New Direction for TPC

by

Michael Stonebraker

DBg Database Group
MIT Computer Science and Artificial Intelligence Lab

Outline

- 1985
- 1985-88
- PAFS
- TPC-H
- The future

1985

- Jim Gray writes debit-credit benchmark
  - And gets his friends to be co-authors
  - Commercial systems do about 25 TPS
    - Obviously inadequate
- Jim Gray starts HPTS
  - Goal is 1000 TPS (x40)

1985-88

- Lots of ideas generated on improving OLTP performance
  - Facilitated by HPTS
- Lots of apples-to-oranges debit-credit benchmarks
  - With conventional vendor marketing spin
- But performance improves by an order of magnitude

Obvious Need for

- A level playing field for debit-credit
- A non-vendor organization to carry debit-credit forward

- Enter TPC and TPC-A

Characteristics of Debit-Credit

- Pressing need
  - for better OLTP performance
- Application focused
  - Cash a check
- Simple
  - 5 commands, 5 pages of specification

- Result was vendor focus and much better OLTP systems

Meta - Characteristics

- Find a Pressing need
- Find a simple Application
- Focus the vendor community
- To provide better Systems

PAFS!

TPC-H (PAFS)

- Application/schema doesn't correspond to an obvious business problem
  - schema seems unnatural
  - see Pat's O'Neil's talk

TPC-H (PAFS)

- Way too many queries (22)
- And queries seem politically gerrymandered
  - Can't use materialized views

TPC-H (PAFS)

- No load component in TPC-H
- Users want the ability to perform incremental/trickle load

TPC-H (PAFS)

- Out-of-box experience awful for most systems
- Data base design way too hard – too many knobs
- And automatic tools don't work very well
- RDBMS considered too hard to use by many

TPC-H (PAFS)

- Scalability over a range of sizes is a big issue
- Ability to add resources on the fly is a big issue

TPC-H (PAFS)

- Nobody recovers from the data base log
- No replication in TPC-H

TPC-H (PAFS)

- Major warehouse vendors (e.g. Teradata, Netezza) ignore TPC-H
- Analysts (Forrester, Gartner) say TPC-H is irrelevant

TPC-H (PAFS)

- Current leaders run on silly hardware configuations
  - E.g. 1 Terabyte of disk for a 30 Gbyte configuration (32 X)

TPC-H

- A failure by PAFS standards
- At the very best is “long in the tooth”
  - Follow-on effort (TPC-DS) is worse by PAFS standards
- And TPC progress is at the speed of molasses

TPC-H

- A failure by PAFS standards
- At the very best is “long in the tooth”
  - Follow-on effort (TPC-DS) is worse by PAFS standards
- And TPC progress is at the speed of very slow molasses
  - E.g. little stomach to fix these issues

TPC-C

- Essentially same comments apply

Summary of TPC

- Is very slow moving
- Seems vendor dominated
  - Political and not user focused
- Not focused on PAFS

So What to Do?

- Go back to your roots
- E.g. PAFS
  - In your traditional market
  - In new markets

Example – One Among Many

- Science applications (e.g. Chemistry, Earth Sciences, Remote Sensing, ….)
- Universally hate current RDBMS

Nearest neighbor queries, time series queries

Snow Cover in the Sierras

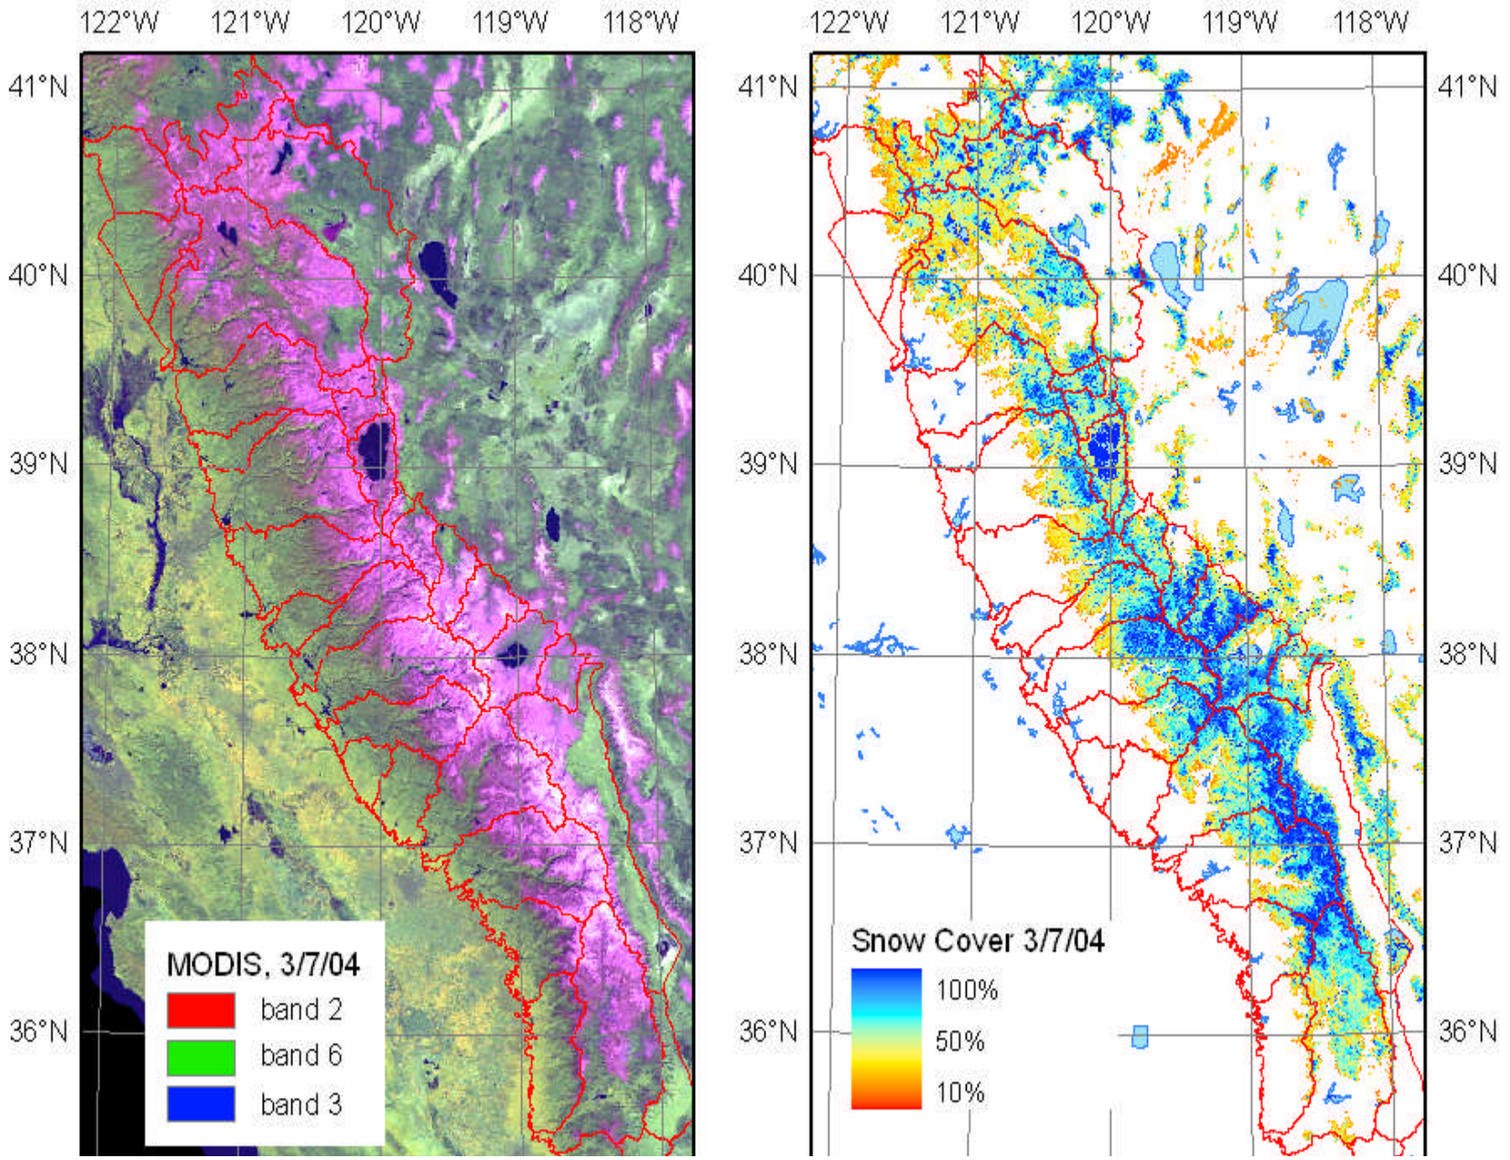

Protein Structure

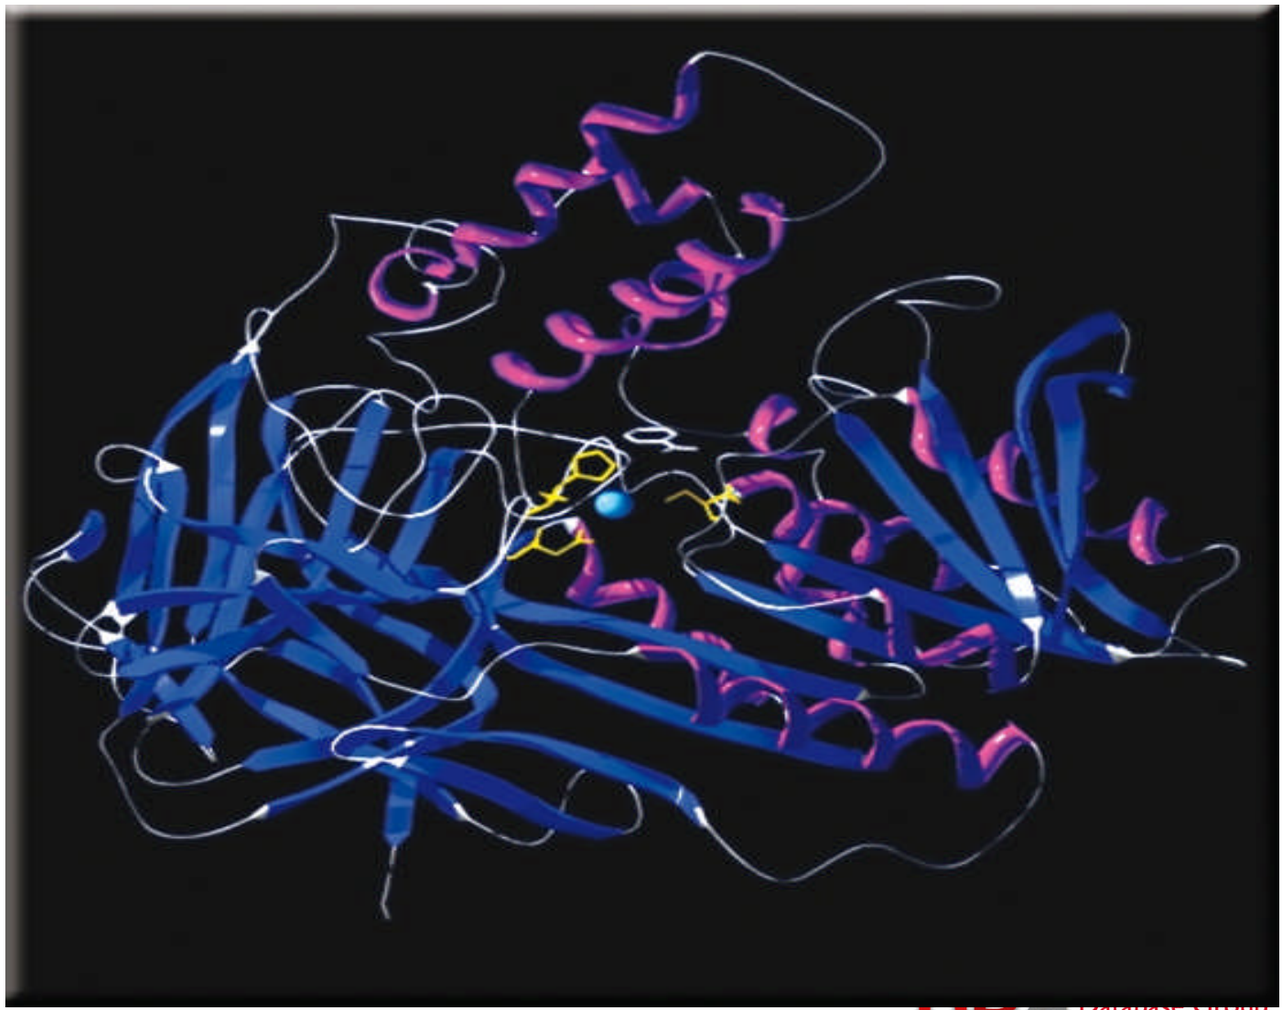

Chromatin Structure

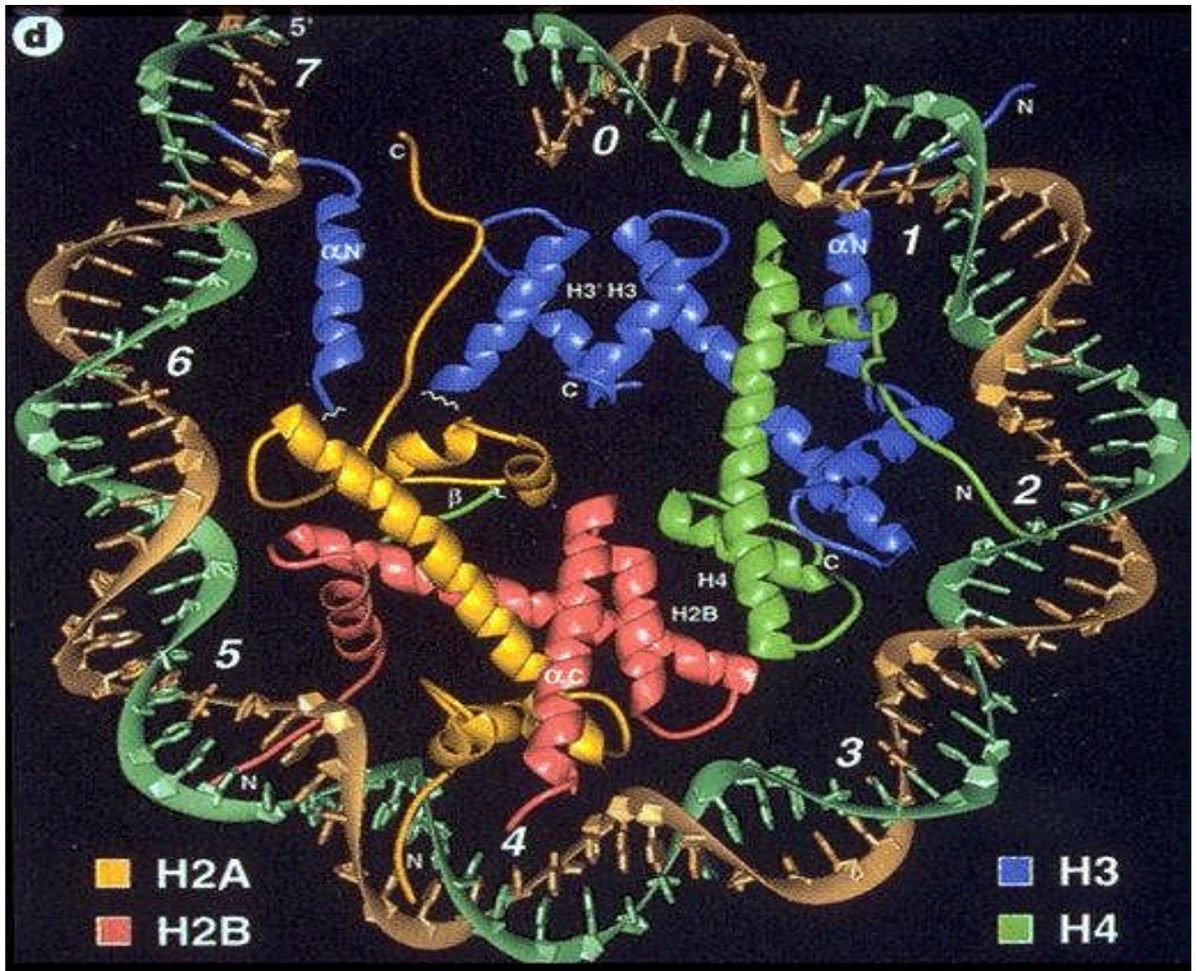

DNA

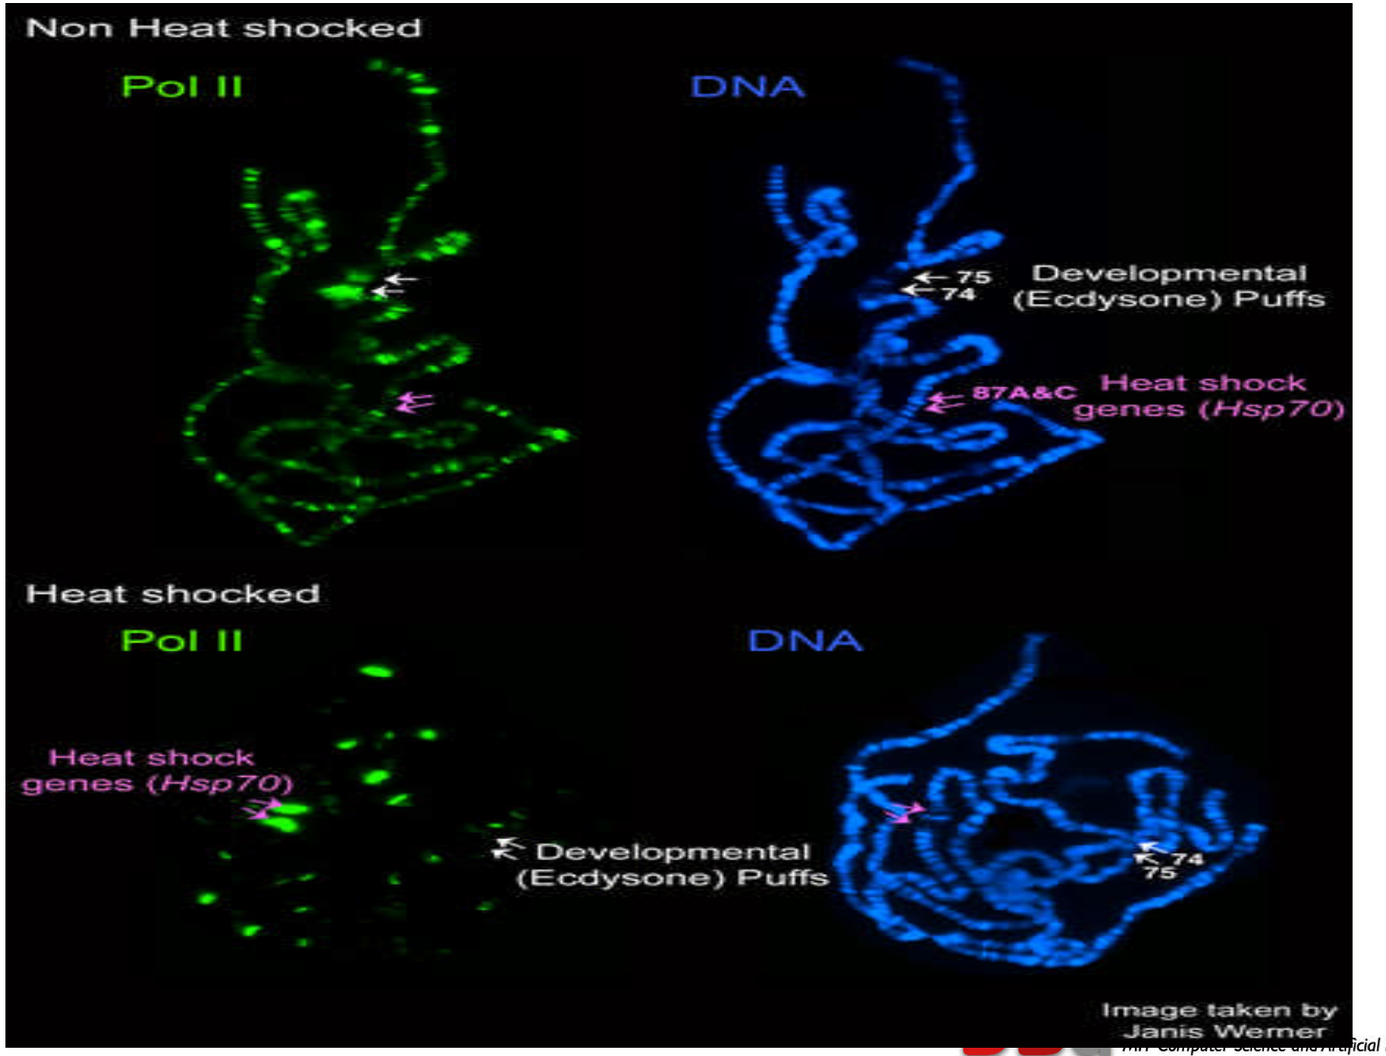
Non Heat shocked
Pol II
DNA
75
74
Developmental
(Ecdysone) Puffs
87A&C
Heat shock
genes (Hsp70)
Heat shocked
Pol II
DNA
Heat shock
genes (Hsp70)
Developmental
(Ecdysone) Puffs
74
75
Image taken by
Janis Werner

Human Genome Matching

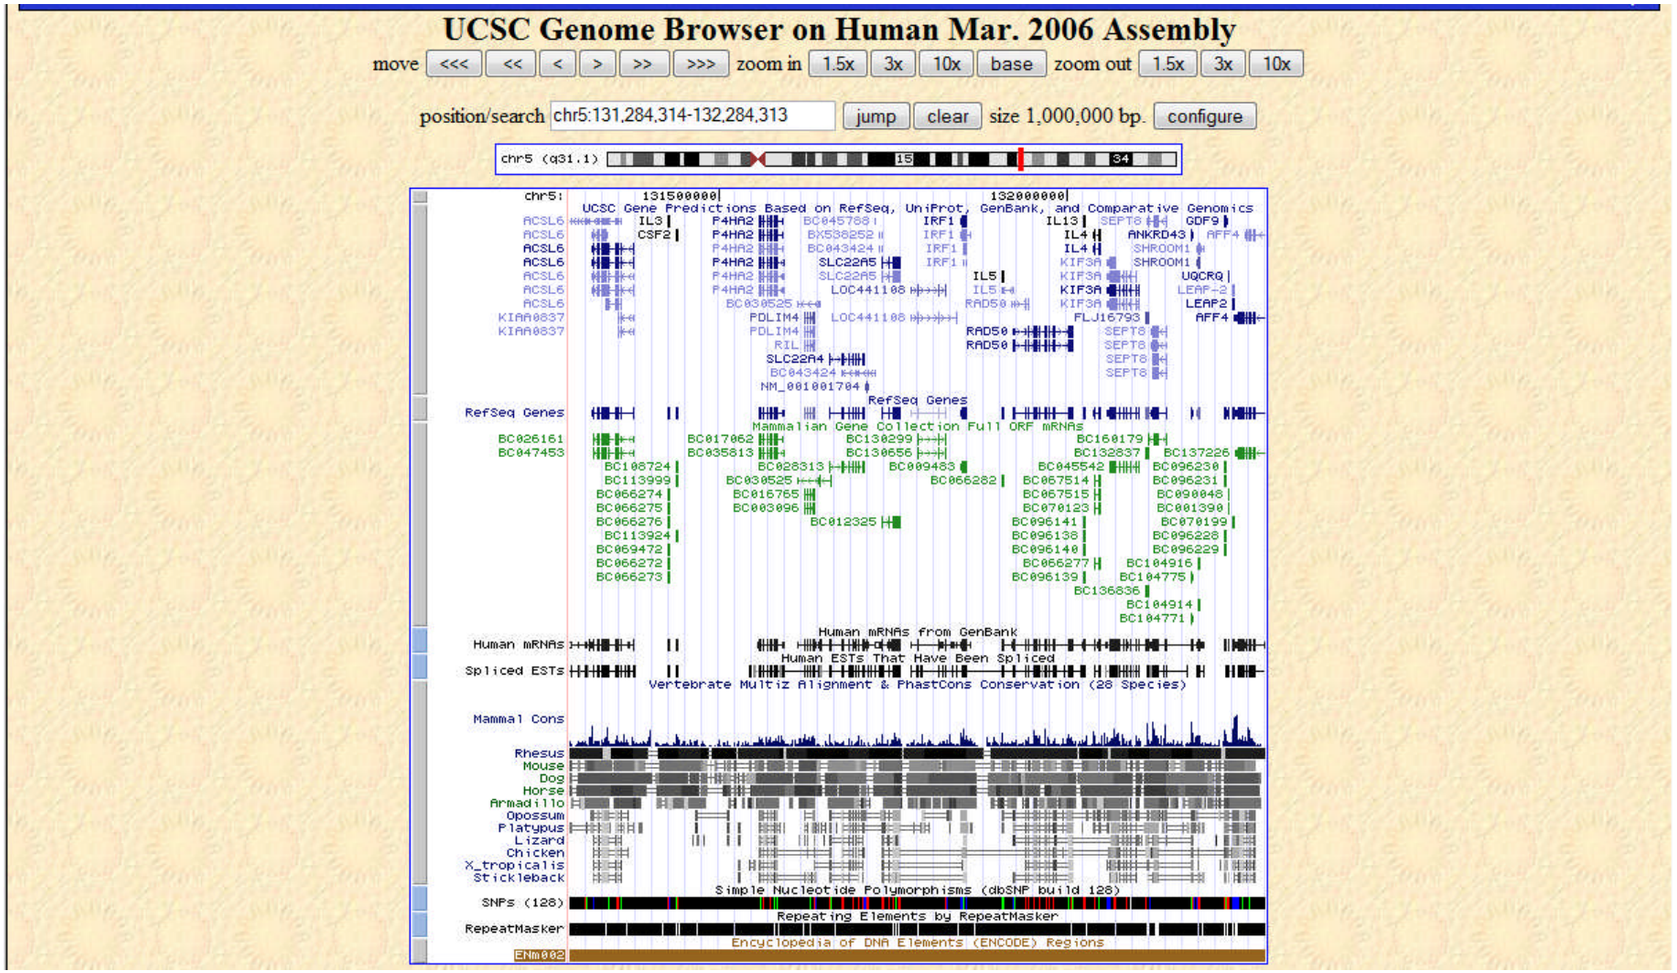

http://genome.ucsc.edu/ENCODE/encode.hg18.html

Why?

- Wrong data model
  - Remote sensing guys want arrays
  - Which are horribly inefficient and usually very unnatural to simulate on top of tables

Why?

- Wrong operations
  - Consider two satellite imagery data sets, one with 50m cells in lat-long and one with 75 meter cells in mercator
  - Need to regrid one to the other as a DBMS operation
    - Regrid needs to be built in

Why?

- Wrong features
  - Need provenance (i.e. ability to tell how a data element was derived)
    - Requires a log of all operations and some provenance-oriented operations
  - And repeatability (i.e. rederive the scientific calculation if necessary)
    - Requires no-overwrite storage and time-travel

Net Result

- Science does not use RDBMS (for anything other than metadata)
- Crying need not being met by current systems!
- A PAFS effort by TPC could change all this!!

Same Story

- In RDF
- In Web 2.0 companies
- In real-time data manipulation
- In Map-Reduce style computing

So What is the Best Route Forward?

- Best benchmarks are written by one person (e.g. debit-credit)
  - Typically in small numbers of days
  - And reviewed by the community in small numbers of weeks
  - And adopted in months (not years or decades)

So What is the Best Route Forward?

- There are lots of academic benchmarks that fit this model and have gained traction, e.g.
  - Linear road (streaming data)
  - MR benchmark (MR vs DBMS)
  - Madden/Abadi RDF benchmark

So What is the Best Route Forward?

- Troll the research world for such things

So What is the Best Route Forward?

- Involve research community in your activities
  - But nobody will do so with your current heavyweight process
  - you will have to violently streamline

So What is the Best Route Forward?

- Switch from a vendor-focus to a user-focus
  - Only way to get PA in PAFS

I.e. It is Time for TPC to Reinvent Itself

- Mantra has to be PAFS
- Streamline process
- Involve research community
- New charter!
- Everybody should do this once a decade – you are a decade late

Otherwise

- TPC will become a legacy world only relevant in some traditional business data processing areas
- i.e. you will walk into the sunset of irrelevance