Online Aggregation

10-Year Feedback

Joseph M. Hellerstein (Berkeley)
Peter J. Haas (IBM)
Helen J. Wang (Berkeley → MSR)

A Demo is Worth a Thousand Words

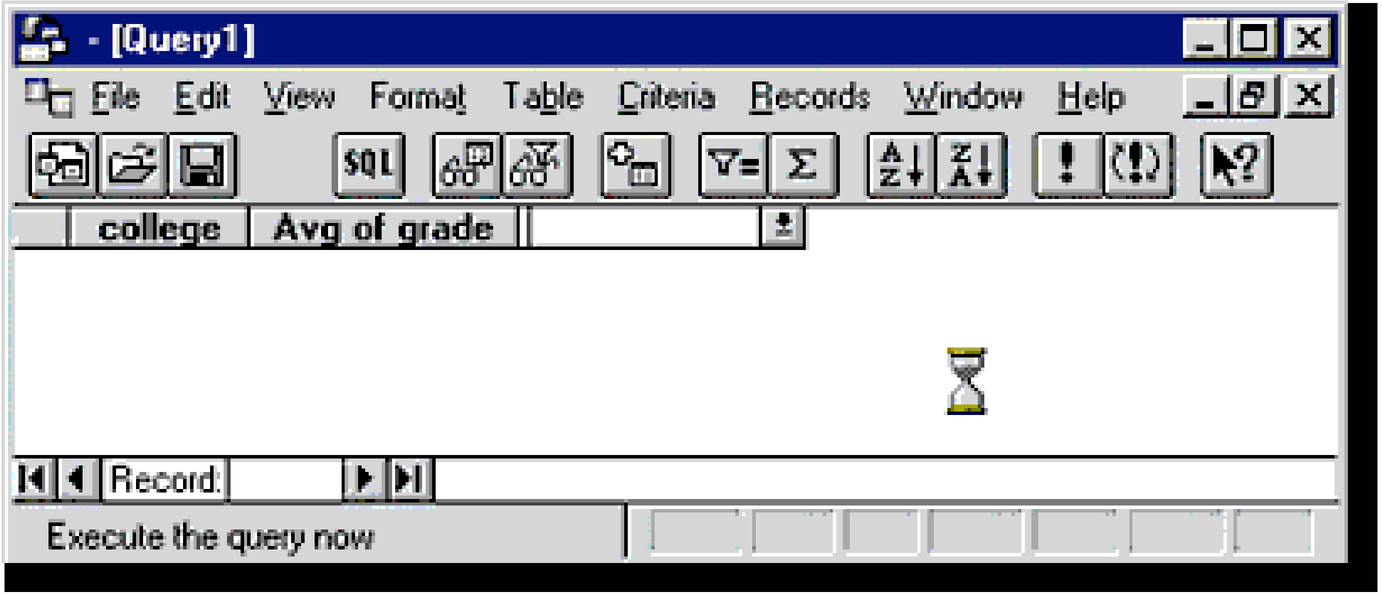

Execute the query now

A Demo is Worth a Thousand Words

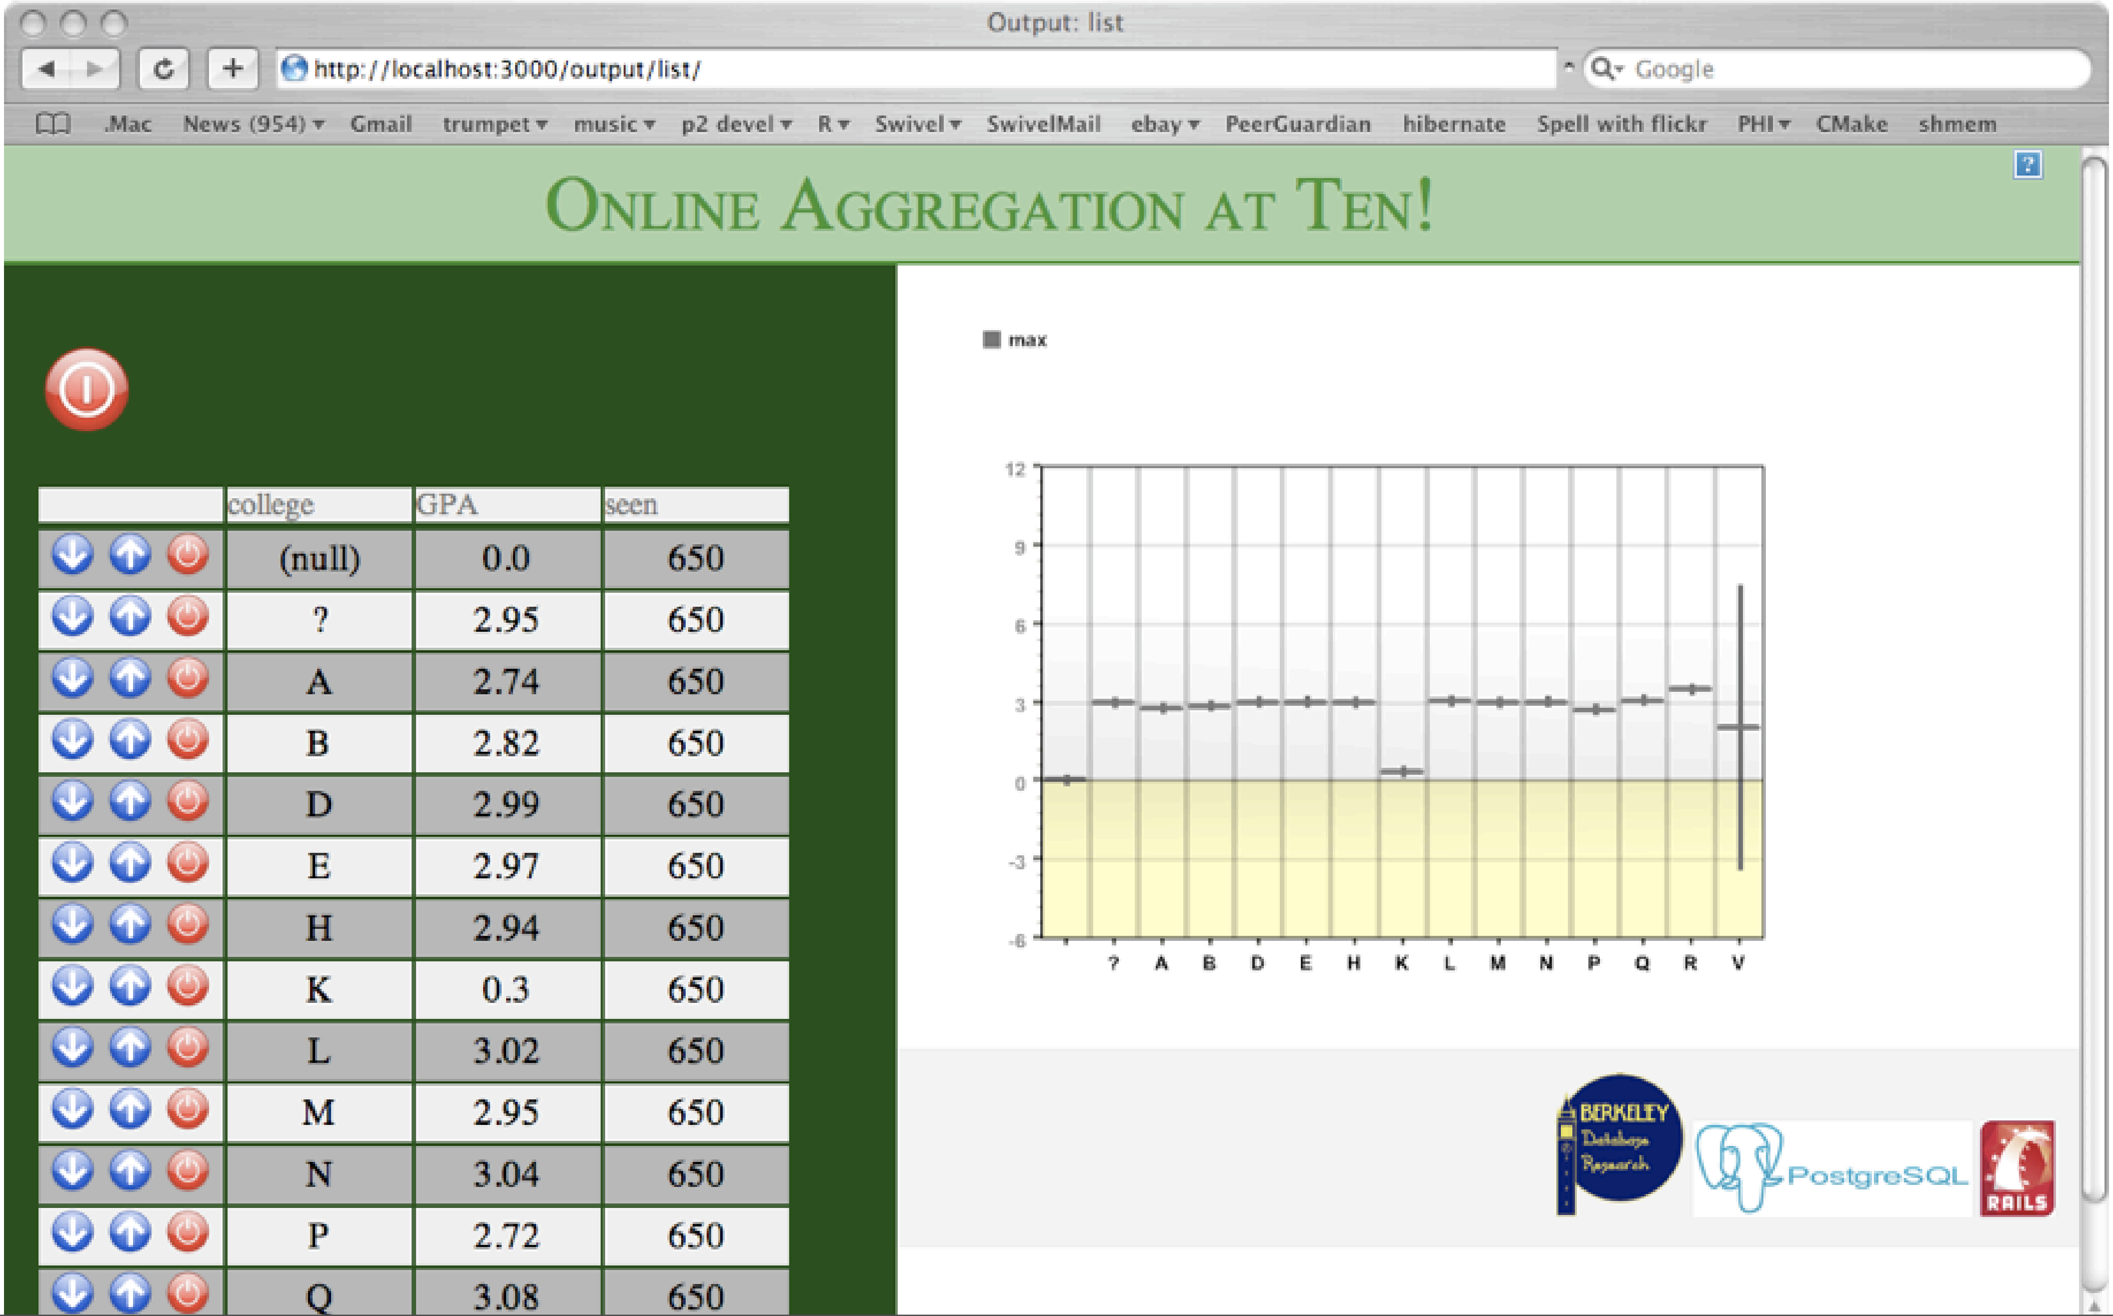

Online Aggregation at Ten!

	college	GPA	seen
⬇ ⬆ ⏻	(null)	0.0	650
⬇ ⬆ ⏻	?	2.95	650
⬇ ⬆ ⏻	A	2.74	650
⬇ ⬆ ⏻	B	2.82	650
⬇ ⬆ ⏻	D	2.99	650
⬇ ⬆ ⏻	E	2.97	650
⬇ ⬆ ⏻	H	2.94	650
⬇ ⬆ ⏻	K	0.3	650
⬇ ⬆ ⏻	L	3.02	650
⬇ ⬆ ⏻	M	2.95	650
⬇ ⬆ ⏻	N	3.04	650
⬇ ⬆ ⏻	P	2.72	650
⬇ ⬆ ⏻	Q	3.08	650

The OA Back-Story

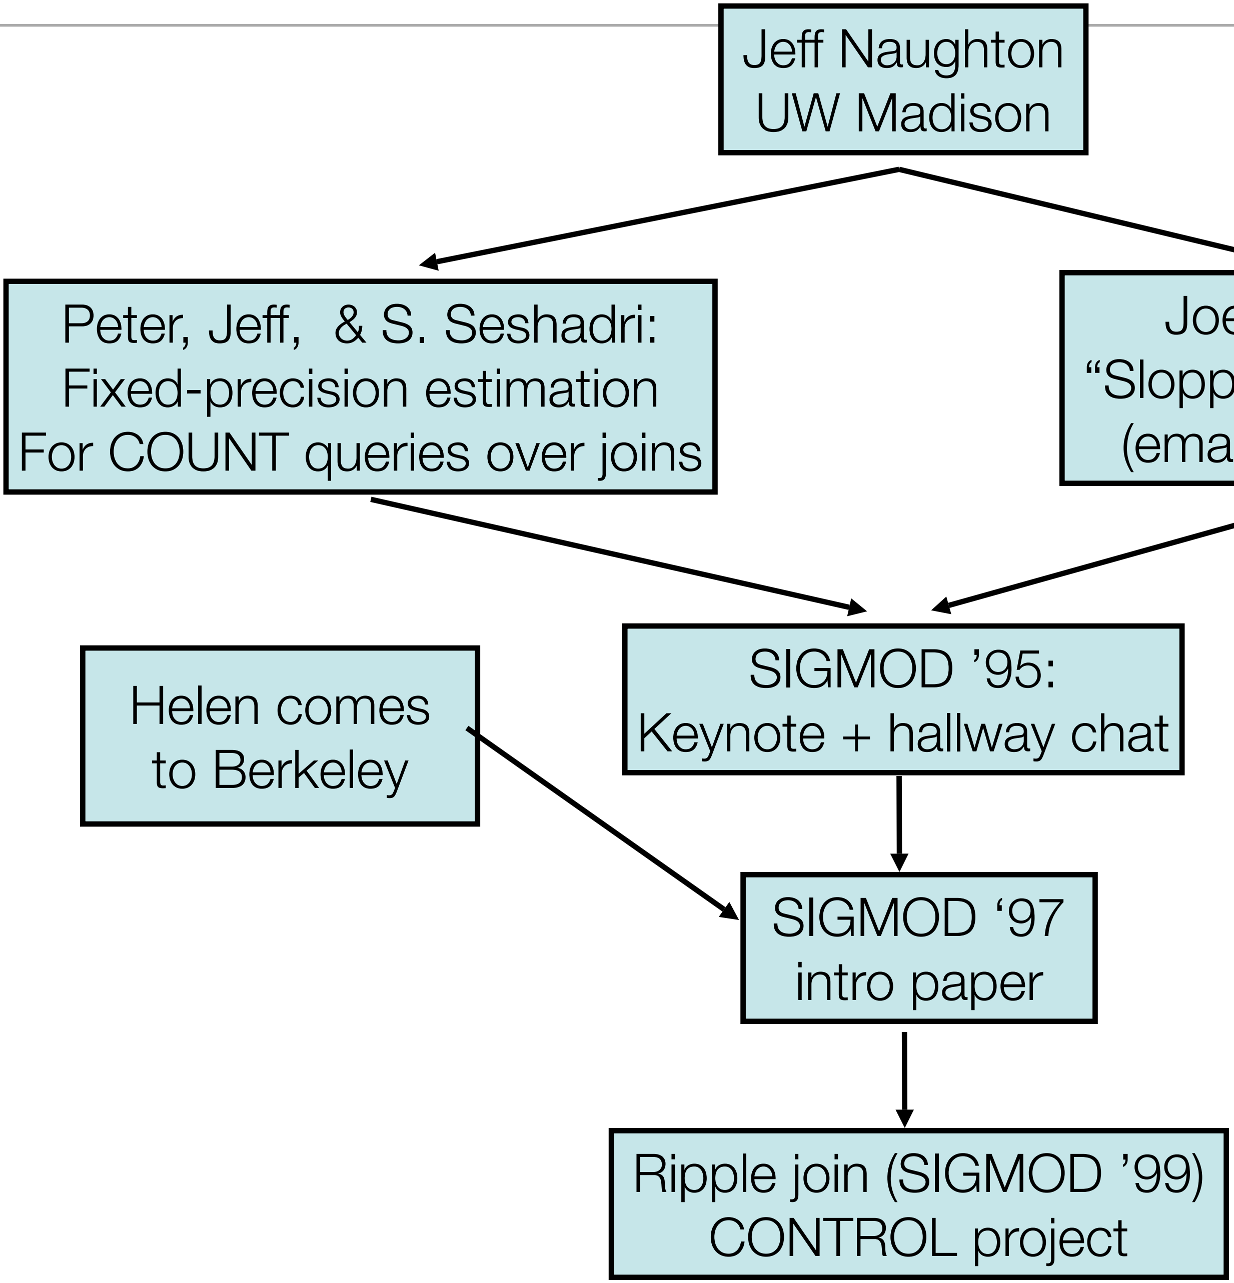

Jeff Naughton
UW Madison

Peter, Jeff, & S. Seshadri:
Fixed-precision estimation
For COUNT queries over joins

Joe and Jeff:
“Sloppy Databases”
(email exchange)

Helen comes
to Berkeley

SIGMOD ’95:
Keynote + hallway chat

SIGMOD ‘97
intro paper

Ripple join (SIGMOD ’99)
CONTROL project

A Three-Way Synthesis

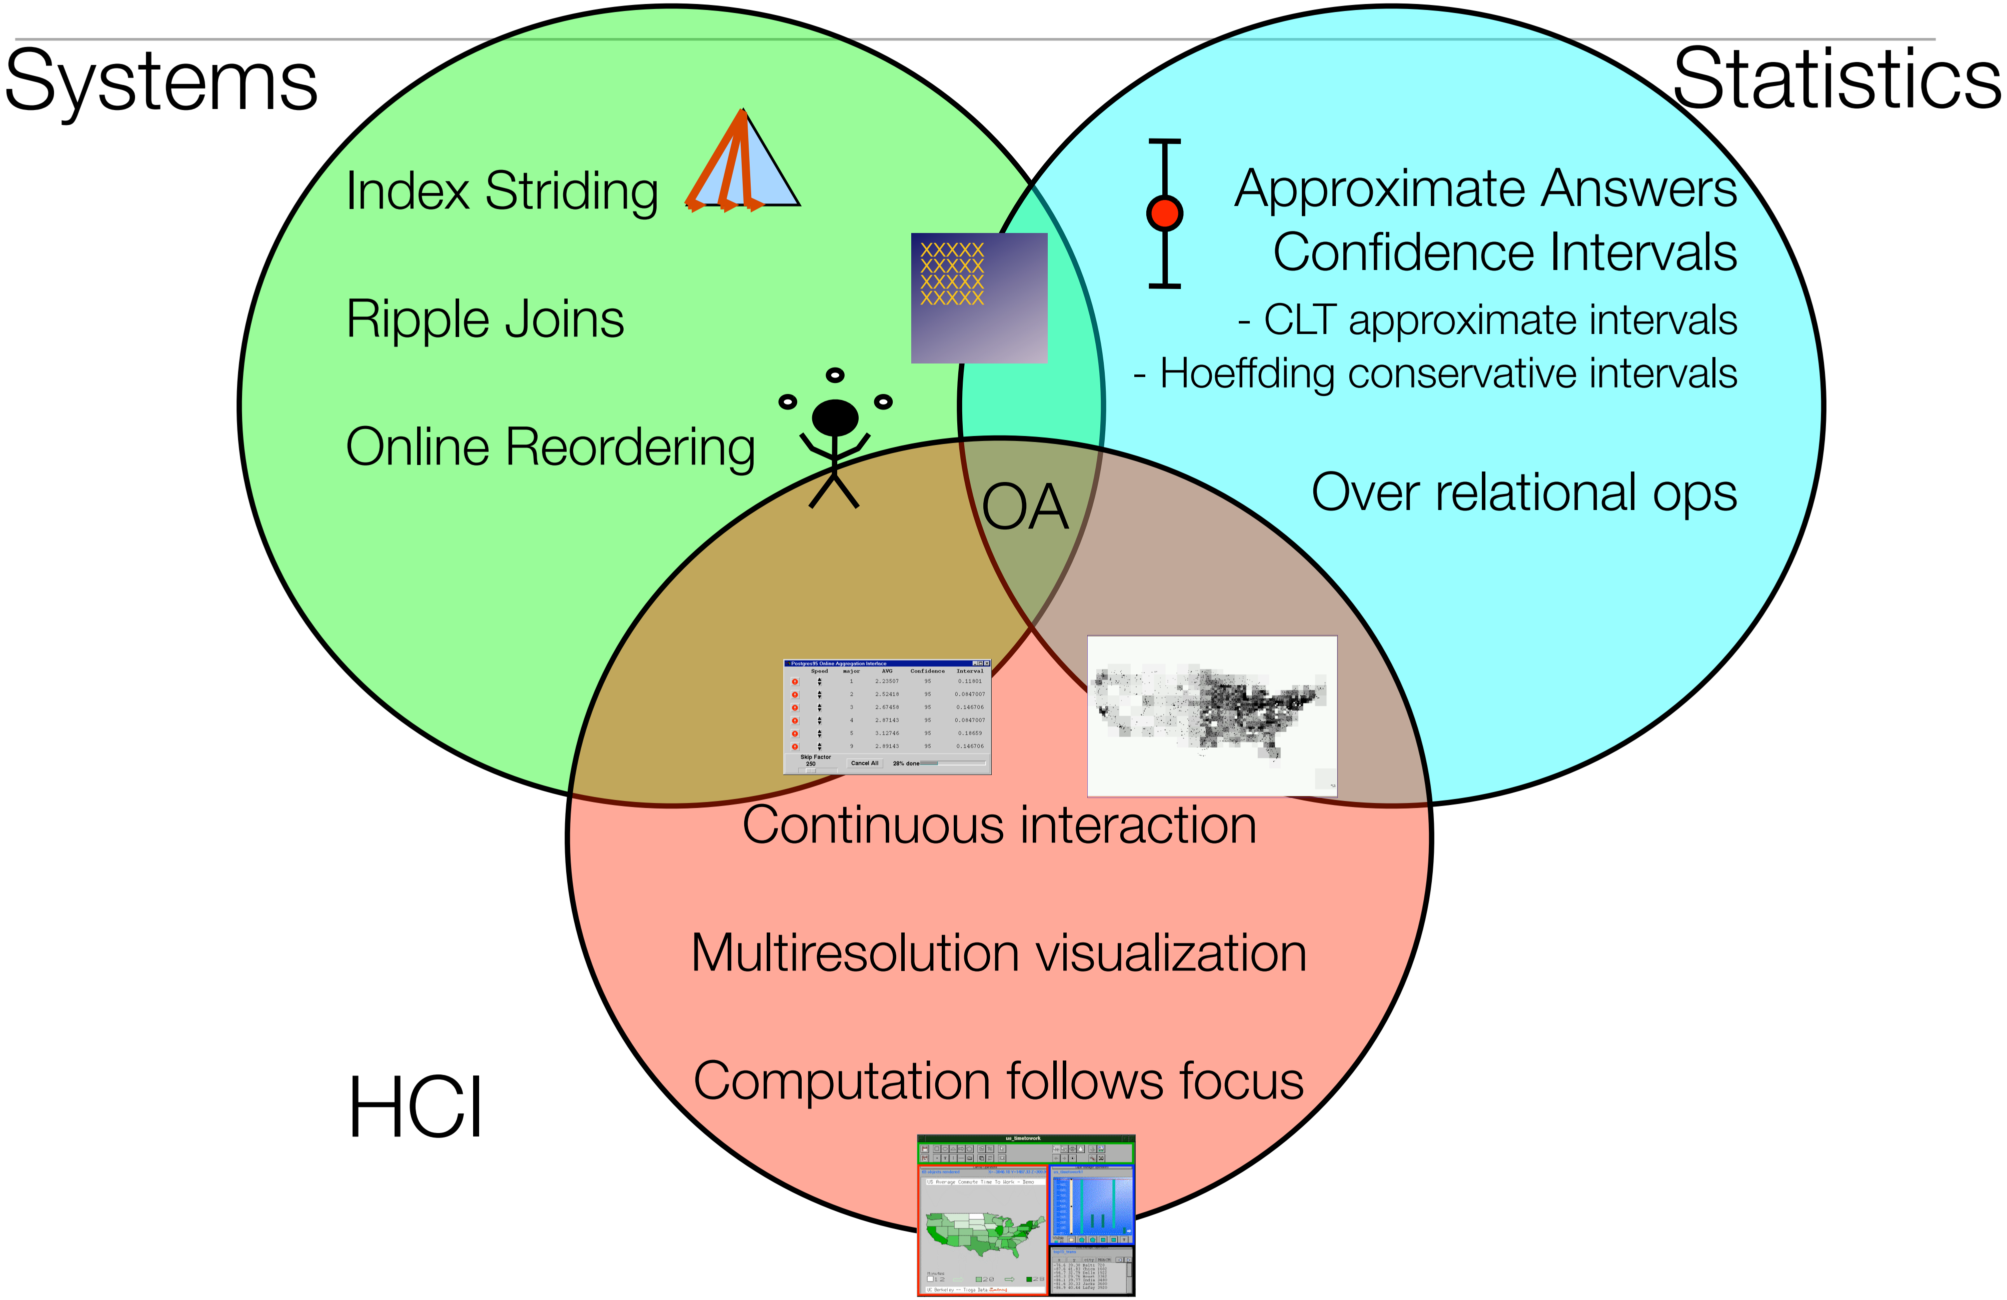

Systems

Statistics

HCI

Index Striding

Ripple Joins

Online Reordering

Approximate Answers

Confidence Intervals

- CLT approximate intervals
- Hoeffding conservative intervals

Over relational ops

OA

Continuous interaction

Multiresolution visualization

Computation follows focus

Postgres95 Online Aggregation Interface

	Speed	major	AVG	Confidence	Interval
		1	2.23507	95	0.11801
		2	2.52418	95	0.0847007
		3	2.67458	95	0.146706
		4	2.87143	95	0.0847007
		5	3.12746	95	0.18659
		9	2.89143	95	0.146706

Skip Factor 250

Cancel All

28% done

Implementations

- Postgres
  - Index stride, ripple join in engine
  - Simple Tcl/Tk interface

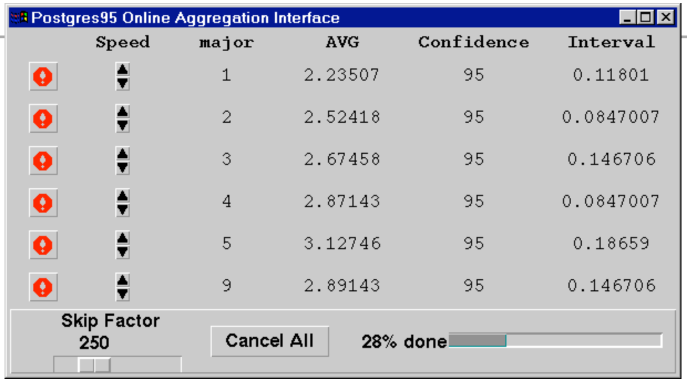

Postgres95 Online Aggregation Interface

Speed	major	AVG	Confidence	Interval
	1	2.23507	95	0.11801
	2	2.52418	95	0.0847007
	3	2.67458	95	0.146706
	4	2.87143	95	0.0847007
	5	3.12746	95	0.18659
	9	2.89143	95	0.146706

Skip Factor
250

Cancel All

28% done

- DB2
  - JDBC client application (differential pacing)
  - Java swing interface

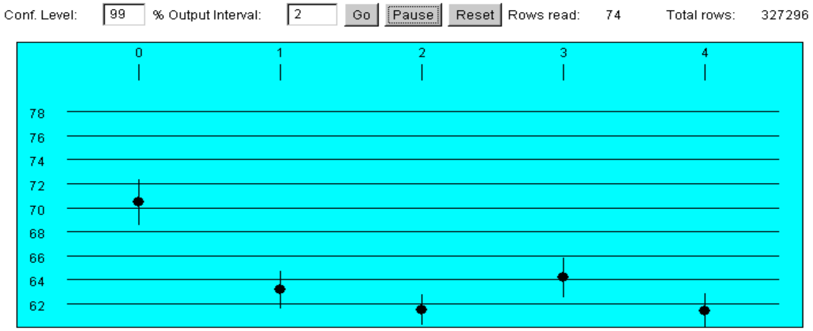

Conf. Level: 99 % Output Interval: 2 Go Pause Reset Rows read: 74 Total rows: 327296

- Informix Universal Server prototype
  - Index stride, ripple join, online reordering
  - Integrated with Metacube OLAP tool
  - Visual-basic + Excel interface

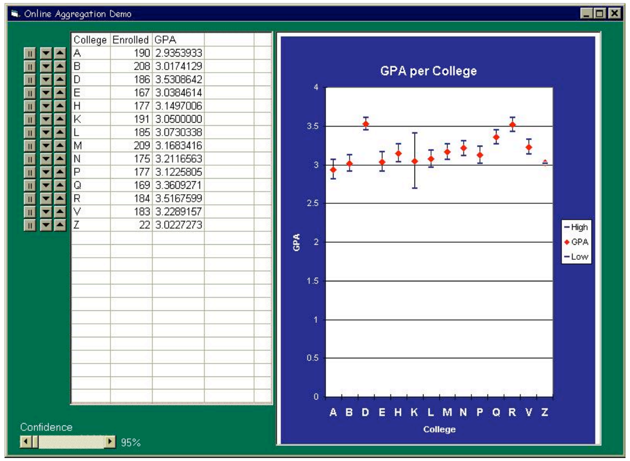

Online Aggregation Demo

College	Enrolled	GPA
A	190	2.9353933
B	208	3.0174129
D	186	3.5308642
E	167	3.0384614
H	177	3.1497006
K	191	3.0500000
L	185	3.0730338
M	209	3.1683416
N	175	3.2116563
P	177	3.1225805
Q	169	3.3609271
R	184	3.5167599
V	183	3.2289157
Z	22	3.0227273

Confidence 95%

Before

Systems

Statistics

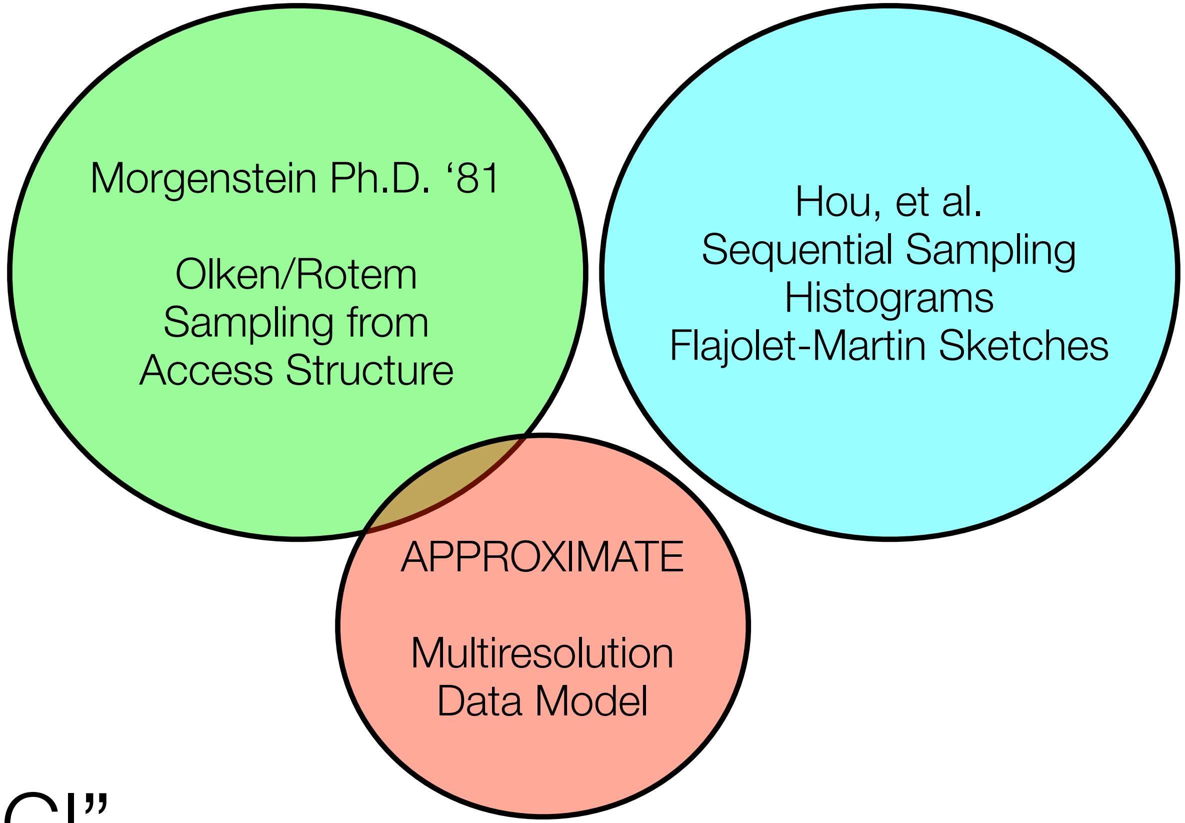

Morgenstein Ph.D. '81

Olken/Rotem
Sampling from
Access Structure

Hou, et al.
Sequential Sampling
Histograms
Flajolet-Martin Sketches

APPROXIMATE

Multiresolution
Data Model

"HCI"

The Last Decade

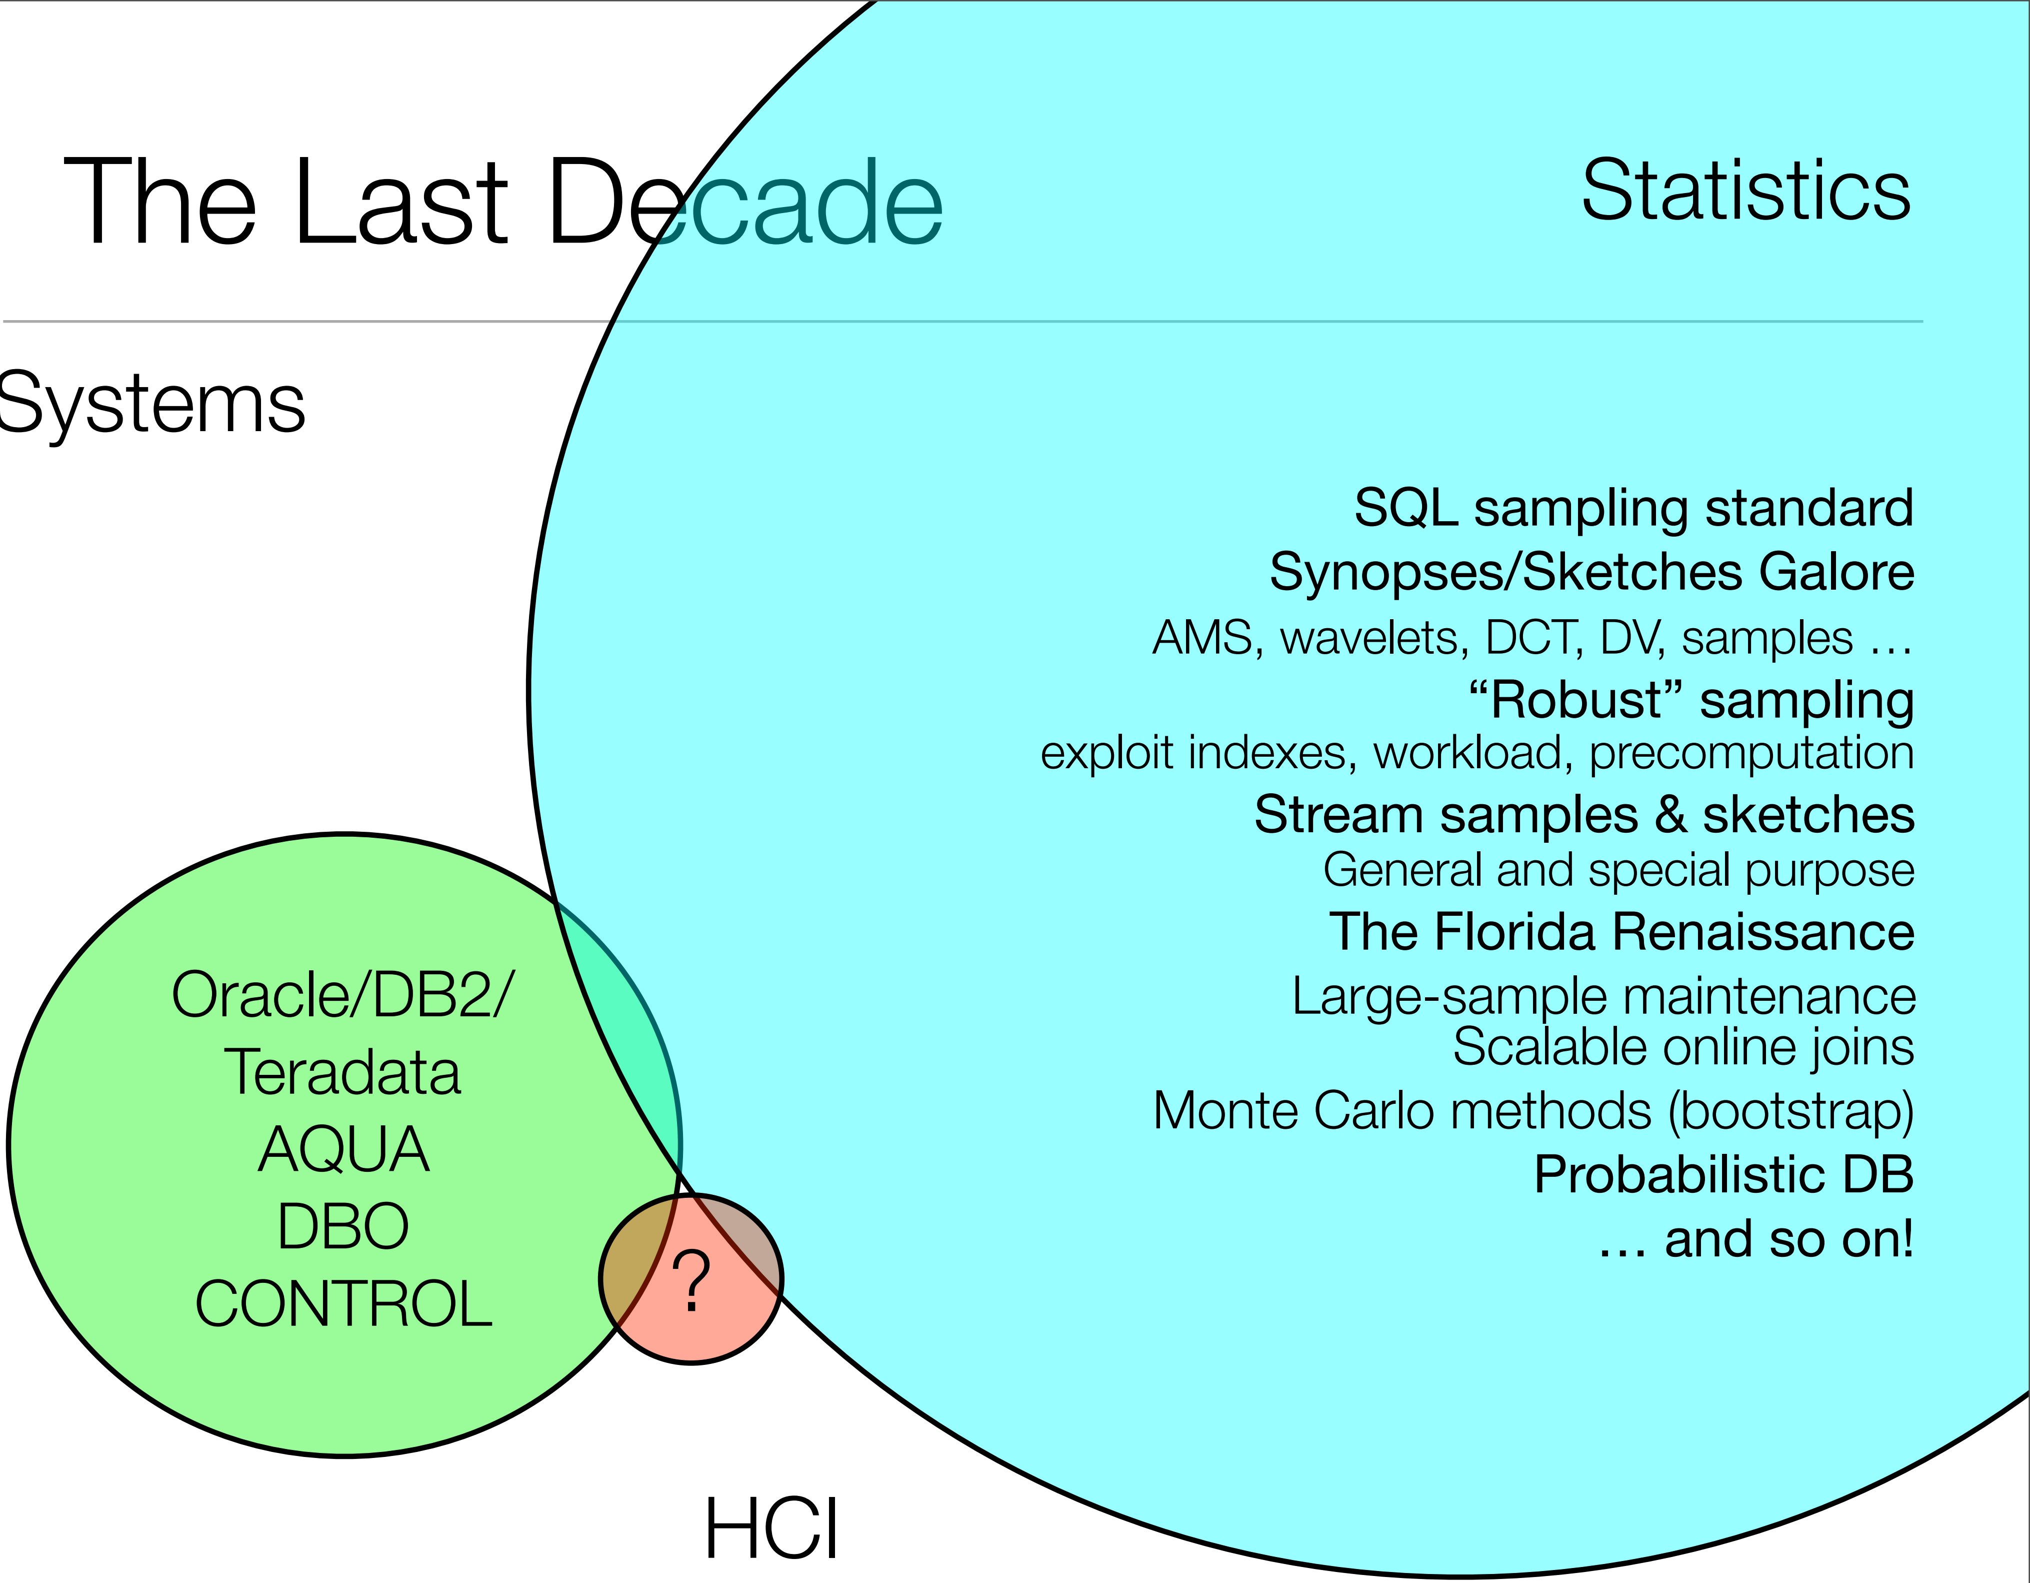

Statistics

Systems

SQL sampling standard

Synopses/Sketches Galore

AMS, wavelets, DCT, DV, samples …

“Robust” sampling

exploit indexes, workload, precomputation

Stream samples & sketches

General and special purpose

The Florida Renaissance

Large-sample maintenance

Scalable online joins

Monte Carlo methods (bootstrap)

Probabilistic DB

… and so on!

Oracle/DB2/
Teradata
AQUA
DBO
CONTROL

?

HCI

The CONTROL Project

- Online Aggregation
  - This paper, Ripple Joins, new confidence intervals
- Clouds
- CARMA
- Online “Enumeration”
  - Scalable Spreadsh
  - Partial Query Resul

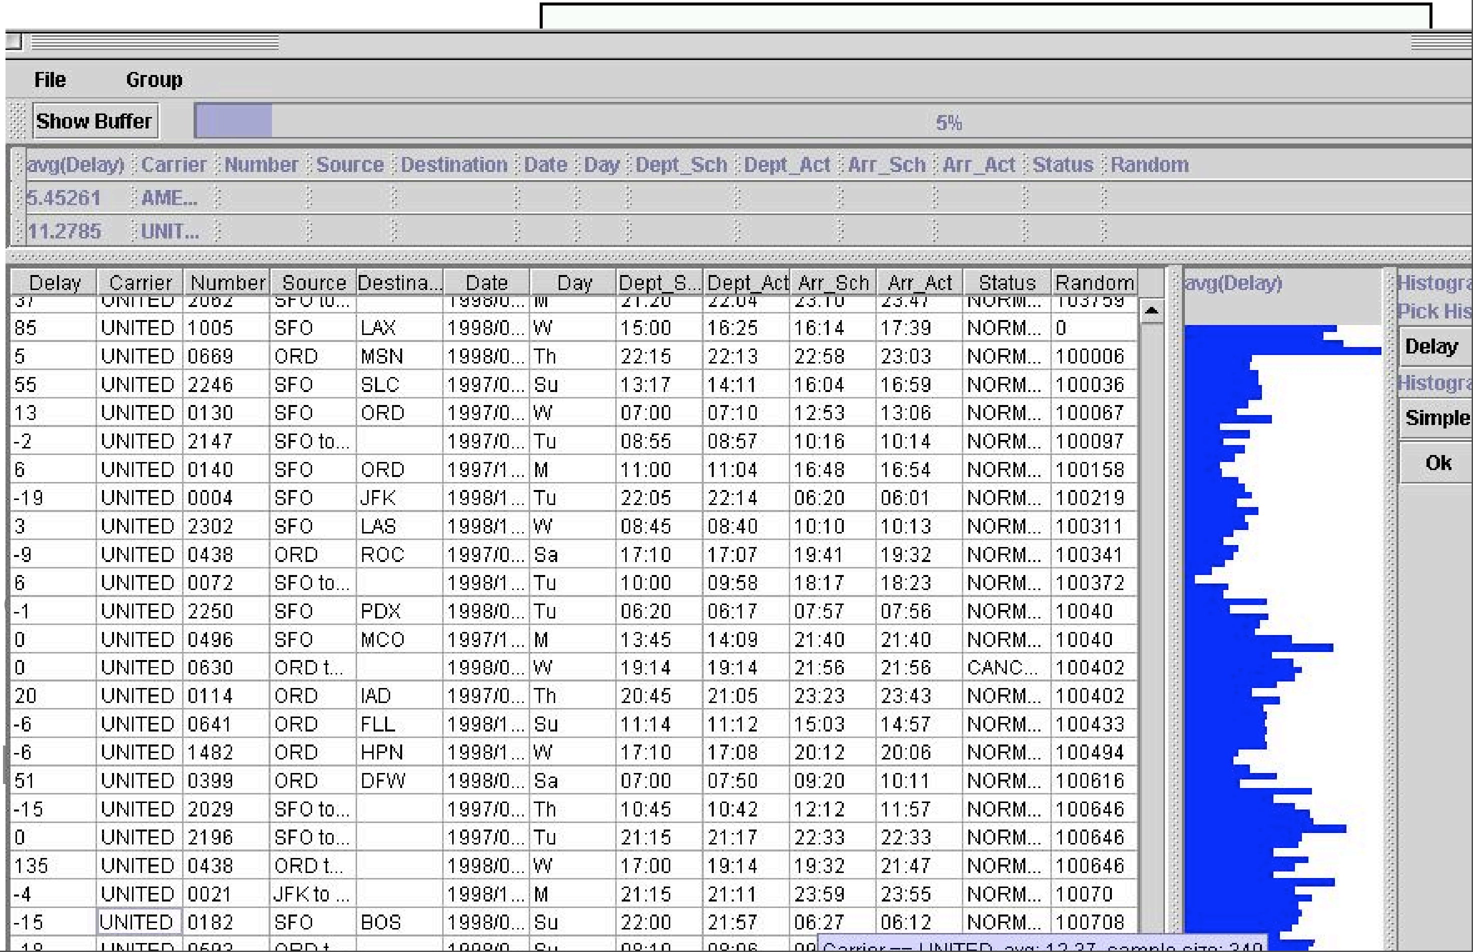

avg(Delay)	Carrier	Number	Source	Destination	Date	Day	Dept_Sch	Dept_Act	Arr_Sch	Arr_Act	Status	Random
5.45261	AME...											
11.2785	UNIT...											

Delay	Carrier	Number	Source	Destina...	Date	Day	Dept_S...	Dept_Act	Arr_Sch	Arr_Act	Status	Random
37	UNITED	2062	SFO to...		1998/0...	M	21:20	22:04	23:10	23:47	NORM...	103759
85	UNITED	1005	SFO	LAX	1998/0...	W	15:00	16:25	16:14	17:39	NORM...	0
5	UNITED	0669	ORD	MSN	1998/0...	Th	22:15	22:13	22:58	23:03	NORM...	100006
55	UNITED	2246	SFO	SLC	1997/0...	Su	13:17	14:11	16:04	16:59	NORM...	100036
13	UNITED	0130	SFO	ORD	1997/0...	W	07:00	07:10	12:53	13:06	NORM...	100067
-2	UNITED	2147	SFO to...		1997/0...	Tu	08:55	08:57	10:16	10:14	NORM...	100097
6	UNITED	0140	SFO	ORD	1997/1...	M	11:00	11:04	16:48	16:54	NORM...	100158
-19	UNITED	0004	SFO	JFK	1998/1...	Tu	22:05	22:14	06:20	06:01	NORM...	100219
3	UNITED	2302	SFO	LAS	1998/1...	W	08:45	08:40	10:10	10:13	NORM...	100311
-9	UNITED	0438	ORD	ROC	1997/0...	Sa	17:10	17:07	19:41	19:32	NORM...	100341
6	UNITED	0072	SFO to...		1998/1...	Tu	10:00	09:58	18:17	18:23	NORM...	100372
-1	UNITED	2250	SFO	PDX	1998/0...	Tu	06:20	06:17	07:57	07:56	NORM...	10040
0	UNITED	0496	SFO	MCO	1997/1...	M	13:45	14:09	21:40	21:40	NORM...	10040
0	UNITED	0630	ORD t...		1998/0...	W	19:14	19:14	21:56	21:56	CANC...	100402
20	UNITED	0114	ORD	IAD	1997/0...	Th	20:45	21:05	23:23	23:43	NORM...	100402
-6	UNITED	0641	ORD	FLL	1998/1...	Su	11:14	11:12	15:03	14:57	NORM...	100433
-6	UNITED	1482	ORD	HPN	1998/1...	W	17:10	17:08	20:12	20:06	NORM...	100494
51	UNITED	0399	ORD	DFW	1998/0...	Sa	07:00	07:50	09:20	10:11	NORM...	100616
-15	UNITED	2029	SFO to...		1997/0...	Th	10:45	10:42	12:12	11:57	NORM...	100646
0	UNITED	2196	SFO to...		1997/0...	Tu	21:15	21:17	22:33	22:33	NORM...	100646
135	UNITED	0438	ORD t...		1998/0...	W	17:00	19:14	19:32	21:47	NORM...	100646
-4	UNITED	0021	JFK to ...		1998/1...	M	21:15	21:11	23:59	23:55	NORM...	10070
-15	UNITED	0182	SFO	BOS	1998/0...	Su	22:00	21:57	06:27	06:12	NORM...	100708

Carrier == UNITED, avg: 12.37, sample size: 240

If you’re so smart, why ain’t you rich?

- Marketplace Challenge
  - Apps + Engine
  - Customer aversion to statistics (see no evil)
  - OLAP postponed this by a decade
  - If you want to rewrite the DB2 engine, you better have a VERY good business case
- Technical Challenge
  - How many DBMS engineers does it take to screw in a confidence interval?

Why the Renaissance?

- The rise of stat in CS
  - KDD, ML, WebSearch
- Market forces
  - DBMS Market Consolidation
  - Software Appliances/Services
  - New “niche” opportunities
- Web expectations
  - Speed, data-rich, rough-and-ready answers
  - App/Engine integration not a barrier

What Next?

- Clearly, tons of community energy on approximation
    - Especially algorithmics
- The fun (for us) is in the integrative work

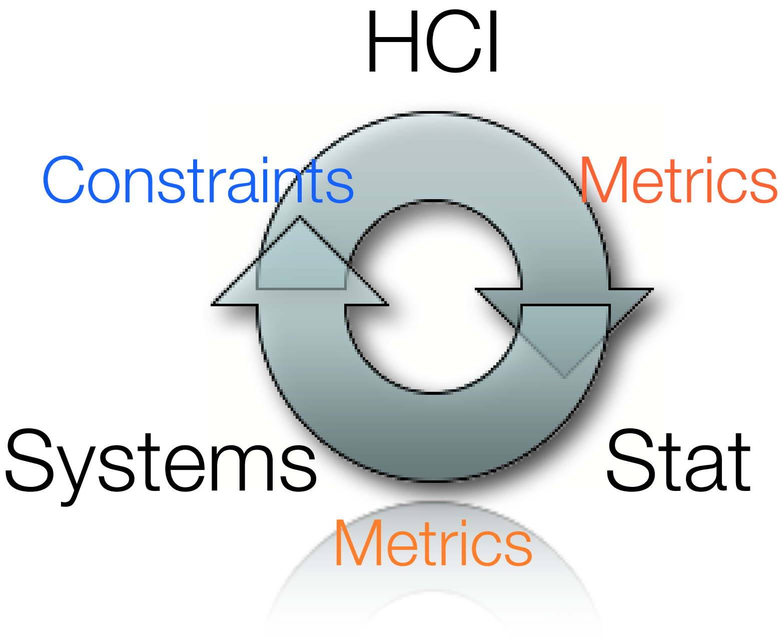

"What unlike thing must meet and mate" -- Melville

With Thanks…

Joe:

Jeff Naughton

Mike Stonebraker

The CONTROL Freaks:

- Ron Avnur
- Andy Chou
- Christian Hidber
- Bruce Lo
- Chris Olston
- Vijayshankar Raman
- Tali Roth

Peter:

Jeff Naughton

Pat Selinger

Bill Cody

S. Seshadri

Ashutosh Singh

Guy Lohman

Vijayshankar Raman

Gang Luo