provides high-performance transactional storage. Compared to system using conventional transaction protocols that require replication with strong consistency, TAPIR-KV has 2x better latency and throughput.

- **11:50-12:05: *Claret: Handling contention in distributed systems with abstract data types*, [Brandon Holt](#)**

Interactive distributed applications like Twitter or eBay are difficult to scale because of the high degree of writes or update operations. The highly skewed access patterns exhibited by real-world systems lead to high contention in datastores, causing periods of diminished service or even catastrophic failure. There is often sufficient concurrency in these applications to scale them without resorting to weaker models like eventual consistency, but traditional concurrency control mechanisms, which operate on low level operations, are unable to detect it.

Instead, we propose that programmers should express their high-level application semantics to datastores through abstract data types (ADTs), letting the programmer and datastore alike reason about their logical behavior, such as operation commutativity. The datastore can use this knowledge to avoid unnecessary conflicts or reduce communication, while the programmer can tune their performance by carefully choosing the right ADTs for their application or implementing custom ones which expose the maximum amount of concurrency.

In this talk, I will describe how ADTs are represented in Claret, our prototype ADT-store, and the ways in which Claret leverages its knowledge to expose concurrency. We show that these ADT-aware optimizations allow transactions to scale, even on benchmarks modelling real-world contention, including an online auction service and a Twitter clone.

- **12:05-12:20: *Subways: A Case for Redundant, Inexpensive Data Center Edge Links*, [Vincent Liu](#)**

As network demand increases, data center network operators face a number of challenges including the need to add capacity to the network. Unfortunately, network upgrades can be an expensive proposition, particularly at the edge of the network where most of the network's cost lies.



This talk presents a quantitative study of alternative ways of wiring multiple server links into a data center network. In it, we propose and evaluate Subways, a new approach to wiring servers and Top-of-Rack (ToR) switches that provides an inexpensive incremental upgrade path as well as decreased network congestion, better load balancing, and improved fault tolerance. Our simulation-based results show that Subways significantly improves performance compared to alternative ways of wiring the same number of links and switches together.

- ***Natural Language Processing (CSE 691)***

- **11:15-11:20: Towards Broad Coverage Language Understanding, [Luke Zettlemoyer](#)**
- **11:20-11:40: *Freebase Semantic Parsing: Question Answering and Information Extraction*, [Eunsol Choi](#)**

We consider the challenge of learning semantic parsers that scale to large, open-domain problems, such as question answering and information extraction with Freebase. In such settings, the sentences cover a wide variety of topics and include many phrases whose meaning is difficult to represent in a fixed target ontology. We develop a parser using a probabilistic CCG to build linguistically motivated logical form meaning representations, and includes an ontology matching model that adapts the output logical forms for each target ontology. Furthermore, when concepts lies outside a fixed target ontology, the parser builds partial analyses that ground as much of an input text as possible. Experiments demonstrate strong performance on benchmark semantic parsing datasets as well as on a new large-scale dataset.

- **11:40-12:00: *Understanding Time and Events in Text*, [Kenton Lee](#)**

The essence of understanding a textual narrative is determining what and when things happened. Towards the goal of automatically constructing timelines from text, we present methods and data for understanding times and events. We first present a context-dependent temporal semantic parser that detects and resolve time expressions such as "two weeks earlier". Our approach produces compositional meaning representations with CCG while considering contextual cues, such as the dependency structure and other mentioned time expressions. Experiments show 13% to 21% improvement over benchmark datasets. We also present a new dataset for event detection and factuality, where we use non-experts to label events that are mentioned and scalar labels indicating the author's certainty in whether the event happened. We provide several models and baselines for this task, highlighting the challenges of automatically labeling event factuality.

- **12:00-12:20: *Question-Answer Driven Semantic Role Labeling: Using Natural Language to Annotate Natural Language*, [Lehung Le](#)**

This talk describes the task of question-answer driven semantic role labeling (QA-SRL), where question-answer pairs are used to represent predicate-argument structure. For example, the verb "introduce" in the previous sentence would be labeled with the questions "What is introduced?", and "What introduces something?", each paired with the phrase from the sentence that gives the correct answer. Posing the problem this



way allows the questions themselves to define the set of possible roles, without the need for predefined frame or thematic role ontologies. It also allows for scalable data collection by annotators with very little training and no linguistic expertise. We gather data in two domains, newswire text and Wikipedia articles, and introduce simple classifier-based models for predicting which questions to ask and what their answers should be. Our results show that non-expert annotators can produce high quality QA-SRL data, and also establish baseline performance levels for future work on this task.

## Session II

- ***Big Data Management and Analytics (CSE 305)***

- **1:30-1:35: Introduction and Overview, [Alvin Cheung](#)**
- **1:35-1:50: *Leveraging Lock Contention to Improve Transactional Application Performance*, [Cong Yan](#)**

Locking is one of the predominant costs in transaction processing. While much work has focused on designing efficient locking mechanisms, not much has been done on understanding applications and using their semantics to improve performance. This paper presents QURO, a query-aware compiler that automatically reorders queries in transaction code to improve performance. Utilizing that certain queries within a transaction are more likely to wait on locks and block the transaction execution, QURO automatically changes the application such that these contentious queries are issued as late as possible. We have evaluated QURO on various transaction benchmarks, and results show that implementations generated by QURO can increase transaction throughput by up to 6.53x, while reduce transaction latency by 85%.

- **1:50-2:05: *Asynchronous and Fault-Tolerant Recursive Datalog Evaluation in Shared-Nothing Engine*, [Jingjing Wang](#)**

We present a new approach for data analytics with iterations. Users express their analysis in Datalog with bag-monotonic aggregate operators, which enables the expression of computations from a broad variety of application domains. Queries are translated into query plans that can execute in shared-nothing engines, are incremental, and support a variety of iterative models (synchronous, asynchronous, different processing priorities) and failure-handling techniques. The plans require only small extensions to an existing shared-nothing engine, making the approach easily implementable. We implement the approach in the Myria big-data management system and use our implementation to empirically study the performance characteristics of different combinations of iterative models, failure handling methods, and applications. Our evaluation uses workloads from a variety of application domains. We find that no single method outperforms others but rather that application properties must drive the selection of the iterative query execution model.

- **2:05-2:20: *Data Transfer between Database Systems using Automatic Connector Generation*, [Brandon Haynes](#)**



Many modern database systems include the ability to import and export in bulk to and from the file system. To date, users who wish to move data between pairs of such systems must write dedicated transfer logic or rely on a common intermediate format (e.g., CSV). Unfortunately, the first approach requires creating connections between every pair of systems, while the second requires an expensive write/read cycle via the file system. In this work we propose an algorithm for the automated generation of connecting operators between arbitrary pairs of compatible database systems, without the need to materialize the underlying data on the file system. Our connectors perform approximately as well as manually-constructed versions, and enable users to easily leverage functionality in external systems for which no transfer mechanism exists.

- **2:20-2:35: *SlimShot: Probabilistic Inference for Web-Scale Knowledge Bases*, [Eric Gribkoff](#)**

We introduce SlimShot (Scalable Lifted Inference and Monte Carlo Sampling Hybrid Optimization Technique), a probabilistic inference engine for Web-Scale knowledge bases. SlimShot converts the inference task to a tuple-independent probabilistic database, then uses a simple Monte Carlo-based inference algorithm with three key enhancements: (1) it combines sampling with safe query evaluation, (2) it estimates a conditional probability by jointly computing the numerator and denominator, and (3) it adjusts the proposal distribution based on the sample cardinality. In combination, these three techniques allow us to give formal error guarantees, and we demonstrate empirically that SlimShot outperforms today's state of the art probabilistic inference engines used in knowledge bases.

- ***Systems, Architecture, and Networking (SANE) 2 (CSE 403)***

- **1:30-1:35: Introduction and Overview, [Dan Ports](#)**
- **1:35-1:55: *Architectural Support for Large-Scale Visual Search*, [Carlo del Mundo](#)**

Emerging classes of computer vision applications demand unprecedented computational resources and operate on tremendous amounts of data. In particular, k-nearest neighbors (kNN), a cornerstone algorithm in these applications, incurs significant data movement. To address this challenge, the underlying architecture and memory subsystems must vertically evolve to address memory bandwidth and compute demands.

To enable large-scale computer vision, we are designing a class of associative memories called NCAMs which encapsulate logic with the memory system to accelerate k-nearest neighbors. The NCAM addresses the shortcomings of current architectures by providing scalable and exact nearest neighbor searching. Finally, we demo an accompanying application, VIRAL, a content-based search system for Video and Image Retrieval.

- **1:55-2:15: *MCDNN: An Execution Framework for Deep Neural Networks on Resource-Constrained Devices*, [Seungyeop Han](#)**

Deep Neural Networks (DNNs) have become the computational tool of choice for many applications relevant to mobile devices. However, given their high memory and



computational demands, running them on mobile devices has required expert optimization or custom hardware. We present a framework that, given an arbitrary DNN, compiles it down to a resource-efficient variant at modest loss in accuracy. Further, we introduce novel techniques to specialize DNNs to contexts and to share resources across multiple simultaneously executing DNNs. Finally, we present a run-time system for managing the optimized models we generate and scheduling them across mobile devices and the cloud. Using the challenging continuous mobile vision domain as a case study, we show that our techniques yield very significant reductions in DNN resource usage and perform effectively over a broad range of operating conditions.

- **2:15-2:35: *Defending Applications Against File System Crashes*, [James Bornholt](#)**

Applications depend on persistent storage to recover state after system crashes. But application writers make implicit assumptions about file system operations, which are violated by performance optimizations and caching behaviors of modern operating systems and hardware. When a system crashes, these mistakes can lead to corrupt application state and, as in several recent examples, catastrophic data loss.

We have developed crash-consistency models for modern file systems, which describe the behavior of a file system across crashes. To put these models into practice, we have developed Ferrite, a toolkit for validating crash-consistency models against real file system implementations. We have used Ferrite to build proof-of-concept tools to verify the crash safety of applications, and to automatically synthesize synchronization code for an application to safely store its state.

- ***Innovations in Mobile Systems (CSE 691)***

- **1:30-1:35: Introduction and Overview, [Shyam Gollakota](#)**
- **1:35-1:50: *Bringing Low-Power to Wi-Fi transmissions*, Bryce Kellogg**

We show for the first time how to generate Wi-Fi transmissions using tens of microwatts of power. This is 10000x lower-power than existing Wi-Fi chipsets and 1000x better than ZigBee and Bluetooth LTE. Since Wi-Fi has traditionally been considered to be high power, this innovation can have a significant impact on the Wi-Fi industry as well as on the IOT and sensor network space.

- **1:50-2:05: *Powering the Next Billion Devices using Wi-Fi*, Vamsi Talla**

We present the first power over Wi-Fi system that delivers power to low-power sensors and devices and works with existing Wi-Fi chipsets. Specifically, we show that a ubiquitous part of wireless communication infrastructure, the Wi-Fi router, can provide far field wireless power without significantly compromising the network's communication performance. Building on our design, we prototype battery-free temperature and camera sensors that we power with Wi-Fi at ranges of 20 and 17 feet respectively. We also demonstrate the ability to wirelessly trickle-charge nickel-metal hydride and lithium-ion coin-cell batteries at distances of up to 28 feet. We deploy our system in six homes in a metropolitan area and show that it can successfully deliver



power via Wi-Fi under real-world network conditions without significantly degrading network performance.

- **2:05-2:20: *Contactless Sleep Apnea Diagnosis on Smartphones*, Rajalakshmi Nandakumar**

We present a contactless solution for detecting sleep apnea events on smartphones. To achieve this, we introduce a novel system that monitors the minute chest and abdomen movements caused by breathing on smartphones. Our system works with the phone away from the subject and can simultaneously identify and track the fine-grained breathing movements from multiple subjects. Building on the above system, we develop algorithms that identify various sleep apnea events including obstructive apnea, central apnea, and hypopnea from the sonar reflections. We deploy our system at the UW Medicine Sleep Center at Harborview and perform a clinical study with 37 patients for a total of 296 hours. Our study demonstrates that the number of respiratory events identified by our system is highly correlated with the ground truth and has a correlation coefficient of 0.9957, 0.9860, and 0.9533 for central apnea, obstructive apnea and hypopnea respectively. Furthermore, the average error in computing of rate of apnea and hypopnea events is as low as 1.9 events/hr.

- **2:20-2:35: *Fine-Grained finger tracking using Active Sonar*, Rajalakshmi Nandakumar**

We present fingerIO, a novel fine-grained finger tracking solution for around-device interaction. FingerIO does not require instrumenting the finger with sensors and works even in the presence of occlusions between the finger and the device. Our evaluation shows that fingerIO can achieve finger tracking with an average accuracy of 8 mm using the in-built microphones and speaker of a Samsung Galaxy S4. It also tracks subtle finger motion around the device, even when the phone is in the pocket. Finally, we prototype a smart watch form-factor fingerIO device and show that it can extend the interaction space to a  $0.5 \times 0.25 \text{ m}^2$  region on either side of the device and work even when it is fully occluded from the finger.

## Session III

- ***Human Computer Interaction and Self-Tracking (CSE 305)***

- **2:40-2:45: Introduction and Overview, [James Fogarty](#)**
- **2:45-2:55: *Barriers and Negative Nudges: Exploring Challenges in Food Journaling*, [Daniel Epstein](#)**

Although food journaling is understood to be both important and difficult, little work has empirically documented the specific challenges people experience with food journals. We identify key challenges in a qualitative study combining a survey of 141 current and lapsed food journalers with analysis of 5,526 posts in community forums for three mobile food journals. Analyzing themes in this data, we find and discuss barriers to reliable food entry, negative nudges caused by current techniques, and challenges with social features. Our results motivate research exploring a wider range of approaches to food journal design and technology.



- **2:55-3:05: *Rethinking the Mobile Food Journal: Exploring Opportunities for Lightweight Photo-Based Capture*, [James Fogarty](#)**

Food choices are among the most frequent and important health decisions in everyday life, but remain notoriously difficult to capture. This work examines opportunities for lightweight photo-based capture in mobile food journals. We first report on a survey of 257 people, examining how they define healthy eating, their experiences and challenges with existing food journaling methods, and their ability to interpret nutritional information that can be captured in a food journal. We then report on interviews and a field study with 27 participants using a lightweight, photo-based food journal for between 4 to 8 weeks. We discuss mismatches between motivations and current designs, challenges of current approaches to food journaling, and opportunities for photos as an alternative to the pervasive but often inappropriate emphasis on quantitative tracking in mobile food journals.

- **3:05-3:25: *Opportunities and Challenges for Self-Experimentation in Self-Tracking*, [Ravi Karkar](#)**

Personal informatics applications support capture and access of data related to an increasing variety of dimensions of everyday life. However, such applications often fail to effectively support diagnostic self-tracking, wherein people seek to answer a specific question about themselves. Current approaches are therefore difficult, tedious, and error-prone. We discuss our ongoing efforts to develop methods for self-experimentation in self-tracking. We examine how self-experimentation situates within existing models of personal informatics processes, discuss our current focus on personal food triggers in patients suffering from Irritable Bowel Syndrome, and highlight open challenges for self-experimentation more broadly.

- **3:25-3:35: *Examining Unlock Journaling with Diaries and Reminders for In Situ Self-Report in Health and Wellness*, [Xiaoyi Zhang](#)**

In situ self-report is widely used in human-computer interaction, ubiquitous computing, and for health and wellness assessment and intervention. Unfortunately, it remains limited by its high burdens. We build upon recent proposals of mobile lockscreen interaction to examine unlock journaling in health and wellness. Specifically, we introduce single-slide unlock journaling gestures appropriate for health and wellness measures. We then present the first field study comparing unlock journaling with traditional diaries and notification-based reminders for self-report of health and wellness measures. We find unlock journaling is less intrusive than reminders, dramatically improves frequency of journaling, and can provide equal or better timeliness. Where appropriate to overall design needs, unlock journaling is an overall promising method for in situ self-report.

- **3:35-3:45: *A Lived Informatics Model of Personal Informatics*, [Daniel Epstein](#)**

Current models of how people use personal informatics systems are largely based in behavior change goals. They do not adequately characterize the integration of self-tracking into everyday life by people with varying goals. We build upon prior work by embracing the perspective of lived informatics to propose a new model of personal



informatics. We examine how lived informatics manifests in the habits of self-trackers across a variety of domains, first by surveying 105, 99, and 83 past and present trackers of physical activity, finances, and location and then by interviewing 22 trackers regarding their lived informatics experiences. We develop a model characterizing tracker processes of deciding to track and selecting a tool, elaborate on tool usage during collection, integration, and reflection as components of tracking and acting, and discuss the lapsing and potential resuming of tracking. We use our model to surface underexplored challenges in lived informatics, thus identifying future directions for personal informatics design and research.

- **Next Generation Programming Tools: Performance (CSE 403)**

- **2:40-2:45: Introduction and Overview, [Emina Torlak](#)**
- **2:45-3:05: *Translating Sequential Java Programs to MapReduce Using Verified Lifting*, [Maaz Bin Safeer Ahmad](#)**

With the advent of "big-data," numerous parallel data processing engines have emerged, with each of these engines offering its own optimizations for different application domains. Choosing the right engine and migrating code across these engines is often a slow and expensive process that require code rewrites and learning new programming languages. Our goal is to leverage program synthesis and verified lifting to automatically extract code fragments written in a general-purpose language and rewrite them to leverage parallel data processing engines. In this talk, I will describe a compiler that automatically identifies code fragments written in sequential Java and rewrite them to be executed on the Hadoop parallel processing engine. Our preliminary results have shown that the rewritten code fragments can improve performance by up to 5x as compared to the original application.

- **3:05-3:25: *Synthesis-Aided Compilers for Low-Power Hardware*, [Ras Bodik](#)**

How do you make a processor ultra-low-power? You strip it of all features designed for programmability, such as hardware-controlled caches. How do you program such a processor? Through expensive manual labor of highly skilled programmers. These questions illustrate that the end of Moore's Law increased the tension between energy efficiency and programmability. Can we hope to attain both? This talk will describe Chlorophyll, a compiler for the a minimalistic low-power spatial architecture that requires partitioning the program into fragments. The Chlorophyll programming model allows programmers to provide human insight by specifying partial partitioning of data and computation. The novelty of the Chlorophyll compiler is that it relies on program synthesis, sidestepping the need to develop optimizations for the unusual architecture. We discuss why Chlorophyll generates code that rivals expert-written code and how Chlorophyll can be ported to other unusual architectures.

- **3:25-3:45: *Herbie: Automatically Fixing Double Trouble*, [Pavel Panchekha](#)**

Scientific and engineering applications depend on floating point arithmetic to approximate real arithmetic. This approximation introduces rounding error, which can



accumulate to produce unacceptable results. While the numerical methods literature provides techniques to mitigate rounding error, applying these techniques requires manually rearranging expressions and understanding the finer details of floating point arithmetic.

We introduce Herbie, a tool which automatically discovers the rewrites experts perform to improve accuracy. Herbie's heuristic search estimates and localizes rounding error using sampled points (rather than static error analysis), applies a database of rules to generate improvements, takes series expansions, and combines improvements for different input regions. We evaluated Herbie on examples from a classic numerical methods textbook, and found that Herbie was able to improve accuracy on each example, some by up to 60 bits, while imposing a median performance overhead of 40%. The first paper on Herbie appeared at PLDI 2015 where it received a Best Paper award.

- ***New Sensing and Interaction Technologies (CSE 691)***

- **2:40-2:45: Introduction and Overview, [Mayank Goel](#)**
- **2:45-3:05: *MagnifiSense: Inferring Device Interaction using Wrist-Worn Passive Magneto-Inductive Sensors*, [Edward Wang](#)**

The different electronic devices we use on a daily basis produce distinct electromagnetic radiation due to differences in their underlying electrical components. We present MagnifiSense, a low-power wearable system that uses three passive magneto-inductive sensors and a minimal ADC setup to identify the device a person is operating.

MagnifiSense achieves this by analyzing near-field electromagnetic radiation from common components such as the motors, rectifiers, and modulators. We conducted a staged, in-the-wild evaluation where an instrumented participant used a set of devices in a variety of settings in the home such as cooking and outdoors such as commuting in a vehicle. Through the various studies, we found that MagnifiSense shows promise as a technique for real world applications that need tracking of an individual's activity throughout the day.

- **3:05-3:25: *Detecting Human-Object Interaction, and Fabricating Paper Interfaces Using UHF RFID*, [Hanchuan Li](#)**

In order to enable unobtrusive human-object interaction detection, we propose a minimalistic approach to instrumenting everyday objects with passive (i.e. battery free) UHF RFID tags. By measuring the changes in the physical layer of the communication channel between the RFID tag and reader (such as RSSI, RF phase, and read rate) we are able to classify, in real-time, tag/object motion events along with two types of touch events. We also describe techniques that allow inexpensive, ultra-thin, RFID tags to be turned into simple paper input devices. We use sensing and signal processing techniques that determine how a tag is being manipulated by the user via an RFID reader. Our techniques provide the capability to sense touch, cover, overlap of tags by conductive or dielectric (insulating) materials, as well as tag movement trajectory. Due to the rapid deployability and low cost of the tags used, we can create a new class of



interactive paper devices that are drawn on demand for simple tasks. These capabilities allow new interactive possibilities for pop-up books and other paper craft objects.

- **3:25-3:45: *Finexus: Tracking Precise Motions of Multiple Fingertips Using Magnetic Sensing*, [Keyu Chen](#)**

With the resurgence of head-mounted displays for virtual reality, users need new input devices that can accurately track their hands and fingers in motion. We introduce Finexus, a multipoint tracking system using magnetic field sensing. By instrumenting the fingertips with electromagnets, the system is able to track fine fingertip movements in real time using only four magnetometers. To keep the system robust to noise, we operate each electromagnet at a different frequency and leverage bandpass filters in order to distinguish signals attributed to individual sensing points for localization. We develop a novel algorithm to efficiently calculate multiple electromagnets' 3D positions from corresponding field strengths. In our evaluation, we report an average accuracy of 0.95 mm (fixed orientation) and 1.33 mm (random orientation) compared to optical trackers. Our real-time implementation enables Finexus to be applicable to a wide variety of high-precision real-time human input tasks.

## Session IV

- ***Computer Graphics and Vision at UW (CSE 305)***

- **3:50-3:55: Introduction and Overview, [Ira Kemelmacher-Shlizerman](#)**
- **3:55-4:15: *Generating Notifications For Missing Actions: Don't forget to turn the lights off!* [Bilge Soran](#)**
- **4:15-4:35: *Time-lapse Mining from Digital Photos*, [Ricardo Martin](#)**

We introduce an approach for synthesizing time-lapse videos of popular landmarks from large community photo collections. The approach is completely automated and leverages the vast quantity of photos available online. First, we cluster 86 million photos into landmarks and popular viewpoints. Then, we sort the photos by date and warp each photo onto a common viewpoint. Finally, we stabilize the appearance of the sequence to compensate for lighting effects and minimize flicker. Our resulting time-lapses show diverse changes in the world's most popular sites, like glaciers shrinking, skyscrapers being constructed, and waterfalls changing course.

- **4:35-4:55: *Digitizing Persona*, [Supasorn Suwajanakorn](#)**

We reconstruct a controllable model of a person from a large photo collection that captures his or her persona, i.e., physical appearance and behavior. The ability to operate on unstructured photo collections enables modeling a huge number of people, including celebrities and other well photographed people without requiring them to be scanned. Moreover, we show the ability to drive or puppeteer the captured person B using any other video of a different person A. In this scenario, B acts out the role of person A, but retains his/her own personality and character. Our system is based on a novel combination of 3D face reconstruction, tracking, alignment, and multi-texture



modeling, applied to the puppeteering problem. We demonstrate convincing results on a large variety of celebrities derived from Internet imagery and video.

- **Next Generation Programming Tools: Correctness (CSE 403)**

- **3:50-3:55: Introduction and Overview, [Ras Bodik](#)**
- **3:55-4:15: *Composing Reliable Distributed Systems with Verdi*, [Zach Tatlock](#)**

On a single day this summer the New York Stock Exchange halted trading, the *Wall Street Journal* website went down, and United Airlines was forced to ground all flights, all due to bugs in their distributed systems. Despite billions of dollars invested in development and testing, even the most highly skilled programmers continue to make disastrous mistakes that bring down critical services. The problem is that distributed systems are difficult to implement correctly because they must handle both concurrency and failure: machines may crash at arbitrary points and networks may reorder, drop, or duplicate packets. This set of behaviors is simply too complex to permit effective testing.

To address this problem, we developed Verdi, a framework for implementing and formally verifying distributed systems implementations. By proving systems correct in Verdi, developers can ensure the absence of the sorts of bugs that caused so much mayhem this summer. We will briefly discuss the first mechanically checked correctness proof for the Raft consensus protocol and the challenges we encountered completing this landmark verification using Verdi.

- **4:15-4:35: *Bagpipe: Verified, Centralized Internet Routing Policies*, [Konstantin Weitz](#)**

Internet Service Providers (ISPs) use the Border Gateway Protocol to exchange path information for routing packets through the Internet. ISPs implement complex security- and business-driven BGP policies by configuring the BGP behavior of their networks. The consequences of misconfiguration have been highly visible, e.g. the worldwide downtime of YouTube in 2008. The problem is that an ISP's routers are each individually configured and must work together to implement the ISP's high level policy. Unfortunately this task has consistently proven incredibly challenging in practice and failures have led to lost time and money.

To address this problem, we built Bagpipe, a tool for automatically verifying that ISP's router configurations correctly enforce its BGP policies. Bagpipe is an efficient, massively parallel BGP configuration checking tool. It provides the first formal semantics of the BGP protocol and BGP policies (based on RFC 4271 and modeled in the Coq proof assistant). It also introduces the empty network reduction, a technique at the core of Bagpipe that enables efficient verification of a restricted but useful subset of BGP policies.

- **4:35-4:55: *Build your own type system for fun and profit*, [Michael Ernst](#)**

Are you tired of null pointer exceptions, unintended side effects, SQL injections, concurrency errors, mistaken equality tests, and other run-time errors? Are your users tired of them in your code? This presentation shows you how to guarantee, at compile



time, that these run-time exceptions cannot occur. You have nothing to lose but your bugs!

Formal verification is often considered an abstruse art: it takes a lot of training to be able to formally verify a program, and it takes even more effort to create a formal verification system. We show that these assumptions are no longer true. Formal verification can be as natural to programmers as type-checking, and even novices can create their own type system to verify important properties of your code.

- **Health Sensing (CSE 691)**

- **3:50-3:55: Introduction and Overview, [Alex Mariakakis](#)**

- **3:55-4:15: *HyperCam: Hyperspectral Imaging and their Applications*, [Mayank Goel](#)**

Emerging uses of imaging technology for consumers cover a wide range of application areas from health to interaction techniques; however, typical cameras primarily transduce light from the visible spectrum into only three overlapping components of the spectrum: red, blue, and green. In contrast, hyperspectral imaging breaks down the electromagnetic spectrum into more narrow components and expands coverage beyond the visible spectrum. While hyperspectral imaging has proven useful as an industrial technology, its use as a sensing approach has been fragmented and largely neglected by the UbiComp community. HyperCam explores an approach to make hyperspectral imaging easier and bring it closer to the end-users. It provides a low-cost implementation of a multispectral camera and a software approach that automatically analyzes the scene and provides a user with an optimal set of images that try to capture the salient information of the scene. We present a number of use-cases that demonstrate HyperCamb's usefulness and effectiveness.

- **4:15-4:35: *SpiroCall: Tracking Lung Function over a Phone Call*, [Elliot Saba](#)**

Cost and accessibility have impeded the adoption of spirometers (devices that measure lung function) outside clinical settings, especially in low-resource environments. Prior work, called SpiroSmart, used a smartphone's built-in microphone as a spirometer. However, individuals in low- or middle-income countries do not typically have access to the latest smartphones. In this talk, we investigate how spirometry can be performed from any phone - using the standard telephony voice channel to transmit the sound of the spirometry effort. We also investigate how using a 3D printed vortex whistle can affect the accuracy of common spirometry measures, consistency, and mitigate usability challenges. Our system, coined SpiroCall, was evaluated with 50 participants against two gold standard medical spirometers. We conclude that SpiroCall has an acceptable mean error with or without a whistle for performing spirometry in developing countries, and advantages of each are discussed.

- **4:35-4:45: *SleepSense: Accuracy of Sleep Stage Mining across Different Sensing Modalities*, [Ruth Ravichandran](#)**

Consuming almost a third of our daily lives, sleep is a significant marker of an individual's health and well-being. Various commercial wearable devices are available



# COMPUTER SCIENCE & ENGINEERING

---

UNIVERSITY *of* WASHINGTON

that track sleep stages and sleep quality using different physiological signals such as body movement, breathing and hear rate. In this talk we explore unique features provided by the three physiological signals for sleep stage mining and compare the compare the results for subsets of these sensing modalities. We also compare the results with a proposed non-contact single sensor solution for tracking all three physiological signals using an off-the-shelf radar module operating at 24GHz.



## Keynote Address

Tuesday, October 20, 2015: 1:00 p.m.

[Paul G. Allen Center](#), Microsoft Atrium

**Noah Smith, Associate Professor**  
**UW Computer Science & Engineering**

---

### Discovering the Social World using Text

#### Abstract

Research in natural language processing is increasingly taking into account the context -- physical and social -- in which linguistic communication occurs. In this talk, I will present examples of models of text data that exploit assumptions about the social world from which those data emerged to further our understanding. First, we'll look at how city-scale linguistic groups in the US influence each other, using a large collection of social media messages. Second, we'll turn to models of how a very small, powerful set of individuals -- the Supreme Court justices -- are influenced by highly-incentivized generators of text.



#### Speaker Bio

Noah Smith is Associate Professor of Computer Science & Engineering at the University of Washington. Previously, he was on the faculty of the Language Technologies and Machine Learning Department in the School of Computer Science at Carnegie Mellon University (2006–2015). His research tackles many problems in natural language processing, machine learning, and their applications, especially to the social sciences. His work and that of his research group, Noah's ARK, has been recognized with a Finmeccanica career development chair at CMU (2011–2014), an NSF CAREER award (2011–2016), a Hertz Foundation graduate fellowship at Johns Hopkins University (2001–2006), numerous best paper nominations and awards, and coverage by NPR, BBC, CBC, the *New York Times*, *Washington Post*, and *Time*.



**Poster and Demo Session**

Tuesday, October 20: 5-7 pm, throughout the Paul G. Allen Center  
Madrona Prize and People's Choice Award Presentation, 7:15pm in the Atrium  
*(Last updated 10/19/2015)*

**PEOPLE'S CHOICE AWARD and SOCIAL IMPACT AWARD for Best Poster:**

Please be sure to vote for your favorite poster in **EACH LAB THAT YOU VISIT**. Ballot boxes are available in each lab, as well as blank ballots. Please vote by writing the number of the poster on your ballot and dropping it into a ballot box. Voting will wrap up by 7:00 p.m.

Please check the following link for the most up-to-date list of posters:

<http://tinyurl.com/CSEposters2015>

**Poster and Demo Session**

**October 20, 2015, 5:00-7:00 p.m.**

unique  
poster #

**Graphics, Vision, Animation, and Games for Learning (291)**

- 1 **Deep Photo Tagging**  
*Hamid Izadinia*
- 2 **Depth from Focus with Your Mobile Phone**  
*Supasorn Suwajanakorn*
- 3 **Digital Pathology: Diagnostic Error Analysis**  
*Ezgi Mercan, Linda G. Shapiro, Tad Brunye, Don L. Weaver, Joann G. Elmore*
- 4 **Floor plan reconstruction**  
*Hamid Izadinia*
- 5 **In-situ CAD**  
*Aditya Sankar, Steve Seitz*
- 6 **Light Field Layer Matting**  
*Juliet Fiss, Brian Curless, Rick Szeliski*
- 7 **Object Detection Demo**  
*Joseph Redmon*
- 8 **Object Discovery and Visual Reasoning**  
*Hamid Izadinia and Fereshteh Sadeghi*
- 9 **Time-lapse mining from Internet photos**  
*Ricardo Martin-Brualla, David Gallup and Steve Seitz*
- 10 **Visual Analogy**  
*Fereshteh Sadeghi*
- 11 **What Makes Tom Hanks Look Like Tom Hanks**  
*Supasorn Suwajanakorn, Steven M. Seitz, Ira Kemelmacher-Shlizerman*



**Computing Accessibility (2nd floor landing)**

- 12 **Universal Play Kiosk**  
*Anat Caspi, Gyorgy Zatloka, Setion Branko, Kevin Porter, Sarah Haase, Milena Batova*
- 13 **AccessComputing: The national alliance for access to computing careers**  
*Richard Ladner*

**Database Systems Lab (405)**

- 14 **A Dynamic Scaling Engine for Data Analytics with Performance Guarantees**  
*Jennifer Ortiz*
- 15 **Asynchronous and Fault-Tolerant Recursive Datalog Evaluation in Shared-Nothing Engines**  
*Jingjing Wang, Magdalena Balazinska, Daniel Halperin*
- 16 **Centiman: Elastic, High Performance Optimistic Concurrency Control by Watermarking**  
*Bailu Ding (Cornell), Lucja Kot (Cornell), Alan Demers (Cornell), Johannes Gehrke (Microsoft)*
- 17 **Data Transfer between Database Systems using Automatic Connector Generation**  
*Brandon Haynes*
- 18 **From Theory to Practice: Efficient Join Query Evaluation in a Parallel Database System**  
*Shumo Chu, Magdalena Balazinska, Dan Suciu*
- 19 **Gaussian Mixture Models Use-Case: In-Memory Analysis with Myria**  
*Ryan Maas*
- 20 **Graphulo: Graph Processing for Accumulo Databases**  
*Dylan Hutchison, Jeremy Kepner, Vijay Gadepally*
- 21 **Leveraging Lock Contention to Improve OLTP Application Performance**  
*Cong Yan, Alvin Cheung*
- 22 **Query engine for high-performance languages**  
*Brandon Myers*
- 23 **Solving Multilinear Polynomials for Data Summarization**  
*Laurel Orr, Magdalena Balazinska, Dan Suciu*
- 24 **Translating Sequential Java Programs to MapReduce Using Verified Lifting**  
*Maaz Ahmad, Alvin Cheung*

**Programming Languages and Software Engineering (407)**

- 25 **Composing Formally Verified Distributed Systems in Verdi**  
*James Wilcox, Doug Woos, & Steve Anton*
- 26 **Data Structure Synthesis**  
*Calvin Loncaric*
- 27 **Dependability Cases for Safety Critical Systems**  
*Stuart Pernsteiner*



- 28 **Formally Verifying Assembly Level Optimizations**  
*Eric Mullen, Daryl Zuniga*
- 29 **Hello Joe: Verified Actor Conversations with Erlang and Dependent Types**  
*Archibald Samuel Elliott, Edwin Brady (University of St. Andrews)*
- 30 **Herbie: Automatically Fixing Double Trouble**  
*Pavel Panchekha, Alex Sanchez-Stern*
- 31 **Incarinate: How Compilers and Vision Might Bring 3D Printing to Life (demo)**  
*Joe Redmon*

## HCI, Accessibility, and Visualization (5th floor landing)

- 32 **A Framework for Self-Experimentation in Personalized Health**  
*Ravi Karkar, James Fogarty, Julie A. Kientz, Sean A. Munson, Roger Vilardaga, Jasmine Zia*
- 33 **AACrobat: Using Mobile Devices to Lower Communication Barriers and Provide Autonomy with Gaze-Based AAC**  
*Alexander Fiannaca, Ann Paradiso, Mira Shah, and Merrie Morris*
- 34 **Being Liked and Being Heard: Use of Visual Feedback in Email Communication**  
*A. Conrad Nied, Margaret Morris, Gary Hsieh, Sean Munson*
- 35 **Designing for user-facing uncertainty in everyday sensing and prediction**  
*Matthew Kay, Tara Kola, Jessica Hullman, Sean Munson, Shwetak Patel, Julie Kientz*
- 36 **Personal guidance on individualized food triggers for IBS patients**  
*Jessica Schroeder, James Fogarty, Sean Munson, Jasmine Zia*
- 37 **Personal Informatics in Everyday Life**  
*Daniel Epstein, James Fogarty, Sean Munson*
- 38 **Repurposing everyday phone interaction**  
*Xiaoyi Zhang, Donny Huang, Laura R. Pina, T. Scott Saponas, James Fogarty, Shyamnath Gollakota*
- 39 **Voyager: Exploratory Analysis via Faceted Browsing of Visualization Recommendations**  
*Kanit Wongsuphasawat, Dominik Moritz, Anushka Anand, Jock Mackinlay, Bill Howe, Jeffrey Heer*

## Systems and Networking, Architecture, and Programming Languages (591)

- 40 **A near-sensor face detection accelerator for low-power camera systems**  
*Amrita Mazumdar, Armin Alaghi, Luis Ceze, Mark Oskin*
- 41 **Apneaapp. Diagnosing sleep apnea**  
*Rajalakshmi Nandakumar*
- 42 **Architectural Support for Large-Scale Visual Search**  
*Carlo C. del Mundo, Vincent T. Lee, Luis Ceze, Mark Oskin*
- 43 **Building Consistent Transactions with Inconsistent Replication**  
*Irene Zhang, Naveen Kr. Sharma, Adriana Szekeres, Arvind Krishnamurthy, Dan R. K. Ports*



- 44 **Claret: Using Data Types for Highly Concurrent Distributed Transactions**  
*Brandon Holt, Irene Zhang, Dan Ports, Mark Oskin, Luis Ceze*
- 45 **Fingerio. Centimeter level finger tracking for around device interaction**  
*Rajalakshmi Nandakumar*
- 46 **Machine-learning Acceleration for Near-Sensor Processing**  
*Sung Kim, Thierry Moreau, Visvesh Sathe, Luis Ceze*
- 47 **New Limits of Knowledge Compilation and Applications to Exact Model Counting**  
*Paul Beame and Vincent Liew*
- 48 **Passive wifi. Wifi at 10000x lower power**  
*Bryce Kellogg*
- 49 **Subways: A Case for Redundant, Inexpensive Data Center Edge Links**  
*Vincent Liu, Danyang Zhuo, Simon Peter, Arvind Krishnamurthy, Thomas Anderson*
- 50 **Transtorm: a benchmark suite for transactional key-value storage systems**  
*Naveen Kr. Sharma, Brandon Holt, Irene Zhang, Dan R. K. Ports, Marcos K. Aguilera*

## AI / Natural Language Processing (609)

- 51 **Event Detection and Factuality Assessment with Non-Expert Supervision**  
*Kenton Lee, Yoav Artzi, Yejin Choi, and Luke Zettlemoyer*
- 52 **Graph-Based Inverse Optimal Control for Robot Manipulation**  
*Arunkumar Byravan, Mathew Monfort, Brian Ziebart, Byron Boots and Dieter Fox*
- 53 **Joint A\* CCG Parsing and Semantic Role Labelling**  
*Mike Lewis, Luheng He and Luke Zettlemoyer*
- 54 **Learning Generalizable Surface Cleaning Actions from Demonstration**  
*Sarah Elliott*
- 55 **Mise en Place: Unsupervised Interpretation of Instructional Recipes**  
*Chloé Kiddon, Ganesa Thandavam Ponnuraj, Luke Zettlemoyer, Yejin Choi*
- 56 **Natural and Visualization-Based Transparency in Situated Language Understanding**  
*Leah Perlmutter, Maya Cakmak*
- 57 **Semantic Role Labeling Improves Semantic Parsing**  
*Ioannis Konstas, Frank Keller*
- 58 **Semantic Role Labeling with Neural Network Factors**  
*Nicholas FitzGerald, Oscar Täckström, Kuzman Ganchev, and Dipanjan Das*
- 59 **Solving geometry problems: combining text and diagram interpretation**  
*Minjoon Seo, Hannaneh Hajishirzi, Ali Farhadi, Oren Etzioni, Clint Malcolm*
- 60 **Towards Never-Ending Object Learning for robots**  
*Yuyin Sun and Dieter Fox*
- 61 **Unsupervised Code-Switching for Multilingual Historical Document Transcription**  
*Dan Garrette, Hannah Alpert-Abrams, Taylor Berg-Kirkpatrick, Dan Klein*
- 62 **Visual Storybooks: Multimodal Schema Learning and Summarization with Photo Albums**  
*Antoine Bosselut, Jianfu Chen, Hannaneh Hajishirzi, Yejin Choi*



## Sensing/Sustainability, HCI, Accessibility (691)

- 63 **Accuray of sleep staging mining across various sensing modalities**  
*Ruth Ravichandran, Hong Lu, Lama Nachman, Julie Kientz, Shwetak Patel*
- 64 **BANDAIDS: Battery Free On-Body Sensing**  
*Josh Fromm, Kyle Lee, Shwetak Patel*
- 65 **BiliCam: Using Mobile Phones to Monitor Newborn Jaundice**  
*Lilian de Greef, Mayank Goel, Min Joon Seo, Eric C. Larson, James W. Stout MD MPH, James A. Taylor MD, Shwetak N. Patel*
- 66 **Finexus: Tracking Precise Motions of Multiple Fingertips Using Magnetic Sensing**  
*Ke-Yu Chen, Shwetak Patel, Sean Keller*
- 67 **HyperCam: Hyperspectral Imaging and their Applications**  
*Mayank Goel, Eric Whitmire, Alex Mariakakis, Scott Saponas, Neel Joshi, Shwetak Patel*
- 68 **Identification of Long Range DNA-DNA Contacts**  
*Jacob Schreiber*
- 69 **Implanted platform for neuromodulation**  
*Brody Mahoney, Vaishnavi Ranganathan, Gregory Moore, Chet Moritz, Bingni Brunton, Polina Anikeeva, Greg Horwitz, Joshua R. Smith*
- 70 **Inferring Power Consumption of Home Electronic Appliances from its EMI**  
*Ke-Yu Chen, Sidhant Gupta, Shwetak Patel*
- 71 **MagnifiSense: Inferring Device Interaction using Wrist-Worn Passive Magneto-Inductive Sensors**  
*Edward Wang, Tien-Jui Lee, Alex Mariakakis, Mayank Goel, Sidhant Gupta, Shwetak Patel*
- 72 **PaperID: A Technique for Prototyping Battery-Free Wireless Interfaces on Paper**  
*Hanchuan Li, Eric Brockmeyer, Elizabeth Carter, Josh Fromm, Scott Hudson, Shwetak Patel, Alanson Sample*
- 73 **RePOV: Using Sensors and Vision to Facilitate Discoveries in Egocentric Videos**  
*Alex Mariakakis, Daniel Avrahami, Sven Kratz*
- 74 **SpiroCall: Tracking Lung Function over a Phone Call**  
*Mayank Goel, Elliot Saba, Maia Stiber, Eric Whitmire, Josh Fromm, Eric C. Larson, Gaetano Borriello, Shwetak N. Patel*
- 75 **UbiComp Lab Overview**  
*UbiComp Lab*



Building Maps

2<sup>nd</sup> floor





Building Maps

4th floor





Building Maps

5<sup>th</sup> floor





Building Maps

6<sup>th</sup> floor





### **Network Access for Allen Center Visitors**

#### **How Do I Connect to the UW Wireless Network?**

The UW's network is named "University of Washington" and is the one you should use when on a UW campus. You may see other networks listed, but they are not supported by UW Information Technology and may be unsafe to use.

#### **Logging In**

No login is required to use UW's wireless service unless you want to connect to services or sites outside the UW.

- If you use your browser to connect to a non-UW Web site, you automatically will be prompted to enter your UW NetID and password.
- If you are using a non-Web service (e.g., a secure terminal or file transfer tool) to communicate outside the UW, you must FIRST run a browser and view a Web page outside the UW to bring up the authentication page.

**UW NetID: event0154**

**Password: cPUa+xWUo+xPAo**

Once you have successfully logged in (authenticated) you will be able to use services outside the UW for up to 12 hours without having to re-authenticate.

---



**University of Washington Computer Science & Engineering  
Recruiting Policy for Employers**

All employment professionals recruiting Computer Science & Engineering students through the CSE department and other on-campus recruiting channels, such as the UW Career Center and other established College or University programs, are required to work within a framework of professionally accepted recruiting, interviewing and selection techniques as stipulated in the [NACE Principles for Professional Conduct](#).

To provide students with sufficient time to complete their employment search and make well informed decisions, we ask employers to abide by the deadlines outlined below. Failure to follow these policies may be detrimental to our students and may result in limiting employer access to departmental recruiting activities.

**Students shall be given a minimum of 2 weeks to decide whether to accept a full time offer or internship offer, or until the following deadlines, whichever is later. Exploding offers and special incentives such as diminishing bonuses or location preferences should not be used to encourage early acceptance of an offer. Signing bonuses should be honored whenever a student accepts an offer.**

*Full-time offers to previous summer interns*

- For full-time offers extended to previous summer interns, students have until November 30 to make a decision.

*New offers for full-time or interns*

- Students extended offers during Autumn quarter have until November 30 to make a decision.
- Students extended offers during Winter break have until February 15 to make a decision.
- Students extended offers during Winter quarter have until February 15 to make a decision.
- Students extended offers during Spring quarter have until April 15 to make a decision.

*Internship offers extended to previous summer interns*

- For new internship offers extended to previous summer interns, students have until February 15 to make a decision.

Employers should be prepared to respond to status inquiries from students within one week of each interview, whether conducted on campus, off campus, or over the phone, for both full-time and internship positions.

\*\*Please note that the time periods above start at the receipt of a written offer for employment to the student.

