

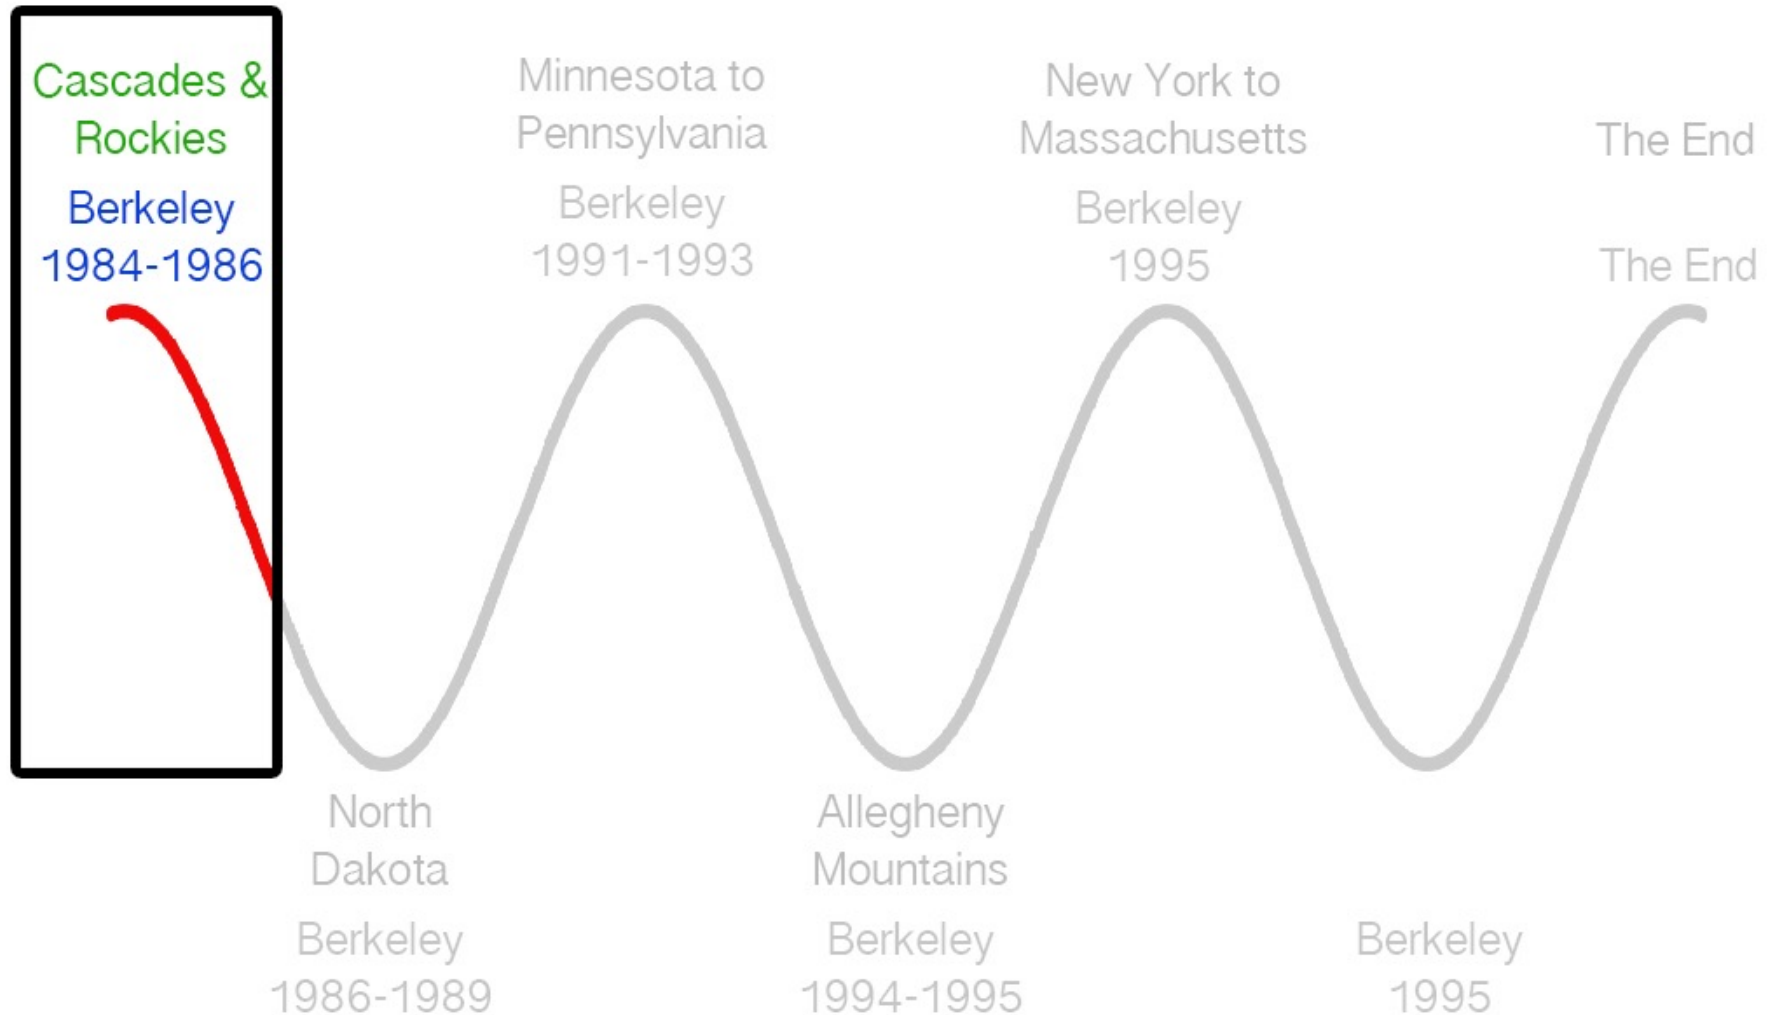
The Land Sharks are on the Squawk Box

(Why Riding a Bicycle Across America
and Building Postgres Have a Lot in
Common)

Kennebago, ME: Summer 1993



Part 1: Off to a Great Start



Anacortes, WA: Day 1 - June 4, 1988



Berkeley, CA: 1984

- Commercial Ingres is 4 years old
 - And much better code than academic Ingres
- Makes no sense to continue prototyping on our code line
- We push the code off a cliff, and start something new
- So what is Postgres to be?

Berkeley, CA: 1984

- Abstract Data Types (ADTs)
 - Ingres, System R were focused on business data processing (competing against IMS and Codasyl)
 - Motivated by a large collection of papers of the form:
 - “we tried RDBMS on X; it didn’t work, so we added Y to the relational model”

Example from GIS

name	age	xmin	xmax	ymin	ymax
Sam	38	100	220	46	87

Find me all employees who overlap a specific (14, 17, 0, 16) rectangle?

Select name where
 $xmin < 17$ and $xmax > 14$ and
 $ymin < 16$ and $ymax > 0$

Messy and very difficult to optimize!!!

Better Solution

name	age	location
Sam	38	Internal representation for a box

Select name where location Overlaps
MakeBox (14, 17, 0, 16)

In General

User-defined data types (Rectangle)

User-defined operators (Overlaps)

User-defined functions (MakeBox)

Requires adding functions to DBMS engine and calling them at appropriate times during query execution...

The Devil is in the Details

- Using new types in indexes (B-trees, hashing)
- New indexes (e.g. R-trees)
- Teaching the query optimizer about new types
- Specifying commutativity rules
- Optimization of “not”
- And others

Winthrop, WA: Day 3



Rainy Pass, WA: Day 3



Washington Pass, WA: Day 3



Berkeley, CA: 1985-86

- Chris Date wrote a pioneering paper on referential integrity
- Basically what to do with foreign-key-primary-key constraints

Berkeley, CA: 1985-86

dname	floor	sq. ft.	budget
Shoe	3	500	40,000
Candy	2	800	50,000

name	dept	salary	age
Bill	Shoe	2000	40
Art	Candy	3000	35
Sam	Shoe	1500	25
Tom	Toy	1000	23

What happens if you delete the Candy department?

Date's Solution

- Cascade the delete
 - Replace dept by “null”
 - Refuse the delete
-
- Plus three corresponding actions on inserts

My Idea: A Single Mechanism

- Referential integrity
- Triggers (make sure a salary update to Sam gets cascaded to Bill)
- Integrity constraints (salaries must be non-negative)

I Hate IF-THEN Rule Systems

- Consider two rules:
 - Make sure George is in the Shoe department
 - Move everybody over 40 from Shoe to Candy
- Order dependent!!!!
- Made worse by SQL being a set-oriented language!
- No mere mortal can understand a 20 line Prolog program!


My Solution

- Update ALWAYS command
- Make sure Bill has the same salary as Sam
 - Update ALWAYS (set salary = E.salary)
where E.name = “Sam” and
name = “Bill”
 - Semantics: the command appears to be
perpetually running
- Somewhat cleaner than an if-then rules
system....

Marias Pass, MT: Day 15



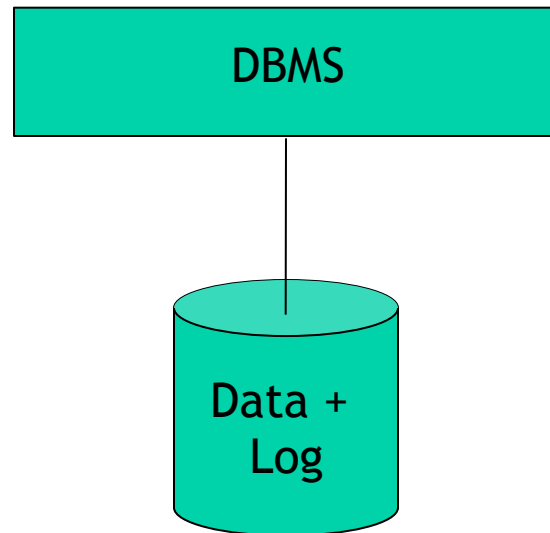
Traditional DBMS Crash Recovery



Two data stores with different access engines, which must be carefully synchronized!!

Log code is complicated and ugly!

Postgres Picture: No Overwrite



Gets rid of crash recovery code!
Provides “time travel” in addition

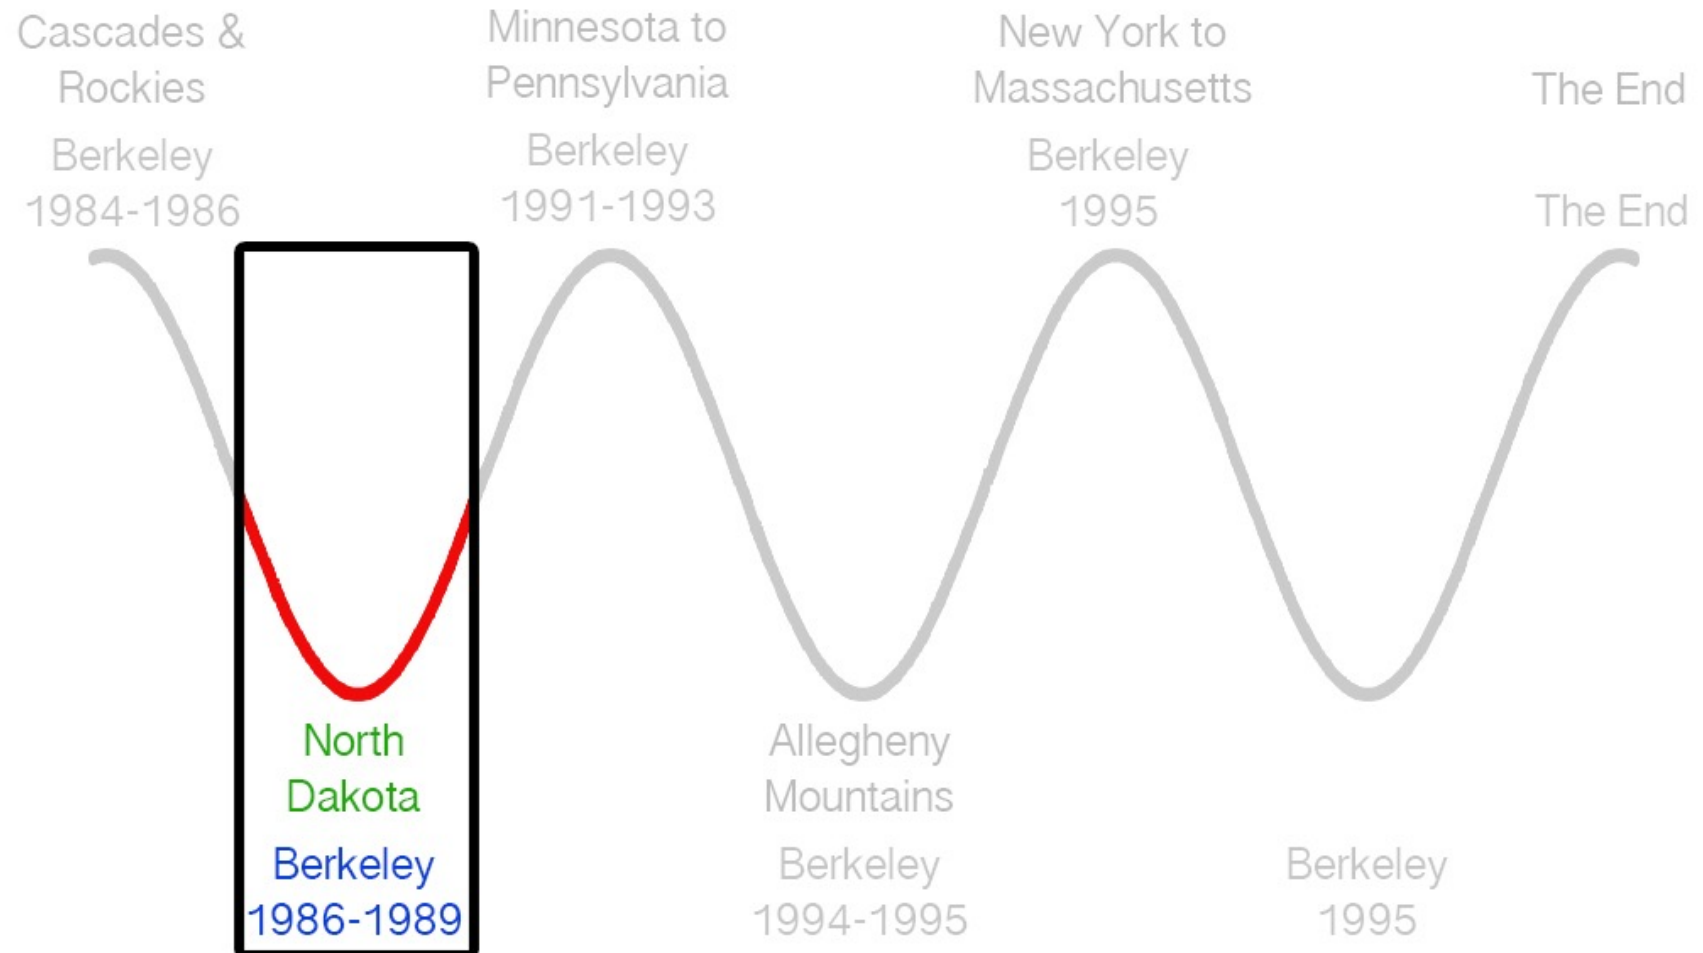
The Devil is in the Details

- Log is “write optimized”
- Data is “read optimized”
- No overwrite must be competitive - tricky placement problem

Postgres Game Plan

- ADTs
- Rules using an “always” command
- No overwrite storage
- There was other (less important) stuff...
 - Complex objects
 - Inheritance

Part 2: Things Aren't Going Well



Drake, ND: Day 26



Berkeley, CA: 1986-87

Reality Sets In

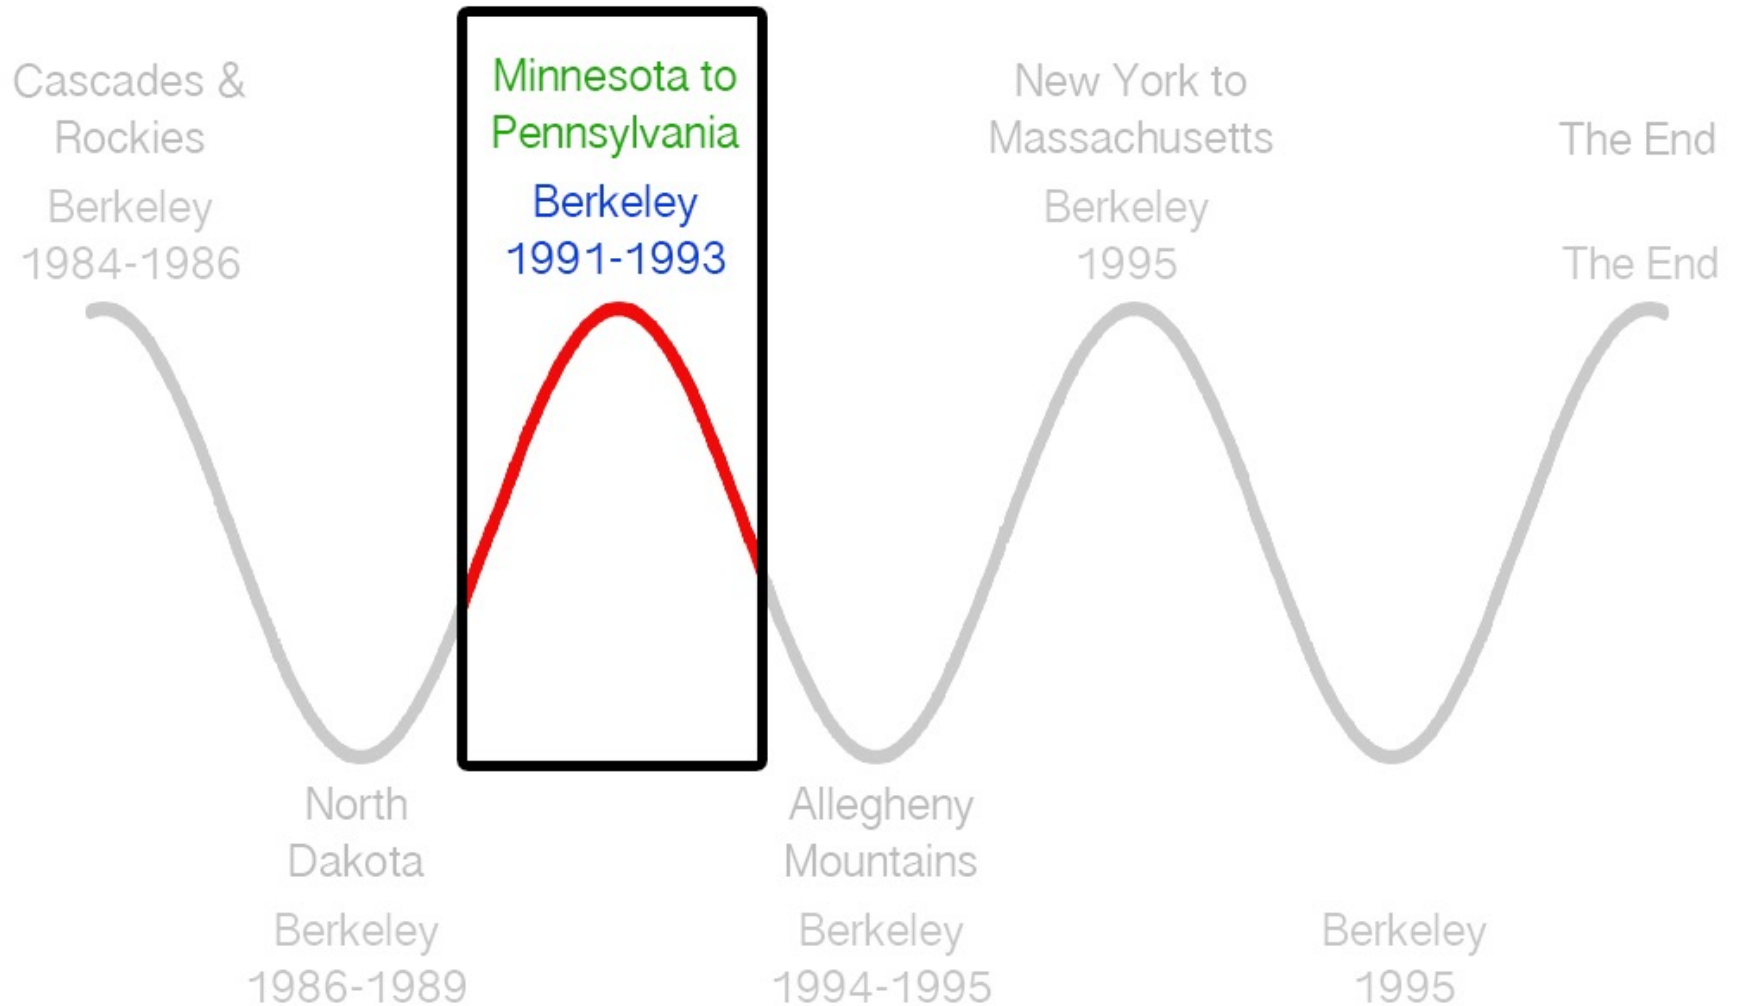
- Lisp is a disaster
 - Lots of code over the cliff
- ADT system works well
 - Multiple semantics for time
- Time travel has promise, but is a tricky to tune
- “Always” doesn’t work
 - Lots of code over the cliff
 - And we go back to a exploring a conventional rule system

Berkeley, CA: 1986-90

Reality Sets In

- The “Wine Connoisseur” and I spend a lot of time on:
 - How complex can complex objects be
 - What to do about time travel, rules, ...
- Next 3 years was a lot of “repair”
 - Think of this as a “slog through the swamp”

Part 3: Another High Point



Carrington, ND: The Next Afternoon Serendipity



Berkeley, CA: 1991-92

Commercial Postgres

- In the DBMS field, the ultimate arbiter of good ideas is the commercial marketplace
- In my opinion, technology transfer is best done through a startup
 - Large companies have their own agenda
 - Read “The Innovators Dilemma” by Clayton Christensen

Berkeley, CA: 1991-92

Commercial Postgres

- Recruited “Mom” and her husband “Short One”
- “EMP1” and “Quiet One” came over from the academic team, joined shortly thereafter by “Triple Rock”
- With the “Tall Shark” as interim CEO and lead investor
- We were off and running – converting the query language from QUEL to SQL, hardening the code and improving performance

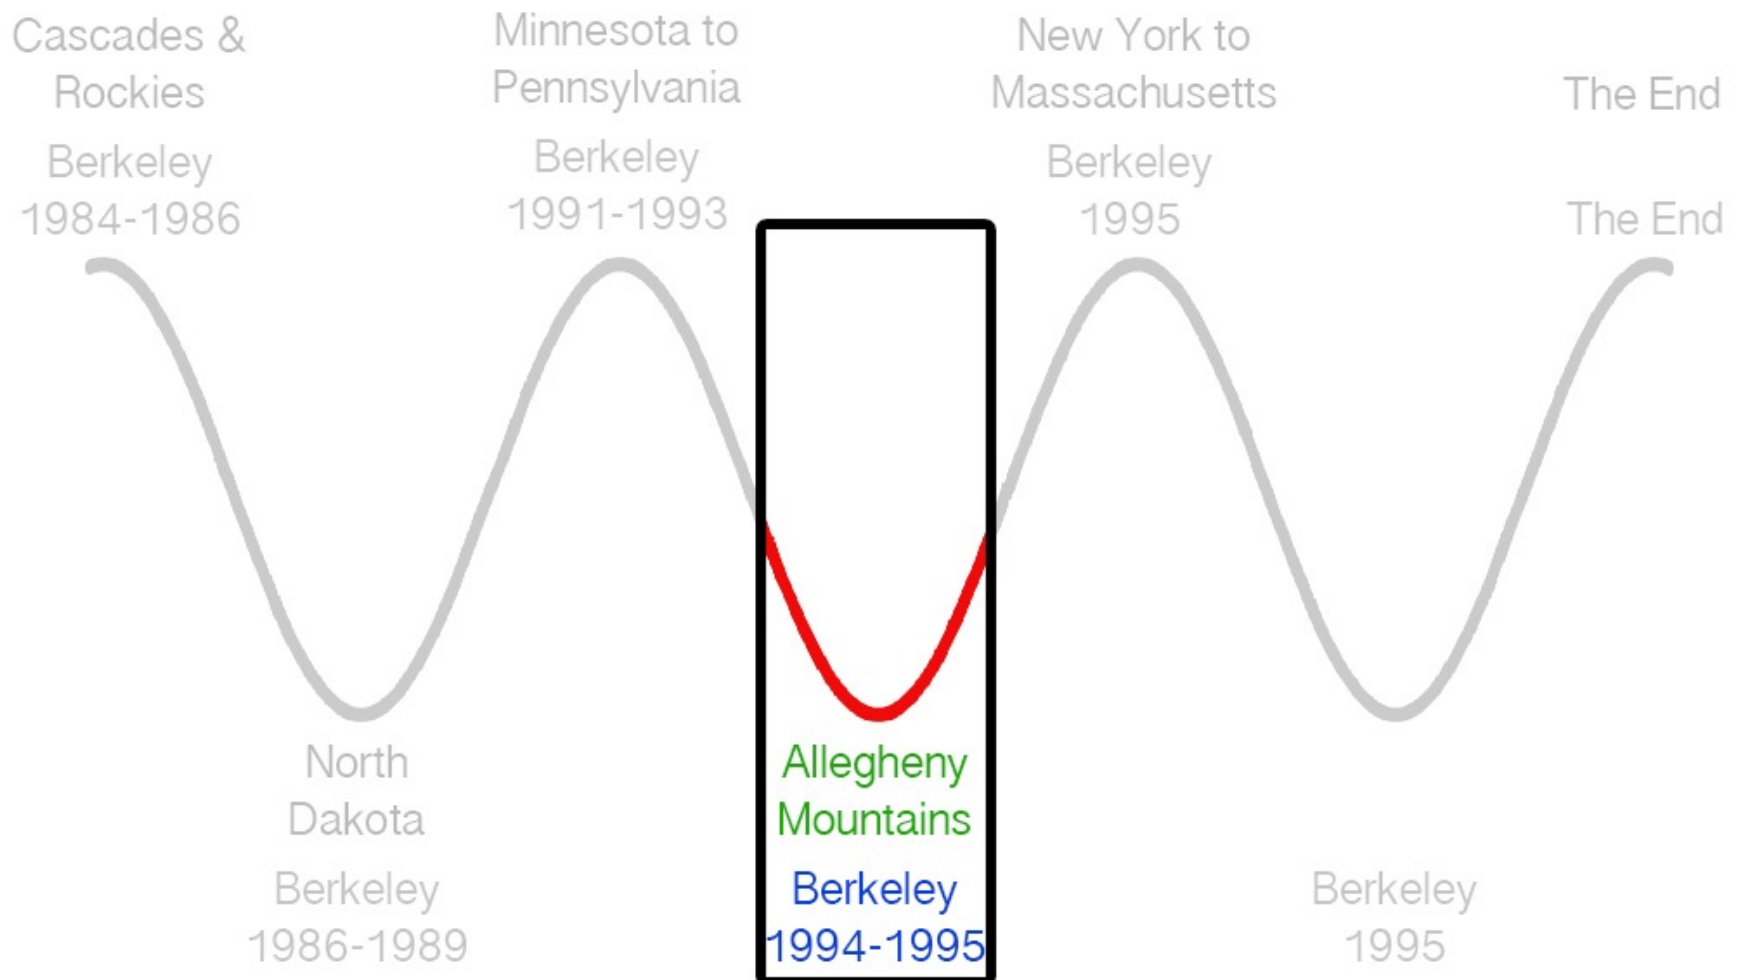
Luddington, MI: Day 38



Berkeley, CA: 1993-94

- After a couple of miscues, we are named Illustra
- Got the first few customers
- Raised more money (see slide 1)
- Hired “Voice of Experience”, “Uptone” and “Smooth” -- a real management team
- The future is looking up....

Part 4: Another Low



Ellicottville, NY: Day 49




Berkeley, CA: 1994 - 1995

Reality Sets In

- Catch-22 between customers and ADT providers
 - Customers wanted ADTs from the large app vendors
 - App vendors wanted a distribution channel
- The dreaded “down round”
 - “a fate worse than death”

Part 5: Looking Good



Troy, NY: Day 56

The Last Hill

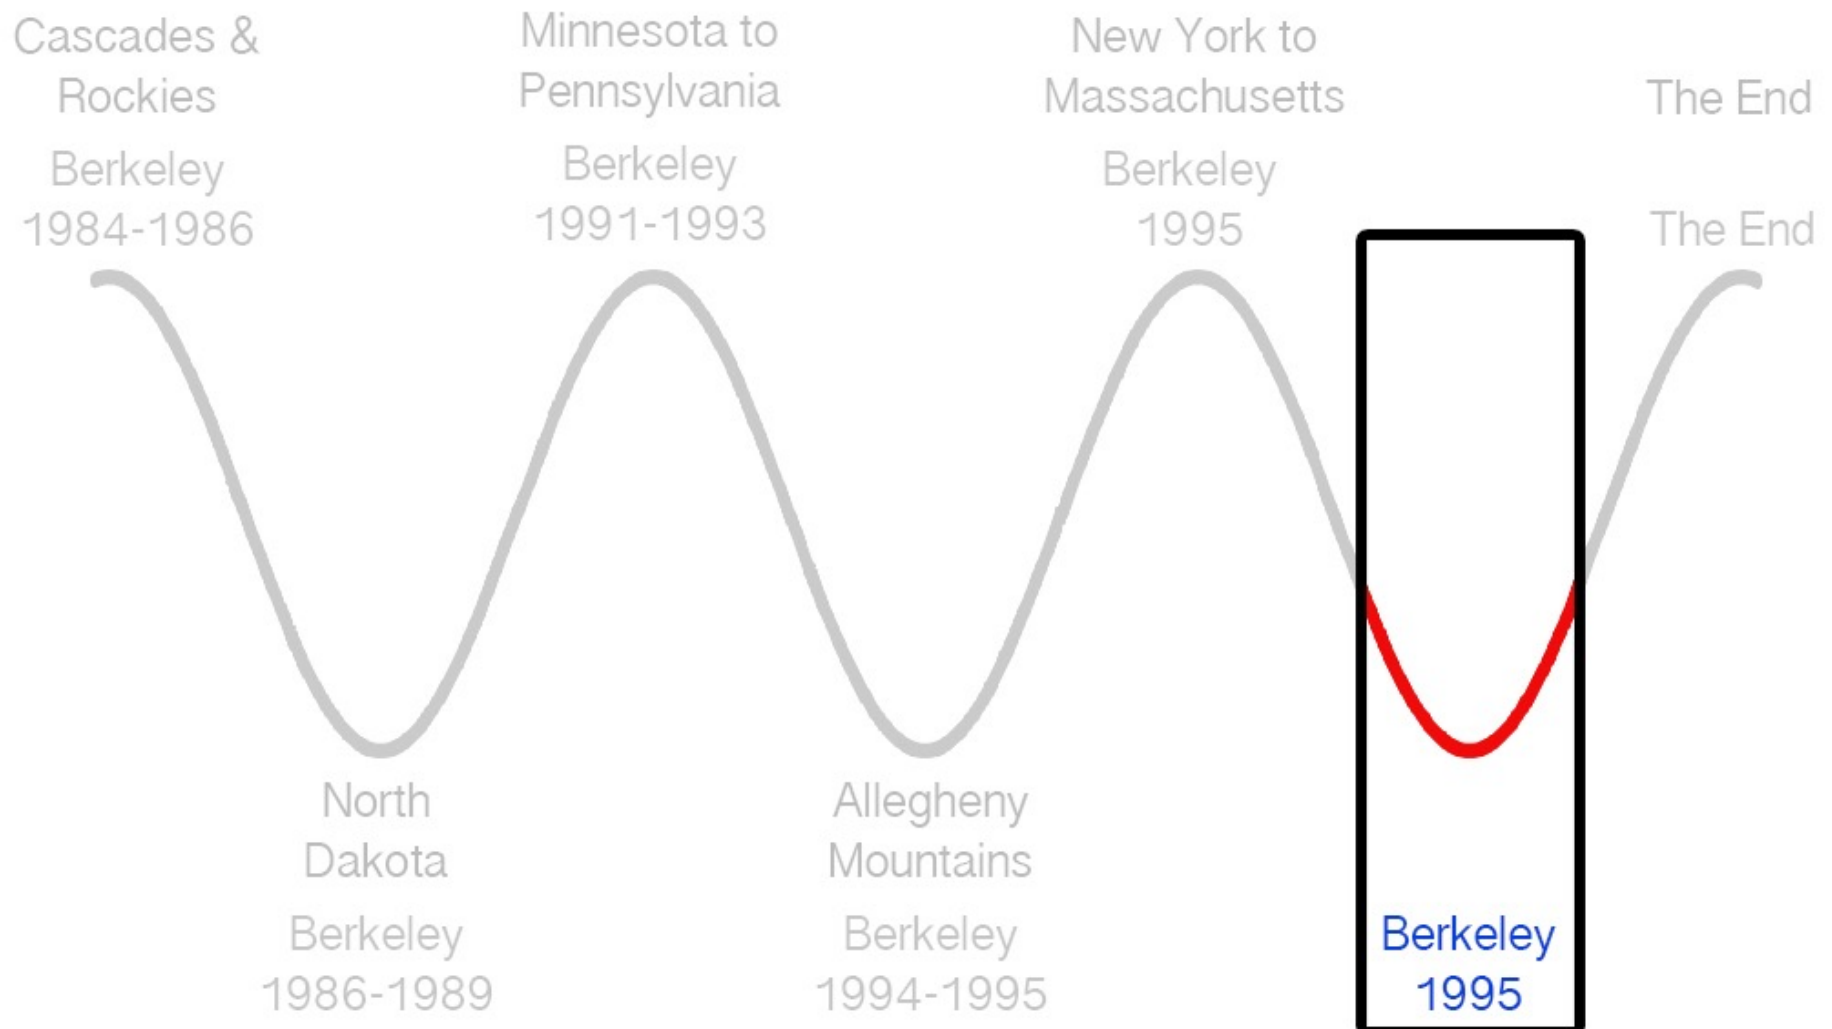


Berkeley, CA: 1995

Serendipity

- The internet is taking off
- “Uptone” turns the company on a dime and we become “the database for cyberspace”
- Get a lot of buzz
- Maybe we have turned the corner

Part 6: More Challenges




Berkeley, CA: 1995

Reality Sets In (Again)

- The internet bake off
- On business data processing data!!!!
 - Transaction processing is not in our “wheelhouse”

Part 7: The Ending



Wollaston Beach, MA: Day 59



Berkeley, CA: February 1996

Serendipity

- Illustra is purchased by one of the elephants
 - Solves the Catch-22 on ADTs
 - Solves the transaction processing problem
- We set about putting Illustra features into their system

Why Tell a Bicycle Story?

Algorithm for a cross-country bike trip:

```
Until (ocean) {  
    get up in the morning;  
    ride east;  
    persevere and overcome any obstacles  
    that arise;  
}
```

A Bit of Abstraction...

```
Until (GOAL) {  
    get up in the morning;  
    ACTION;  
    persevere and overcome any obstacles  
    that arise;  
}
```

And a Macro

Make It Happen

Why Would Anybody Want to Ride Across America?

- Long and hard
- Periods of depression and elation
- Lots of monotony (slogging through the swamp)

My Career

Make It Happen (PhD) – 5 years

Make It Happen (tenure) – 5 years

Make It Happen (bike ride) – 2 months

There is an obvious conclusion to be drawn...

Building Postgres Required Make It Happen

Have a good idea (or 2 or 3)

Make it Happen (prototype)

Start a company

Hire superb implementors

Hire talented execs

Make It Happen (product)

(Have a good idea then

Make It Happen for 10 years)

Why would anybody in their right mind do this?

My Personal Thoughts On Good Ideas

- I have never had any good ideas by communing with nature on the mountaintop
- Instead, I recommend:
 - talking to users
 - talking to more users
 - hanging around a critical environment of smart people (like MIT or Berkeley)
- For my take on current DBMS ideas: see my test-of-time talk in ICDE '15

Closing Thoughts

- Where does Make It Happen come from?
 - I have no idea!
- What about serendipity?
 - Really important in North Dakota!
 - Really important to Illustra!
 - Often a very significant factor
 - As is anti-serendipity!!!!

Boston, MA: Present Day

- Our bicycle (Boston Bound) sits in our basement gathering dust
 - It has not been ridden since Wollaston Beach
- I still love to take on physical challenges (climbing the NH 4000 footers)
- Illustra was successfully integrated into the Informix code base and is still available

Boston, MA: Present Day

- Postgres ADTs have been added to most RDBMSs
- The biggest impact of Postgres came from the shepherding of the open source Berkeley system
 - Pure serendidity!
- Postgres code has formed the starting point for several DBMS startups

Boston, MA: Present Day

I am still building new DBMSs (“one size fits one”):

C-Store/Vertica	column store for the data warehouse market
H-Store/VoltDB	main-memory row store for the OLTP marketplace
SciDB/Paradigm4	array engine for the complex analytics marketplace

Each a decade (and counting)

Could have used any of these in this talk!

The Following 39 Berkeley Students and Staff Wrote Postgres (with data model collaboration from Larry Rowe)

Jeff Anton, Paul Aoki, James Bell, Jennifer Caetta,
Philip Chang, Jolly Chen, Ron Choi, Matt Dillon,
Zelaine Fong, Adam Glass, Jeffrey Goh, Steven Grady,
Serge Granik, Marti Hearst, Joey Hellerstein,
Michael Hirohama, Chin-heng Hong, Wei Hong,
Anant Jhingren, Greg Kemnitz, Marcel Kornacker,
Case Larsen, Boris Livshitz, Jeff Meredith, Ginger Ogle,
Mike Olson, Nels Olsen, Lay-Peng Ong, Carol Paxson,
Avi Pfeffer, Spyros Potamianos, Sunita Surawagi,
David Muir Sharnoff, Mark Sullivan, Cimarron Taylor,
Marc Teitelbaum, Yongdong Wang, Kristen Wright, and
Andrew Yu

None of This Would be Possible Without:

- Incredibly talented system programmers
 - People like “EMP1,” “Triple Rock” and “Quiet One”
- Start-up executives
 - People like “Uptone,” “Smooth” and “Voice-of-experience”
 - “Cueball”
- The land sharks
 - Especially “The Believer”

None of This would be Possible Without My Wife and Family



In Memory Jim Gray

