



Microsoft Research

Faculty Summit

2014 15TH ANNUAL

Hot Topics

Jeannette Wing
Corporate Vice President
Microsoft Research





Microsoft Research
Faculty
Summit
2014 15TH ANNUAL

The Second Age of Computing is Ending

Doug Burger, Microsoft Research



A New Age of Specialization

Analog Specialization

100 BC – 1936

Antikythera mechanism, slide rule

1st Age

Invention

1936 – 1975

Instruction sets, VM, OoO, caches

Integration

1975 - 1990

RISC, single-chip CPUs, integrated FPUs, caches

Clock Frequency (+ ILP)

1990 - 2005

Deep pipelines, speculation, large caches

Multicore

2005 - 2016

1 to 24 cores, on-chip networks

2nd Age

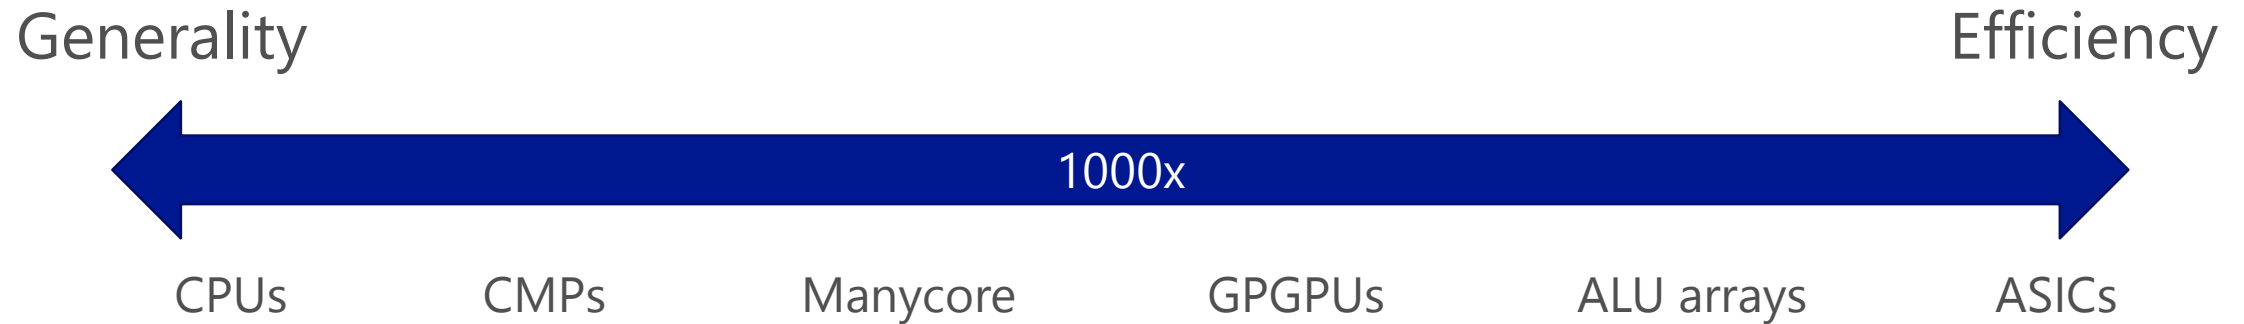
Hardware Specialization

2016 - ?

Programmable logic, rapid ASICs, new archs, statistical models of computation


3rd Age

Challenge of HW Specialization



- Cloud: Two main challenges for specialization
 - Want homogeneous (to the extent possible) server infrastructure
 - Need five years of stability for ASICs (2 to design, 3 for use), software changes monthly
- Client:
 - Area is precious, must be both general and efficient
 - “Uncanny valley” between CPUs and ASICs (where accelerators go to die)

Catapult – A Specialization Fabric for the Cloud



Demonstrated Bing ranking in half the servers for <math><10\%</math> power increase, able to run much larger ranking models in equivalent latency. [ISCA2014]

Programmable SW fabric

Programmable HW fabric

Looking Forward

Hardware Specialization	2016 - ?	Programmable logic, rapid ASICs, CGRAs
Approximate Computing	2020 - ?	Fuzzy computation, analog (NPU's)
Neural Computing	2025 - ?	Doing computation with neuromorphic hardware, as opposed to running neural code in software
Quantum Computing	2025 - ?	Next "digital" paradigm?
Biological Computing	2030 - ?	Interfacing to DNA, statistical computation
Ecological Computing	2040 - ?	Defining ecological outcomes, using ecologies for computation



Microsoft Research
Faculty
Summit
2014 15TH ANNUAL

Quantum Leaps in Computing

Krysta Svore
QuArC Group, Microsoft Research



Chinese Hackers Pursue Key Data on U.S. Workers



BITS BLOG
Adidas Joins Wearable Stampede With Fitness Tracker

BITS BLOG
IBM Wants to Invent the Chips of the Future, Not Make Them

BITS BLOG
Personal Computing's Big Three Get a Little Bigger

BITS BLOG
Microsoft Taking the Right to be Fi

AEROSPACE BIOMEDICAL

COMPUTING // HARDWARE

NEWS
Quantum Chip Help

Experimental chip does

EMAIL PRINT



Photo: Jonathan Matthews/University of
BY ANNE-MARIE CORLEY // SEPT
3 September 2009—Modern cryp
have in factoring huge numbers,
computer finds factors easily. Te

ALL NEW 2014 CTS SEDAN

DIVERSIFY YOUR PORTFOLIO.

LEARN MORE BUILD YOUR OWN

TECHNOLOGY

Microsoft Makes Bet Quantum Computing Is Next Breakthrough

By JOHN MARKOFF JUNE 23, 2014

- EMAIL
- FACEBOOK
- TWITTER
- SAVE
- MORE



Top Stories
This article and others like it are part of our new subscription.
Learn More »

SANTA BARBARA, Calif. — Modern computers are not unlike the looms of the industrial revolution: They follow programmed instructions to weave intricate patterns. With a loom, you see the result in a cloth or carpet. With a computer, you see it on an electronic display.


Now a group of physicists and computer scientists funded by Microsoft is trying to take the analogy of interwoven threads to what some believe will be the next great leap in computing, so-called quantum computing.

If the scientists are right, their research could lead to the design of computers that are far more powerful than today's supercomputers and could solve problems in fields as diverse as chemistry, material science, artificial intelligence and code-breaking.



Michael Freedman, Sankar Das Sarma and Chetan Nayak proposed a computing model in 2005 that can be used to construct qubits, the foundation of quantum computing.
Emily Berl for The New York Times

Quantum Computer



Shop
viron
b

...sited with NASA are forming a laboratory
y means of computers that use the unusual
Their quantum computer, which performs
s of times faster than existing
be in active use in the third quarter of this



The machine does not fit the conventional concept of a quantum computer, but makes use of quantum effects

Quantum “Requirements”

Quantum hardware architecture:

Architect a scalable, fault-tolerant, and **fully programmable quantum computer**

Quantum software architecture:


Program and **compile complex algorithms** into optimized, target-dependent (quantum and classical) instructions

Quantum algorithms:


Design real-world **quantum algorithms** for small-, medium- and large-scale quantum computers

Quantum Hardware Technologies


Ion traps




NV centers



Super-conductors




Quantum dots



Linear optics



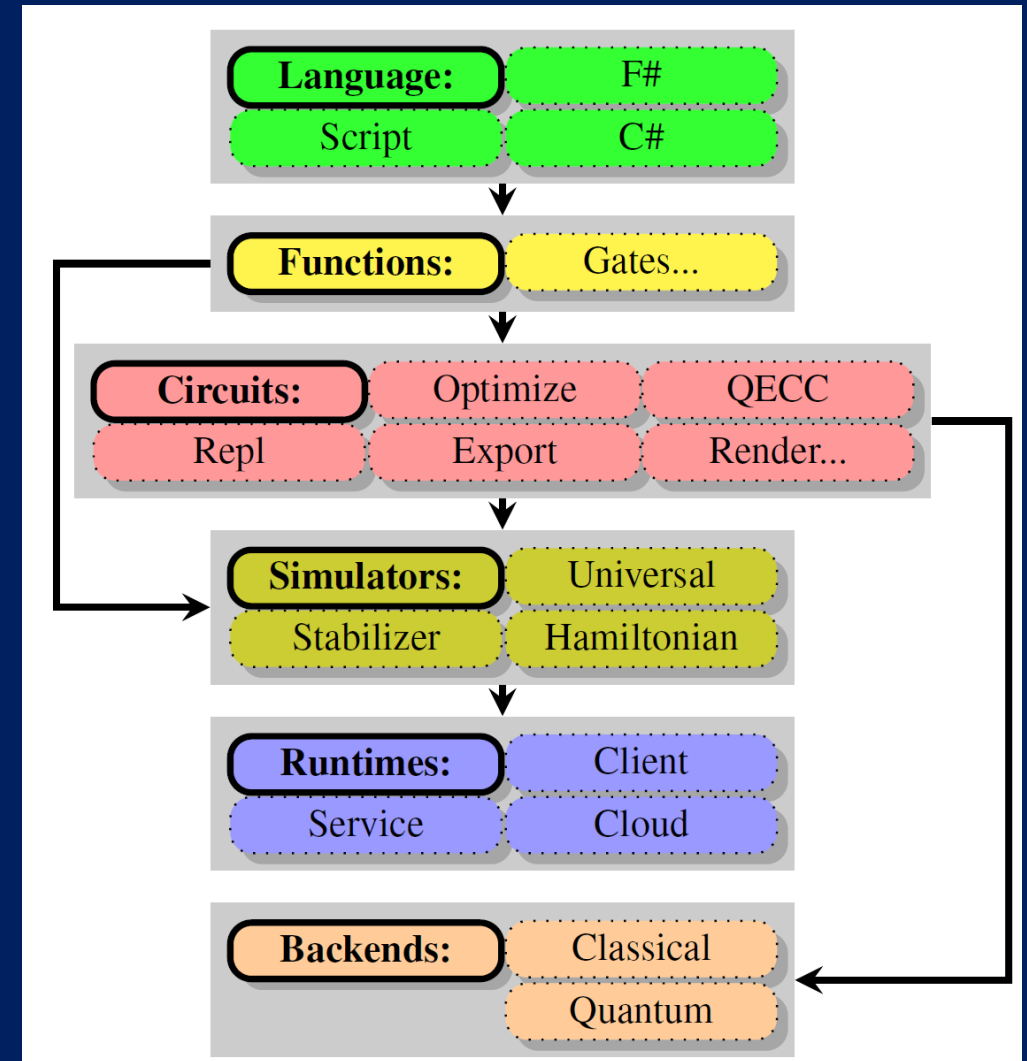
Topological



LIQ*U*i| \rangle : Quantum Software Architecture

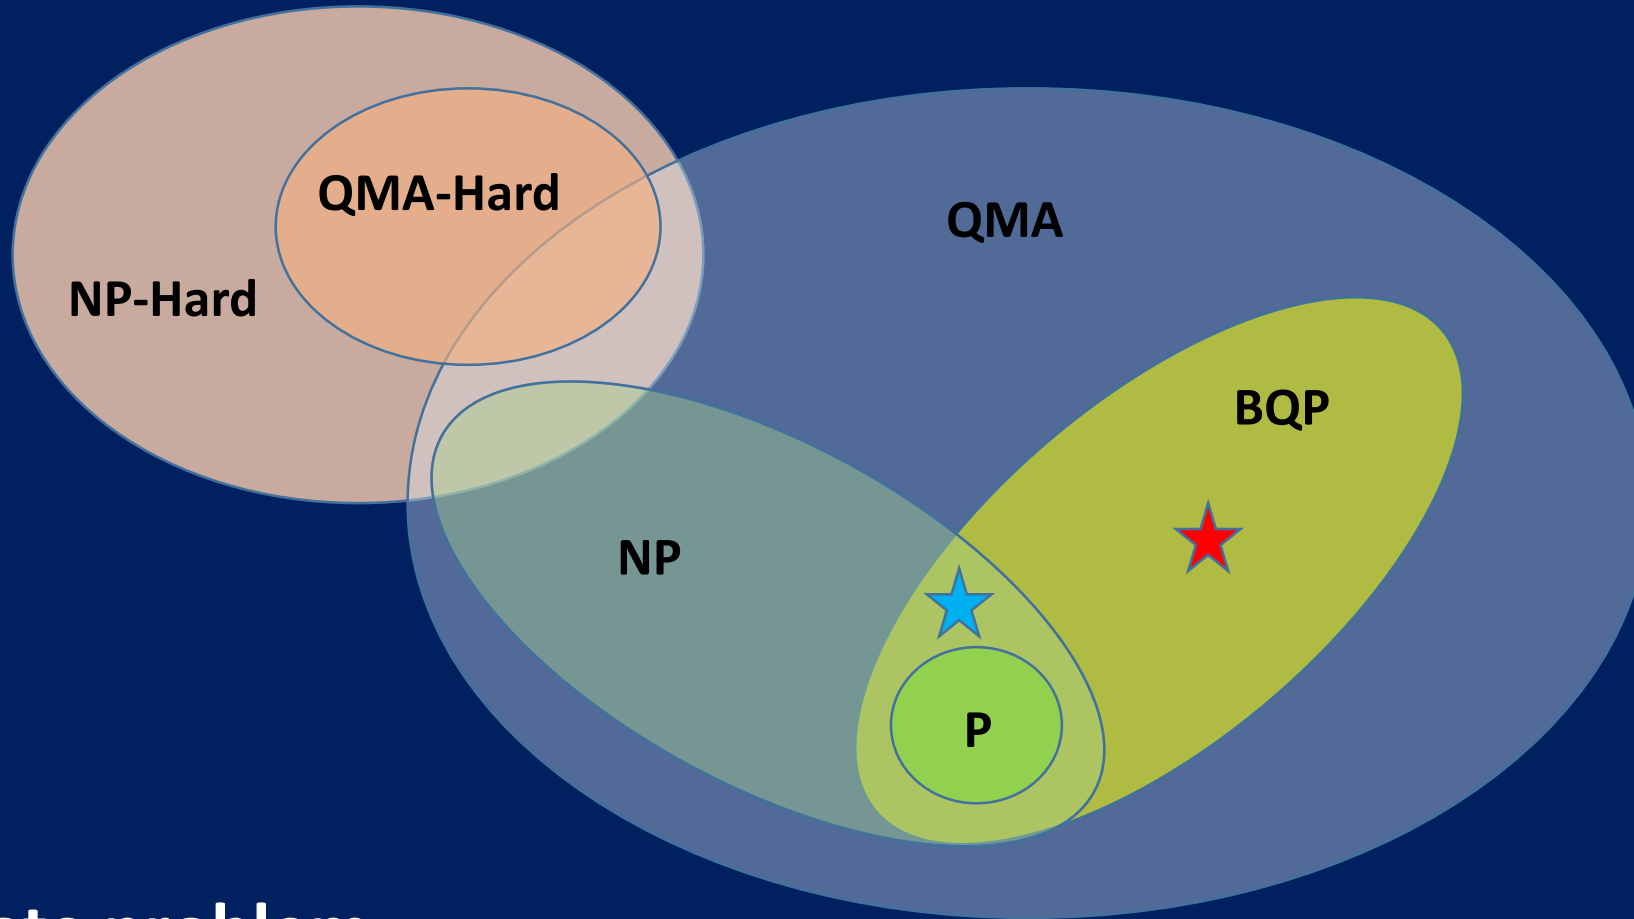
- Enables **easy programming** and simulation of complex quantum circuits
- Allows **retargeting of circuits** for various purposes: simulation, rendering, optimization, noise modeling, and export
- Provides **state-of-the-art quantum circuit simulation** tools

DEMO: Tuesday July 15, 2:15pm



The LIQ*U*i| \rangle platform
[Wecker, Svore, 2014]

Why Quantum Compute?



Ultimate problem:

Develop quantum algorithms whose complexity lies in $BQP \setminus P$

Quantum Algorithm “Wins”

Quantum simulation (1982)

- Simulate physical systems in a quantum mechanical device
- Exponential speedups



Solving Linear Systems of Equations (2010)

- Applications shown for electromagnetic wave scattering
- Exponential speedups



Shor's Algorithm (1994)

- Breaks RSA, elliptic curve signatures, DSA, El-Gamal
- Exponential speedups



Quantum Simulation for Quantum Chemistry

Ultimate problem:

Simulate molecular dynamics of *larger* systems or to *higher accuracy*

Want to solve system *exactly*

Current solution:

33% supercomputer usage dedicated to chemistry and materials modeling

Requires simulation of exponential-size Hilbert space


Limited to 50-70 spin-orbitals classically

Quantum solution:

Simulate molecular dynamics using *quantum simulation*

Scales to 100s spin-orbitals using only 100s qubits

Runtime recently reduced from $O(N^{11})$ to $O(N^4) - O(N^6)$



[Poulin et al., 2014]

Application: Nitrogen Fixation

Ultimate problem:

Find catalyst to convert nitrogen to ammonia at room temperature

Reduce energy for conversion of air to fertilizer

Current solution:


Uses Haber process developed in 1909

Requires high pressures and temperatures

Cost: 3-5% of the world's natural gas production (1-2% of the world's annual energy)

Quantum solution:

~ 100-200 qubits: Design the catalyst to enable inexpensive fertilizer production



Application: Carbon Capture

Ultimate problem:


- Find catalyst to extract carbon dioxide from atmosphere
- Reduce 80-90% of emitted carbon dioxide

Current solution:

- Capture at point sources
- Results in 21-90% increase in energy cost

Quantum solution:

- ~ 200-400 qubits: Design a catalyst to enable carbon dioxide extraction from air



Quantum Algorithm Opportunities


Quantum simulation

- Extend q. chem. method to solid state materials
- E.g., high temp. superconductivity
- ~ 2000 qubits; linear or quad. scaling



Machine learning

- Clustering, regression, classification
- Polynomial speedups to date
- Can we harness interference to produce better inference models?



Cryptography

- RSA, DSA, elliptic curve signatures, El-Gamal
- What questions should we pose to a quantum computer?



Tuesday July 15, 3:00 - 4:30PM, Cascade Room

<http://research.microsoft.com/groups/quarc/>

<http://research.microsoft.com/en-us/labs/stationq/>



ksvore@microsoft.com



Microsoft Research

Faculty Summit

2014 15TH ANNUAL



“Mobilizing” Healthcare

Desney Tan, Microsoft Research



Invisible Sensors

Track Real-Time Performance
Warn of Potential Hazards
Allow Behavior Optimization



Your vehicle's instrument panel is equipped with this cluster or one very similar to it. The U.S. Heavy-Duty Automatic Transmission version is shown here.

The instrument panel cluster includes these key features:

- A. Transmission Temperature Gauge (if equipped)
- B. Tachometer
- C. Speedometer, Odometer, Transmission Gear Selector
- D. Oil Pressure Gauge
- E. Fuel Gauge
- F. Voltmeter
- G. Engine Coolant Temperature Gauge
- H. Driver Information Center (DIC)

Warning and indicator lights			Mark	Name	Page	Mark	Name	Page
	Brake system warning light	3-19	CRUISE	Cruise control indicator light	3-24			
	Seatbelt warning light	3-13	SET	Cruise control set indicator light	3-24			
	Front passenger's seatbelt warning light	3-13		Low fuel warning light	3-20			
	SRS airbag system warning light	3-14		Low tire pressure warning light (U.S.-spec. models)	3-16			
	Front passenger's frontal airbag ON indicator	3-14		Shift-up indicator light (STI)	3-23			
	Front passenger's frontal airbag OFF indicator	3-14		Driver's control center differential auto indicator light (STI)	3-24			
	CHECK ENGINE warning light/ Malfunction indicator light	3-15		Auto (+) mode indicator light (STI)	3-24			
	Charge warning light	3-16		Auto (-) mode indicator light (STI)	3-25			
	Oil pressure warning light	3-16		Driver's control center differential indicator and warning lights (STI)	3-25			
	Rear differential oil temperature warning light (STI)	3-16		REV indicator light (STI)	3-25			
	ABS warning light	3-18		Sport (S) mode indicator light (STI)	3-23			
				Headlight indicator light	3-24			

On-Board Diagnostics (OBD)
displayed on Information Dashboard

A silhouette of a person in a yoga pose, with hands clasped above their head, set against a sunset over a body of water. The person is sitting on a mat, and the background shows a calm sea and distant mountains under a clear sky.

Help People Live Healthier,
for Longer...

...and Cheaper

US spends ~\$2.8 trillion/year on Healthcare
(5% GDP in 1960, 18% today, 37% by 2050)
 $\frac{3}{4}$ spent managing chronic conditions



We cannot improve what we cannot measure

Digitize → Compute → Detect, diagnose, manage, predict, prevent disease

Intuitive Medicine → Precision Medicine

Diagnoses: Loose subjective observation → Precise objective data

Treatment: Therapies with unclear efficacy → Evidence-based treatment



Continuously Capture Health-Relevant Data

Smaller, cheaper sensors/devices in convenient (desirable) form factors
Operate with scarce resource (power, network, interaction/attention)

Impossibly Small Devices (Mon 12:30pm Cascade)



Explosion of Wearable Biosensing Devices

But most fail to engage and deliver real improvement

Research must deliver Deeper Innovation

Treat body as prolific generator of Big Data

Discover (or invent) new metrics (i.e. the "New Vital Signs")

Create Actionable Insights

Interlink devices and aggregate data to facilitate discovery

Explore privacy, security, reliability

Design engaging experiences that motivate real change

Connected Devices (Mon 2:15pm Cascade)

Genomics Software Revolution (Mon 2:15pm Baker)


Making Sense of Billion Sensors (Tues 10:15am Kodiak)

Science in the Cloud (Tues 12:45pm Rainier)


etc...

We are critical to the Healthcare Revolution. Let's do it!

Predict
Diagnose
Treat
Monitor
Inform



Crowds are People Too!



Mary L. Gray, Senior Researcher, Microsoft Research, New England Lab
Associate Professor@The Media School, Indiana University

@maryLgray

Crowds are Important

We rely on crowds to get work done

- Human computation (*VizWiz*)
- Citizen Science (*Eyewire*)

As crowdsourcing advances automation/AI...

- Crowdwork *could* (should and must) offer opportunities for the future of employment



How Do We Study Crowds?

Crowdsourcing is a complex technological AND social system that combines:

- Human labor
- Social connections
- Economic transactions
- Myriad international laws
- Distributed computation



What Tools Do We Need?

- Computer Science maps the *nodes* (people) and *edges* (relationships) in a network
- Ethnography studies the *variety of nodes* (individuals, institutions) and *meaning of edges* (motivations, hierarchies, power dynamics)
 - **in context** and **over time**
(spoiler alert! I'll come back to why these matter)

Requires an Interdisciplinary Team, too!



Siddharth Suri
CS (co-PI)



Sara Kingsley
Labor Studies
/Econometrics



Shoaib Ali
Development Studies



Deepti Desai
Sociology



Gregory Minton
Mathematics



Kate Miltner
Media Studies



Rajesh Patel
Engineer/Systems Builder



Research Methodology: Integrating 5 data sources

Ethnographic fieldwork

- 8 months of in India and U.S.
 - (July 2013- March 2014)
- 46 of 116 interviews completed in India
- 75 interviews scheduled for the U.S.

Ongoing Survey work (goal ~3,000 total)

- Amazon.com's Mechanical Turk (AMT)
- Microsoft's Universal Human Relevance System (UHRS)
- MobileWorks, a startup with a social and entrepreneurial mission
- Amara.org, non-profit, peer-driven translation service

Metadata

- Workflows on AMT, UHRS, and MobileWorks, Amara.org

System-level measurements

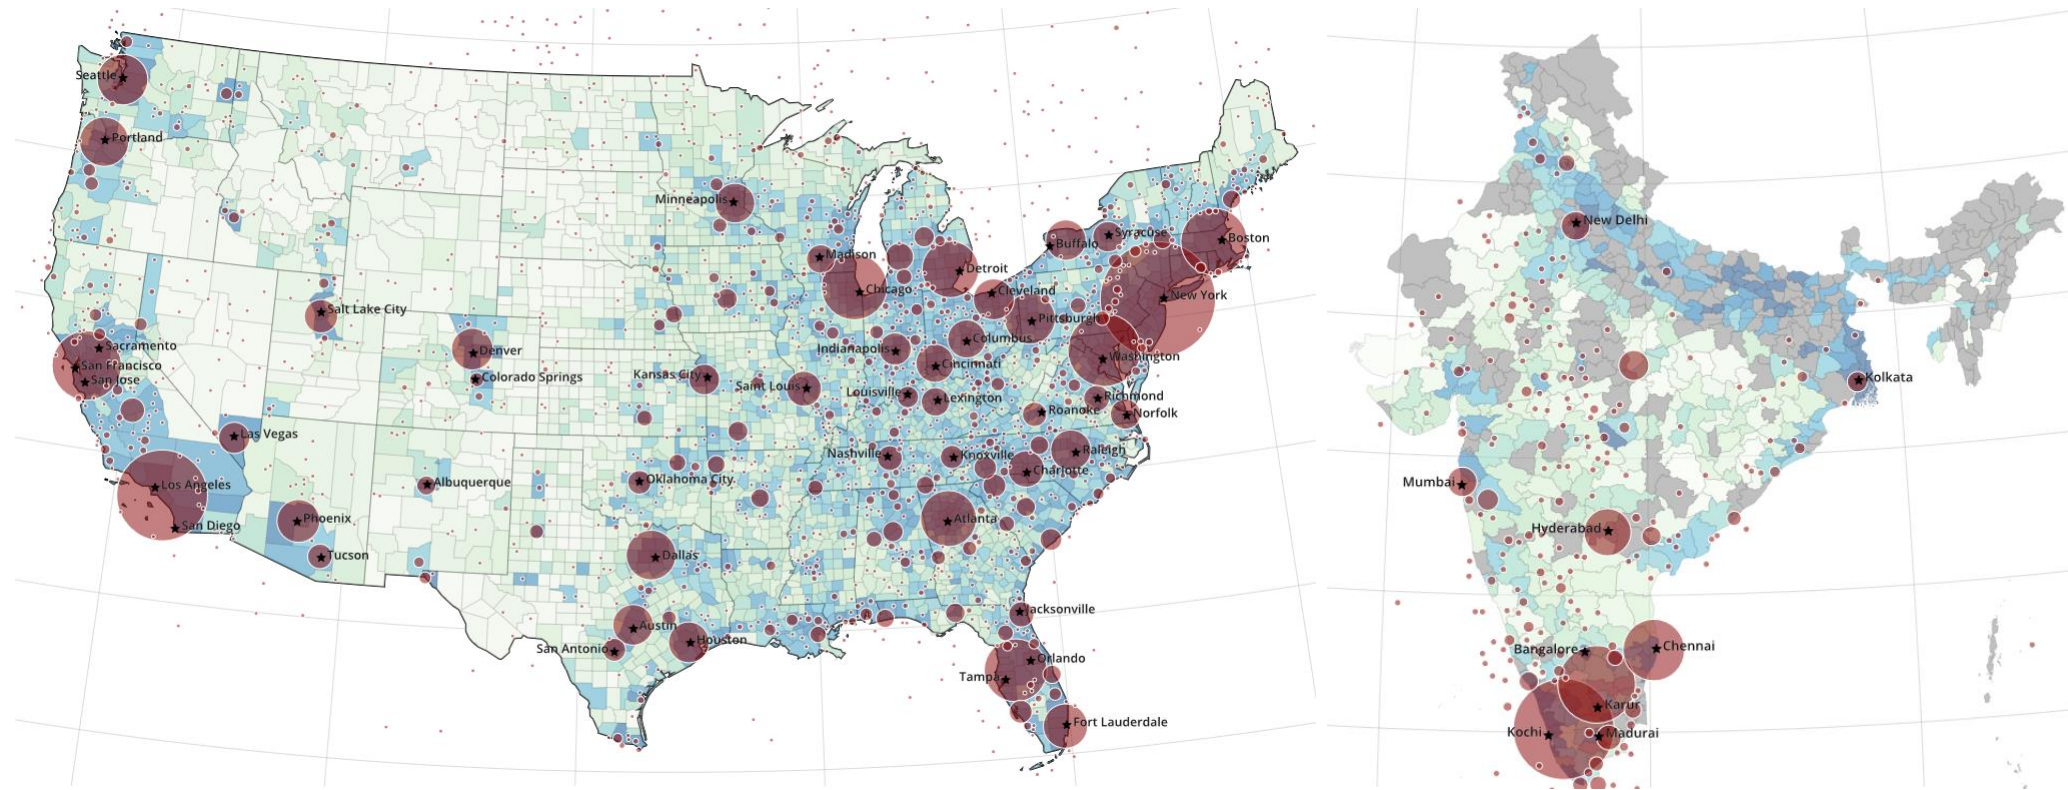
- Geolocation tasks
- Sampling of worker purchases
- Recruitment

Textual analysis

- Worker discussion forums
- Industry rhetoric
- Related industries (piecework, temp work, BPOs)



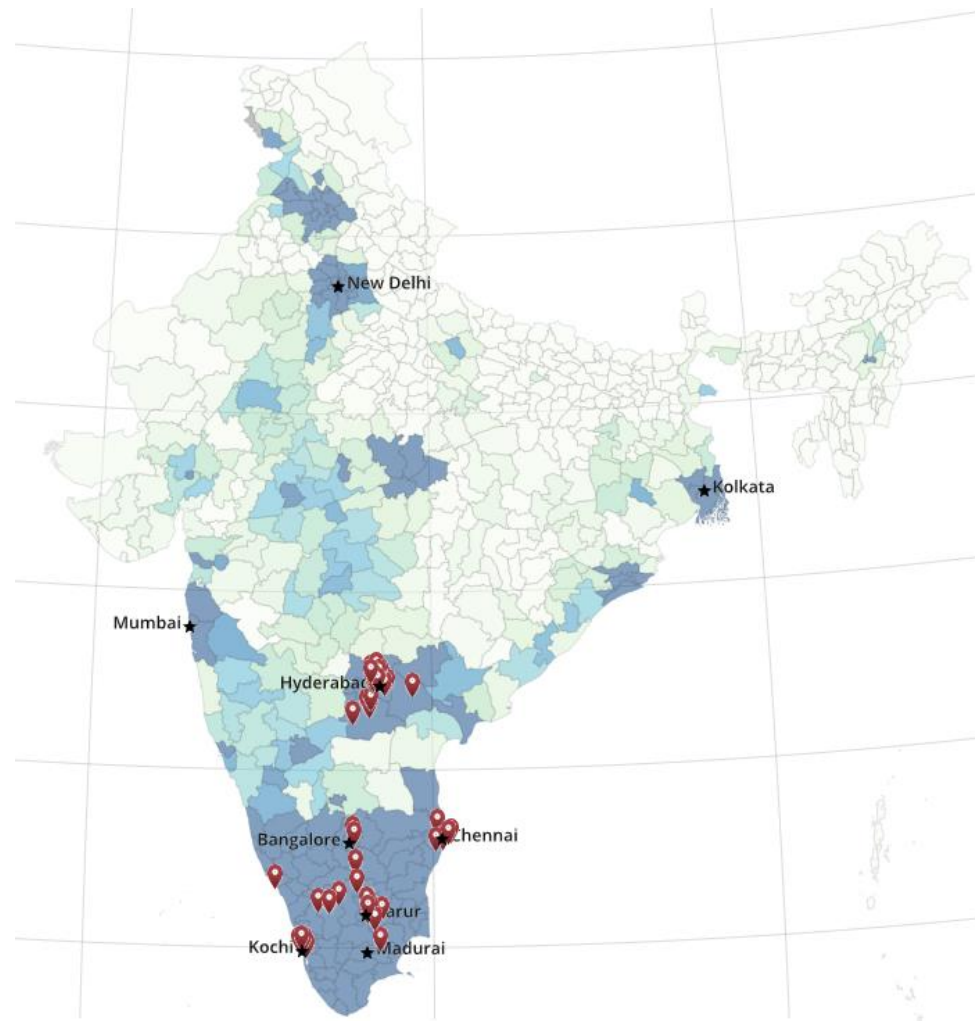
Dataset: Mapping the Crowd



Self-reported locations for about 10,000 participants in a map task (HIT) on AMT. Coloration of counties/districts is by *population density*.



Dataset: Interview Locations



Locations of the 48 interviews and 8 months of fieldwork conducted to date.

Coloration is by estimated density of workers on AMT. This estimate is based on our map task (HIT).

Crowdworkers Collaborate

People connect outside the systems that we can effectively monitor from afar

- Talk via phone
- post on message boards
- help each other sign up
- collect each other's paychecks
- work in family units

“Crowds” are (also) networks of individuals interacting in particular places (social contexts)



2 Key Challenges to Studying Crowds as People

Social media research (tends to) zoom too far out

- Treat crowds as big lumps of aggregated human behavior
- Can't see crowds as people, **interacting** in complicated ways, both on and offline, over time

Uncharted ethics of studying social media as human subjects research

- Supporting women working in the home w/o family's consent
- Helping workers when they lose their accounts
- Protecting respondents' privacy by separating metadata and survey data



Ethnographic Approach and CS Can Help Each Other (A LOT)!

**Must understand people's social lives and identities
to build better tools for them**

To build better tools for the future of work, technologies must learn to anticipate that:

- People will (always) do unexpected, novel things with technologies
- Our social identities—friend, family member, citizen—affect our work lives
- Social contexts+identities **significantly** shape success of technologies (i.e., Muslim workers offset discrimination in the workforce; women able to contribute from the home)

Direct, sustained ethnographic interaction w/ people offers CS 3 things:

1. Detailed attention to social dimensions of everyday life that likely matter most (income, employment, schooling, gender, religion, neighborhood dynamics, etc.)
2. Analysis of social connections that we can't see in the system
3. Sense of how social needs (i.e., employment) co-evolve, w/ tech, over time



Thanks!

Crowdworkers everywhere

Team @MSRcrowdstudy

research.microsoft.com/crowdwork





Microsoft Research

Faculty
Summit

2014 15TH ANNUAL