

*LARGE SCALE MACHINE LEARNING WITH THE
SIMSQL SYSTEM*

**Chris Jermaine
Rice University**

**Many Current and Past Rice Team Members
Also, Peter J. Haas at IBM Almaden**

This talk is About

- Programming environments/execution platforms for big ML

- I'm a database guyTM

- I'm a database guyTM
- What does that mean?

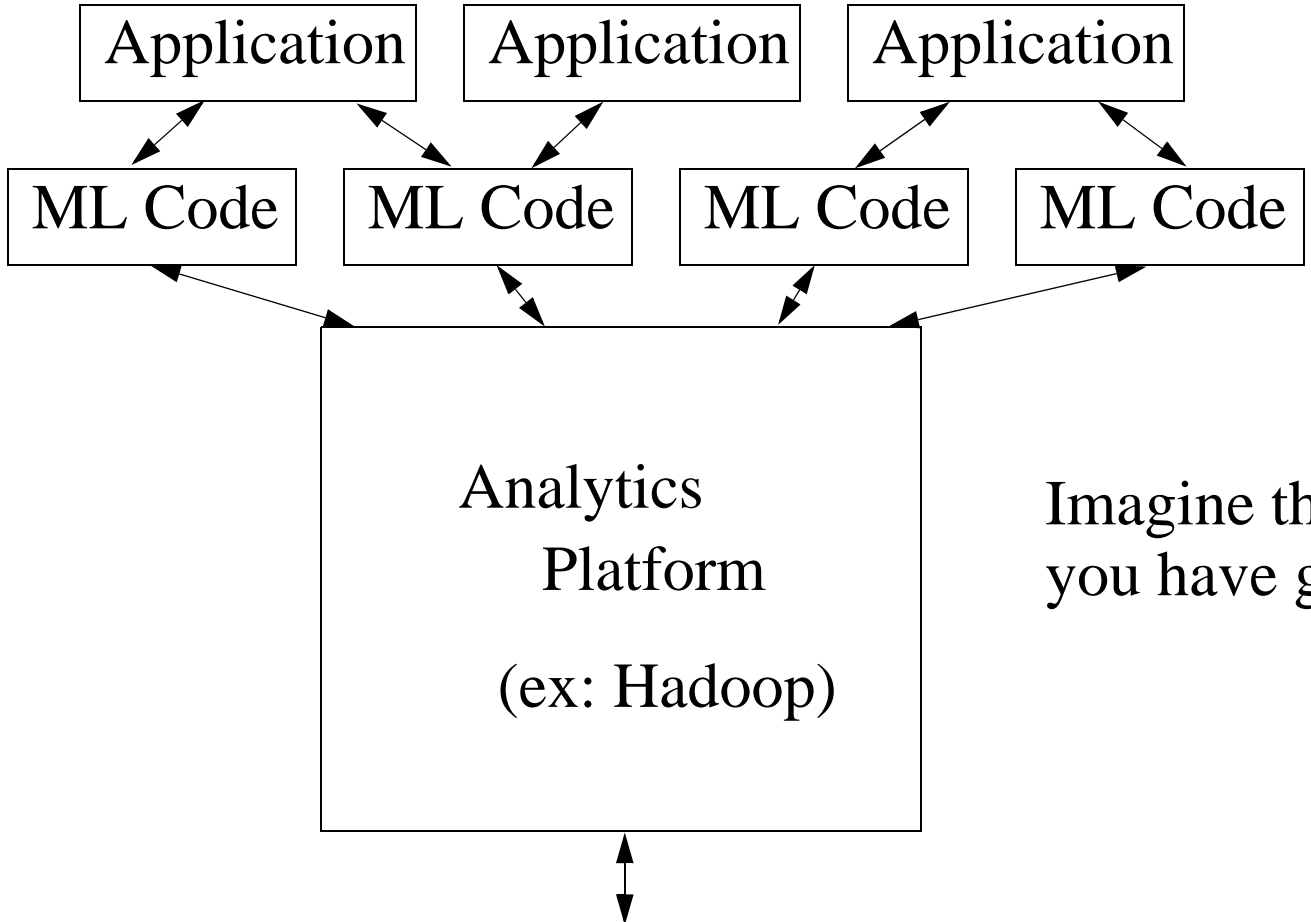
- I'm a database guyTM
- What does that mean?
- To me, means that I worship at the church of **data independence**
 - Now what in the heck does *that* mean?

- I'm a database guyTM
- What does that mean?
- To me, means that I worship at the church of **data independence**
 - Now what in the heck does *that* mean?
- Means that when one designs a data-processing system...
- It should strive for the following ideal:
 - Coder specifies **what** the computation result should be, not **how** to get there
 - System itself figures out the **how** (the “declarative” paradigm)
 - Means code can be independent of data format, size, schema, processing hardware
 - Same code runs on one box with a GPU and on a 1000-machine cluster


Why Are Declarative and Data-Driven Code?

Declarative Arg. 1: One code, many backends

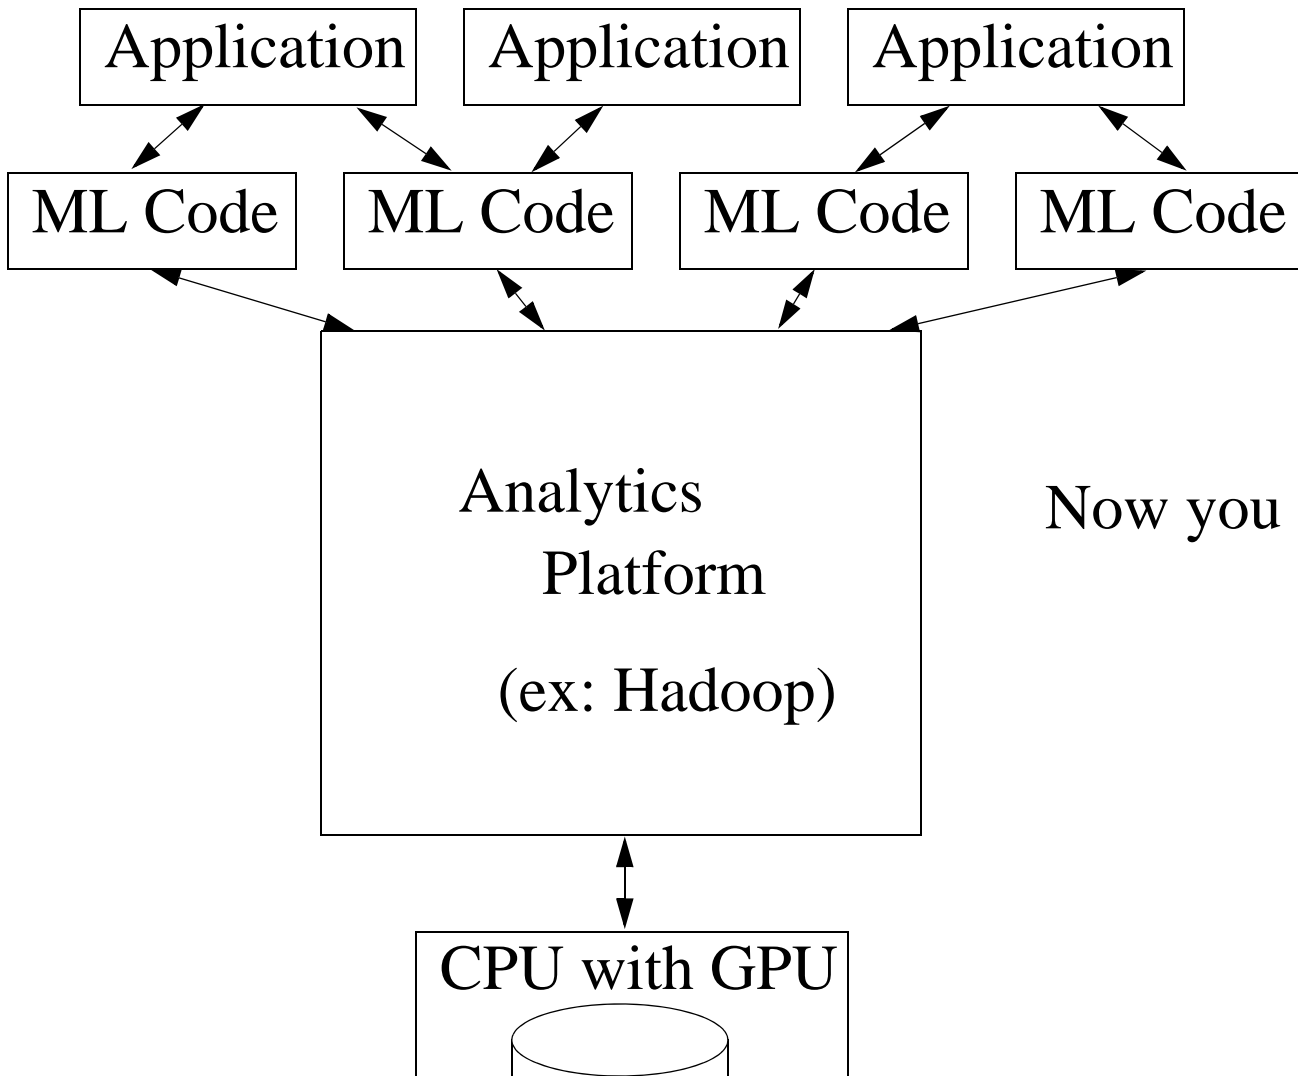
Declarative API: One code, many backends



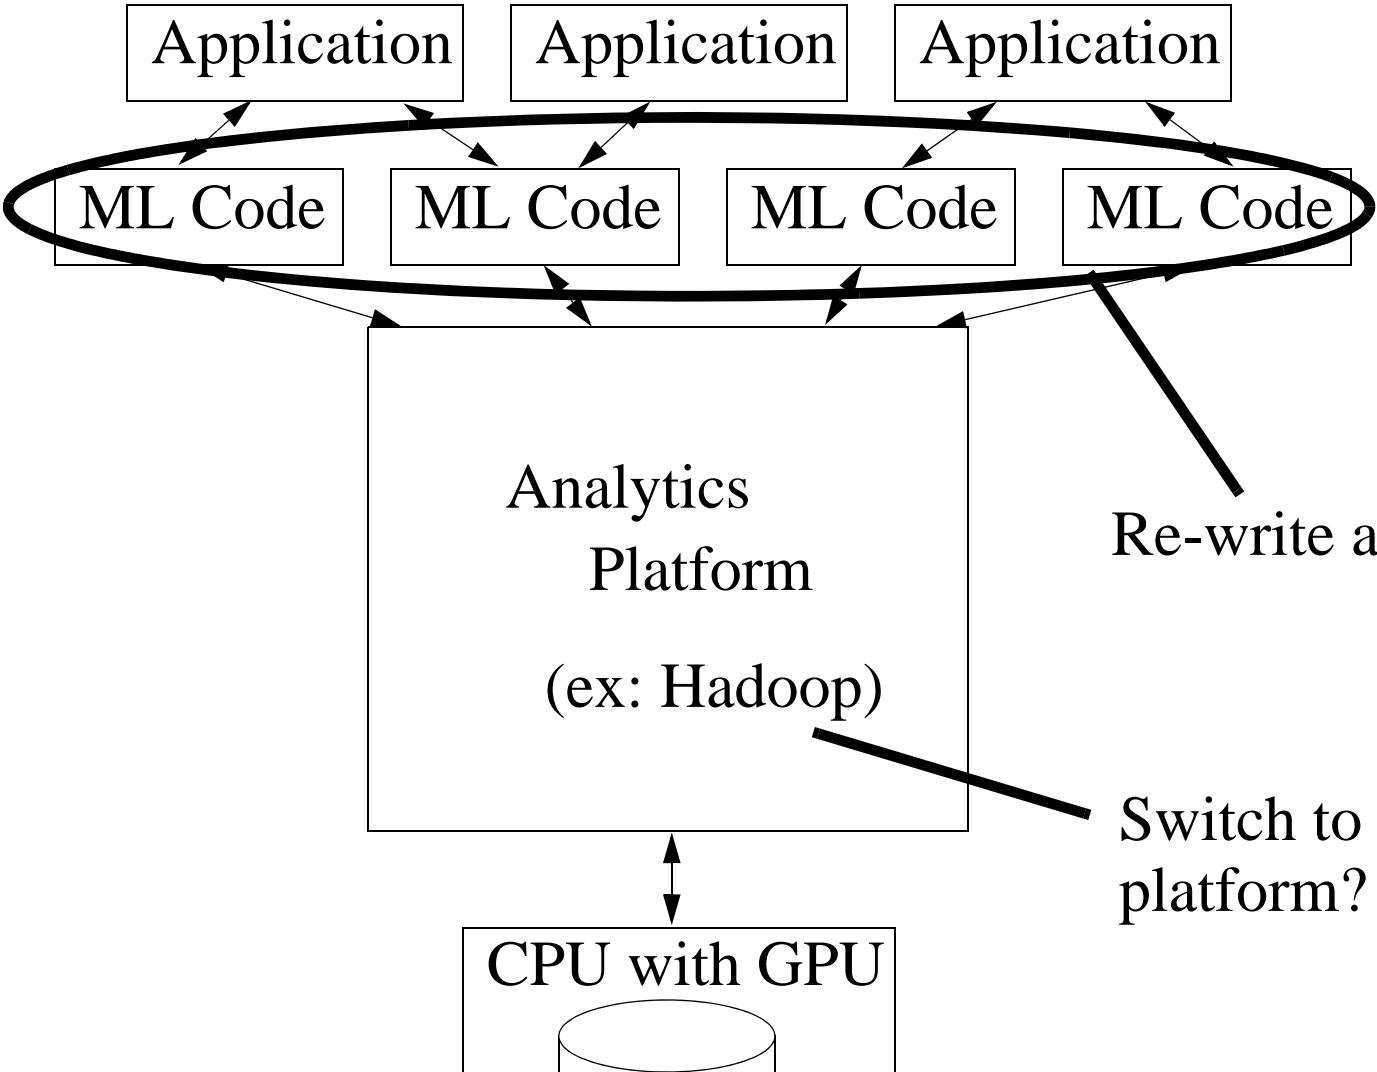
Imagine that you have got this...



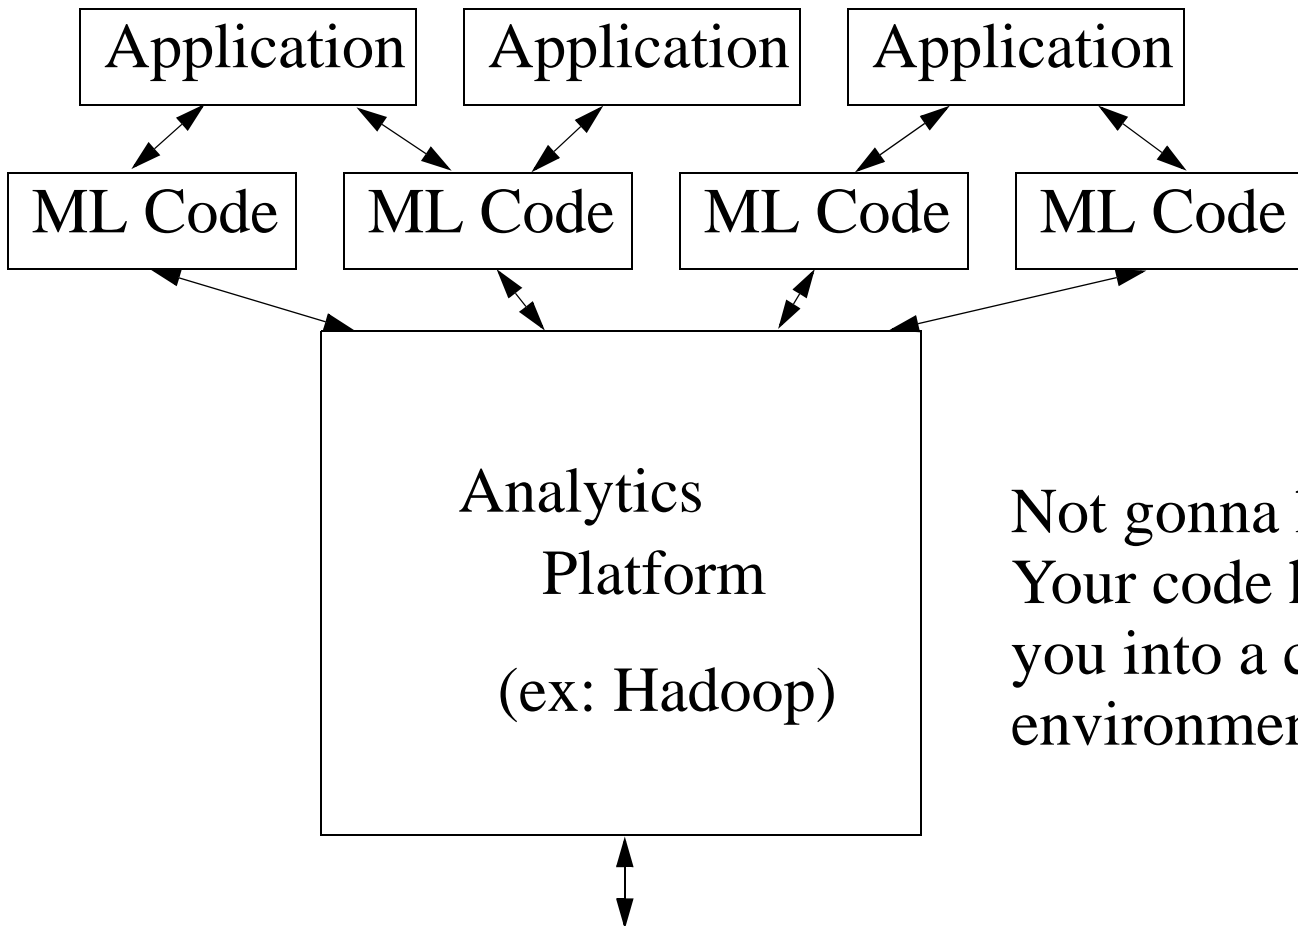
Declarative ML: One code, many backends



Declarative ML: One code, many backends



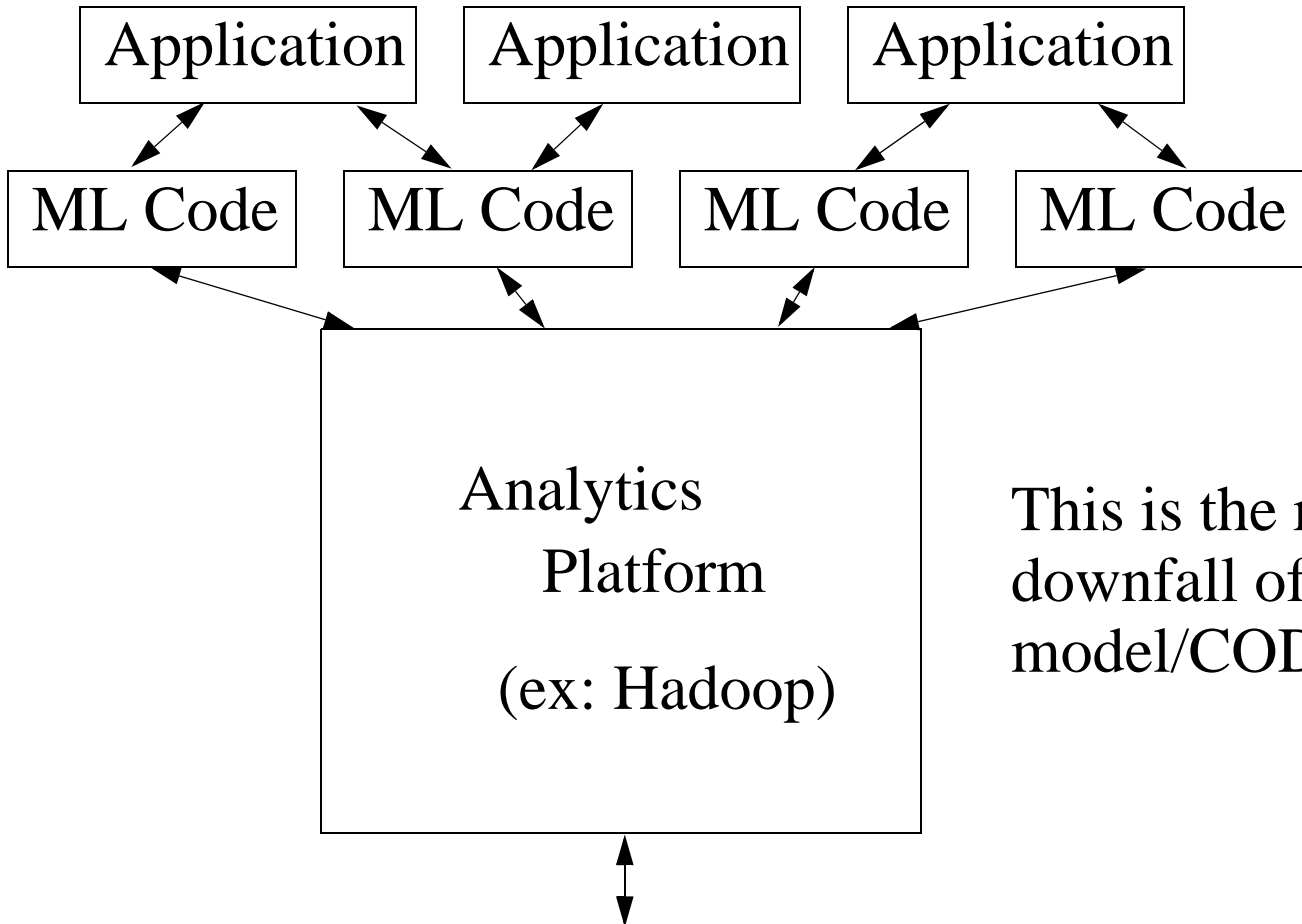
Declarative API: One code, many backends




Not gonna happen!
Your code has locked
you into a compute
environment...



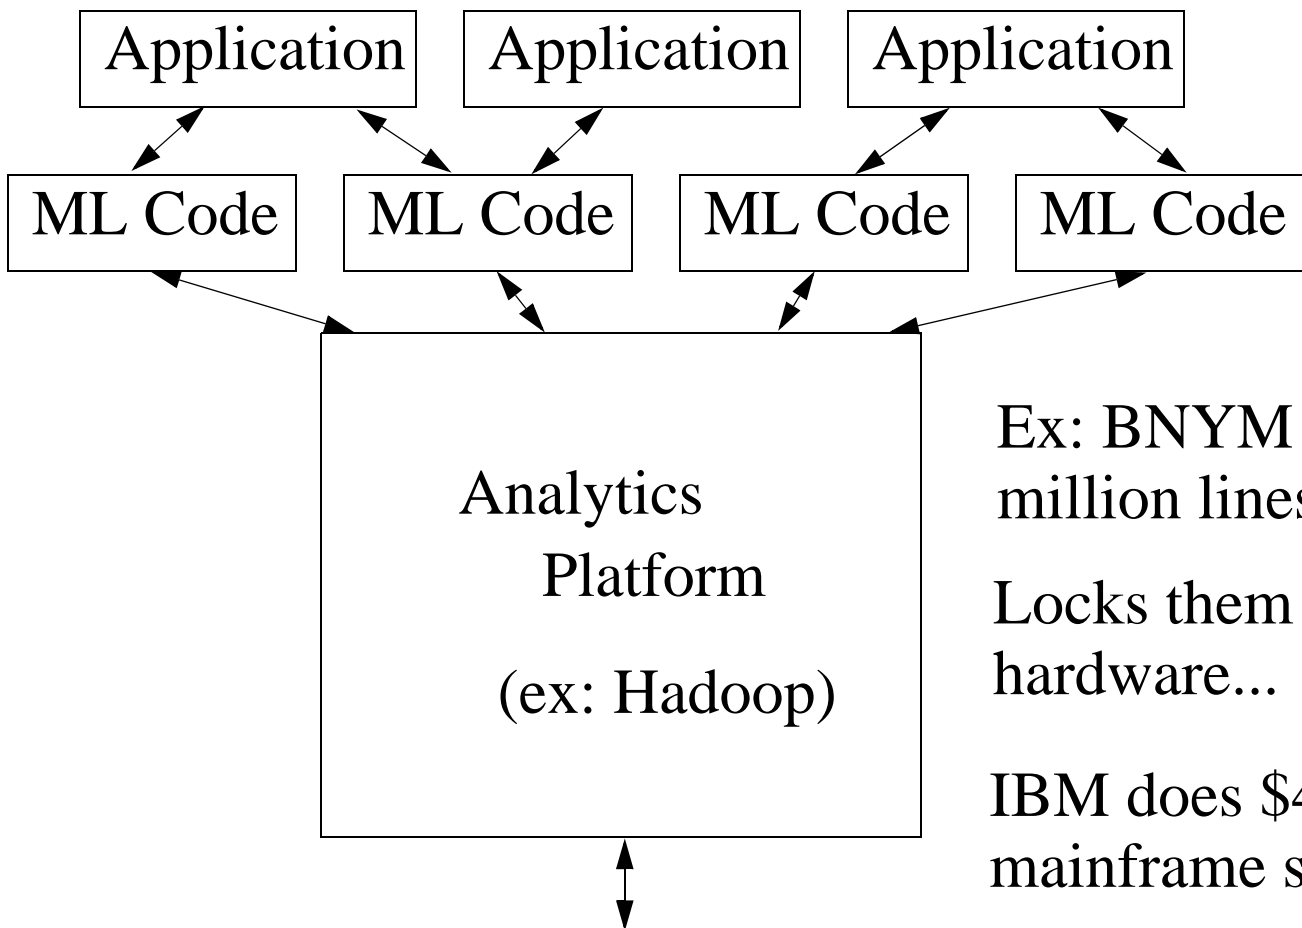
Declarative API: One code, many backends



This is the reason for downfall of network model/CODASYL




Declarative Analytics: One code, many backends



Ex: BNYM runs 343 million lines of COBOL

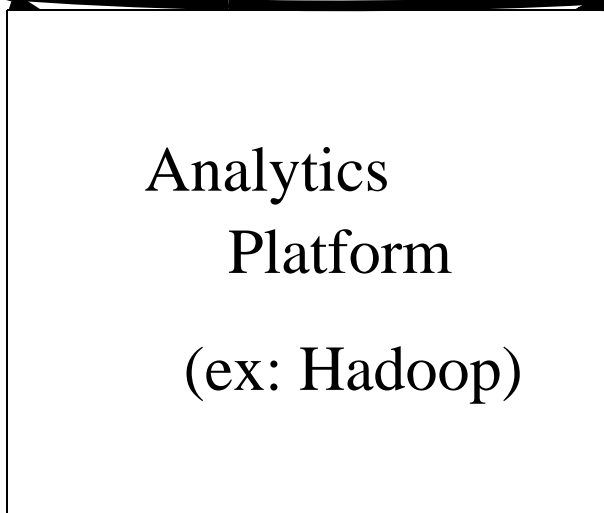

Locks them into 60's hardware...

IBM does \$4B plus in mainframe sales!



Declarative Alg 2: Procedural Form Algorithms

Declarative ML: From Custom ML Algorithms



Imagine many codes have

Map:

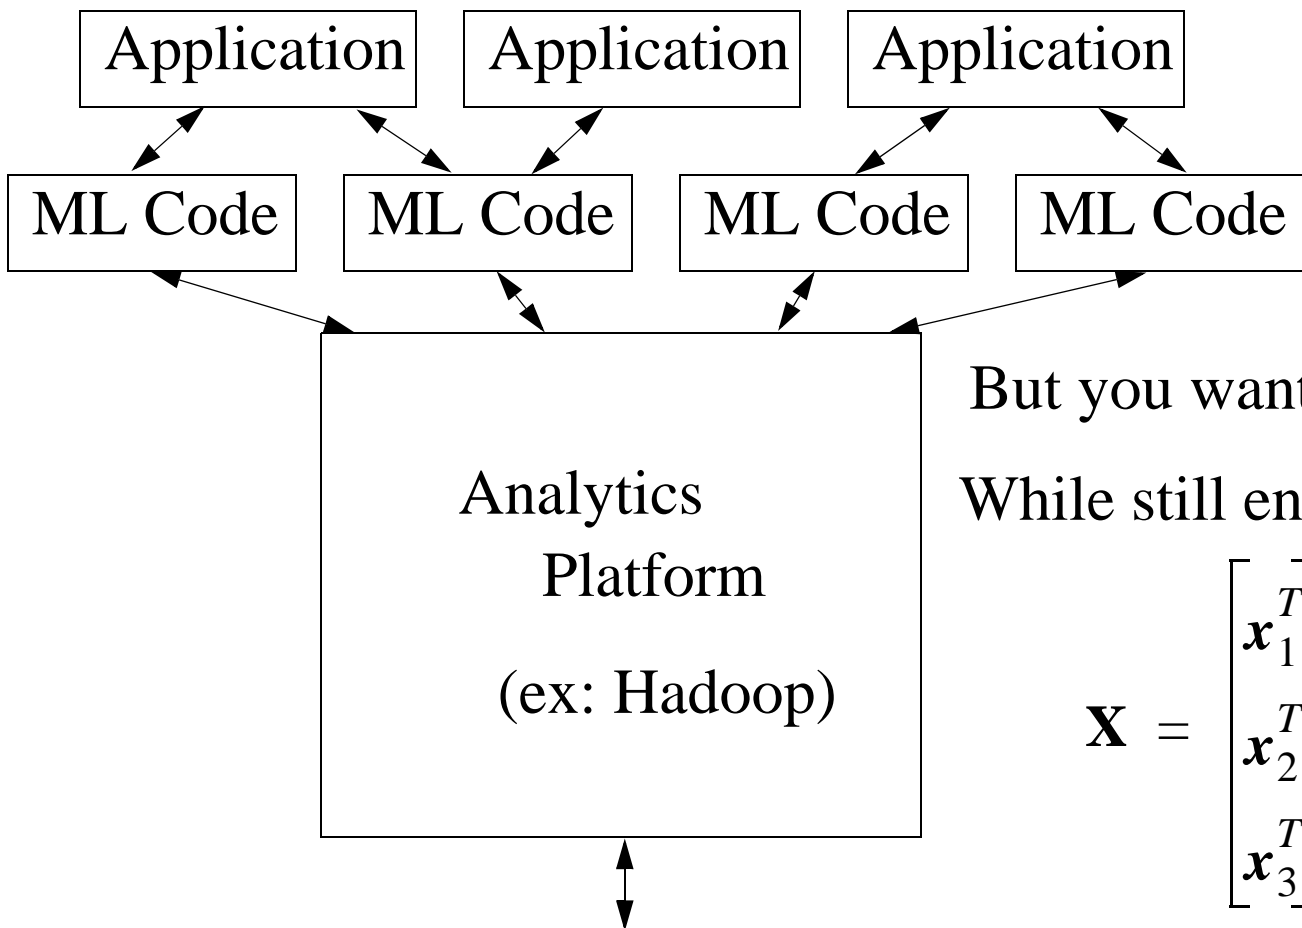
$$\mathbf{x} \rightarrow \mathbf{x}\mathbf{x}^T$$

Reduce:

$$\sum \mathbf{x}\mathbf{x}^T$$




Declarative Analytics: From Algorithms to Applications



But you want this:


While still enough RAM

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1^T \\ \mathbf{x}_2^T \\ \mathbf{x}_3^T \end{bmatrix}$$



Distributed compute cluster

Declarative Analytics: From Custom Algorithms




But you want this:

When RAM fills,

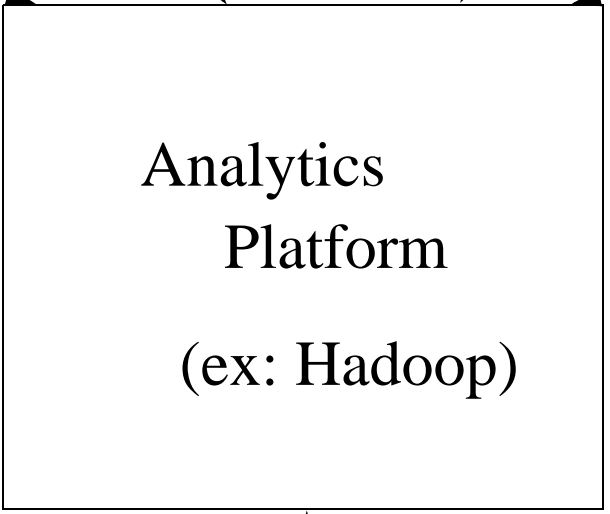

$$tot = tot + \mathbf{XX}^T$$

Output one *tot* from each machine and sum


Much faster!



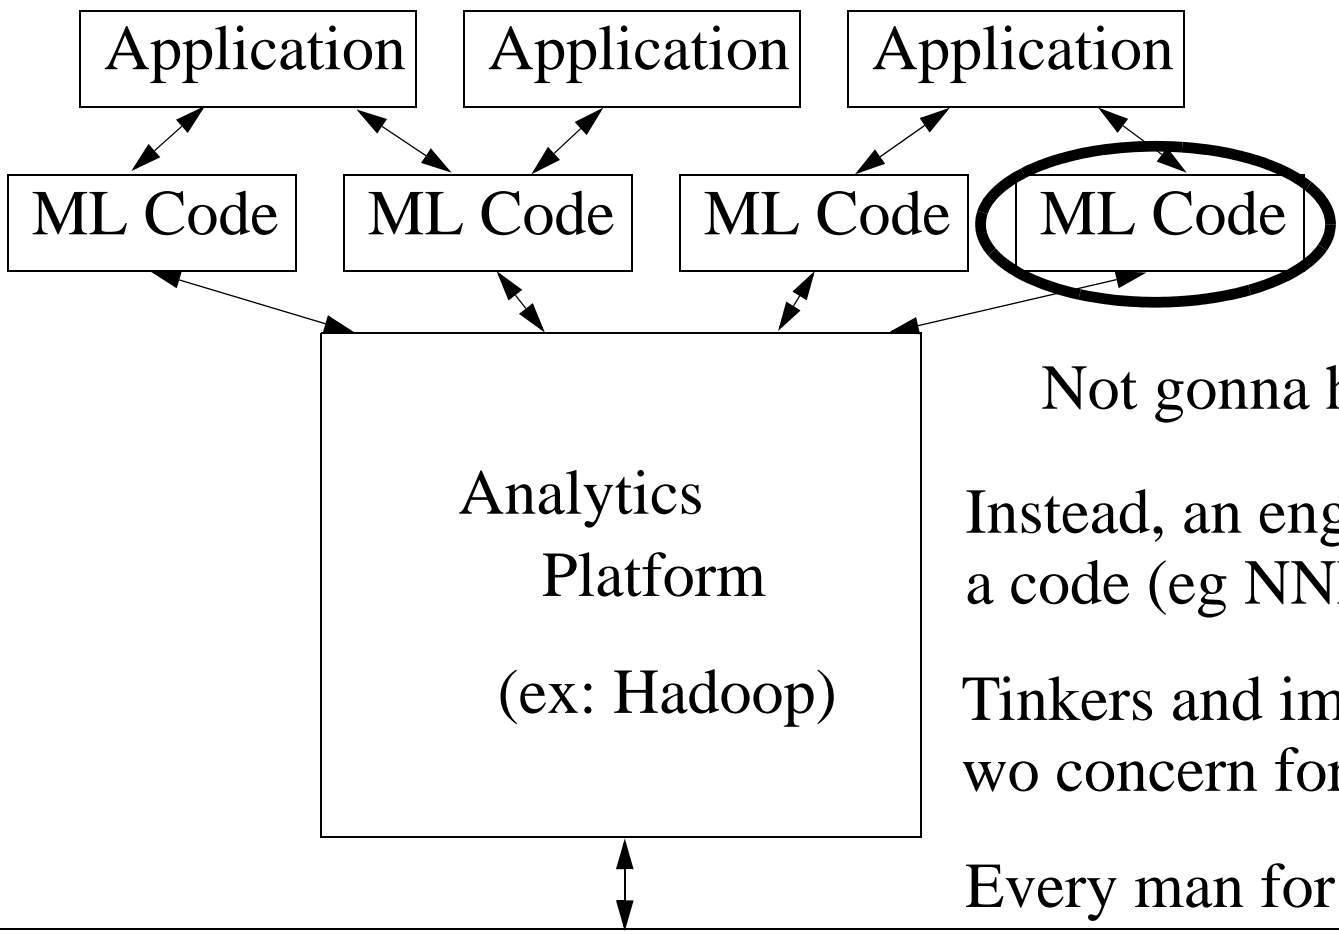
Declarative Analytics: From Algorithms



With non-declarative you need to search your code base for this pattern and refactor the code!



Declarative ML: Freedom from Algorithms



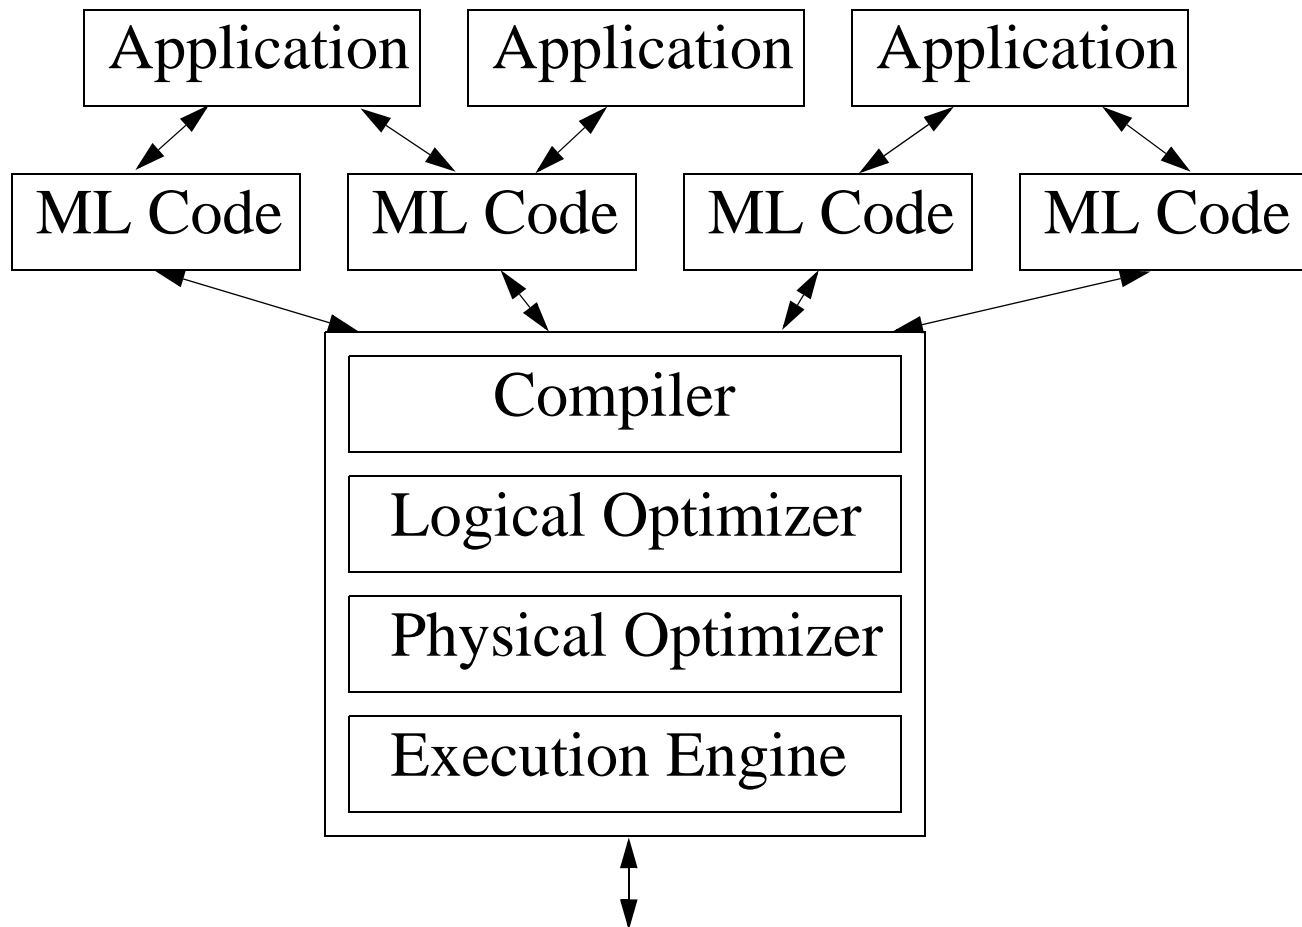
Not gonna happen...
Instead, an engineer owns
a code (eg NNMF)
Tinkers and improves it
w/o concern for other code
Every man for himself!



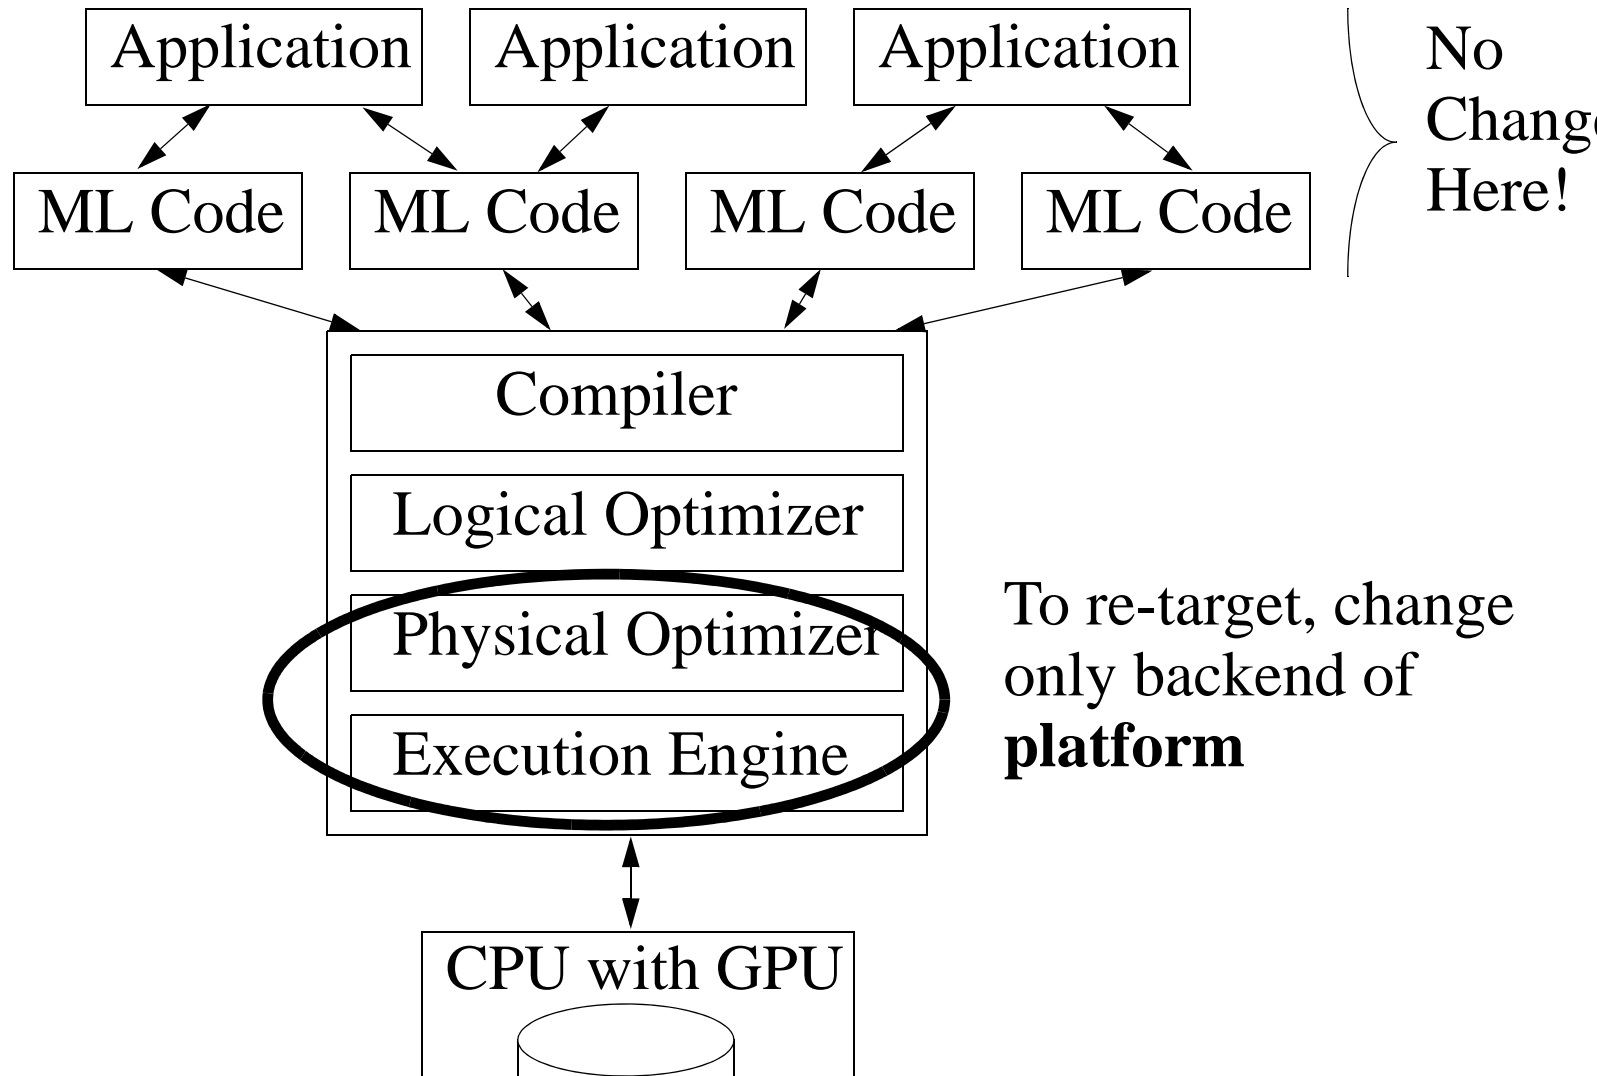
Distributed compute cluster

Machine Better in System Based on Data and:

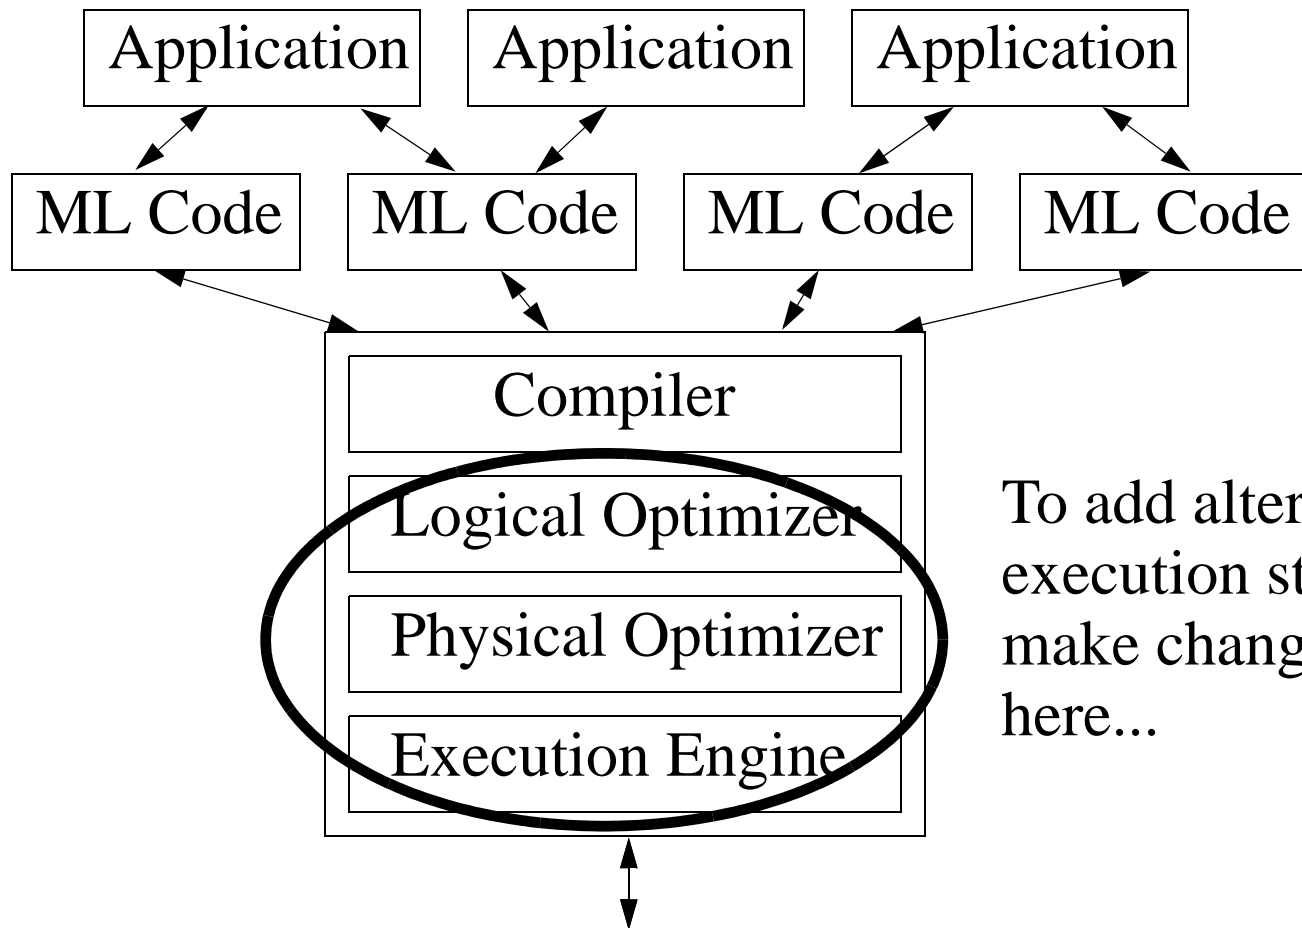
Machine Compiler in System Based on Data and ML



Machine Compiler in System Based on Data and



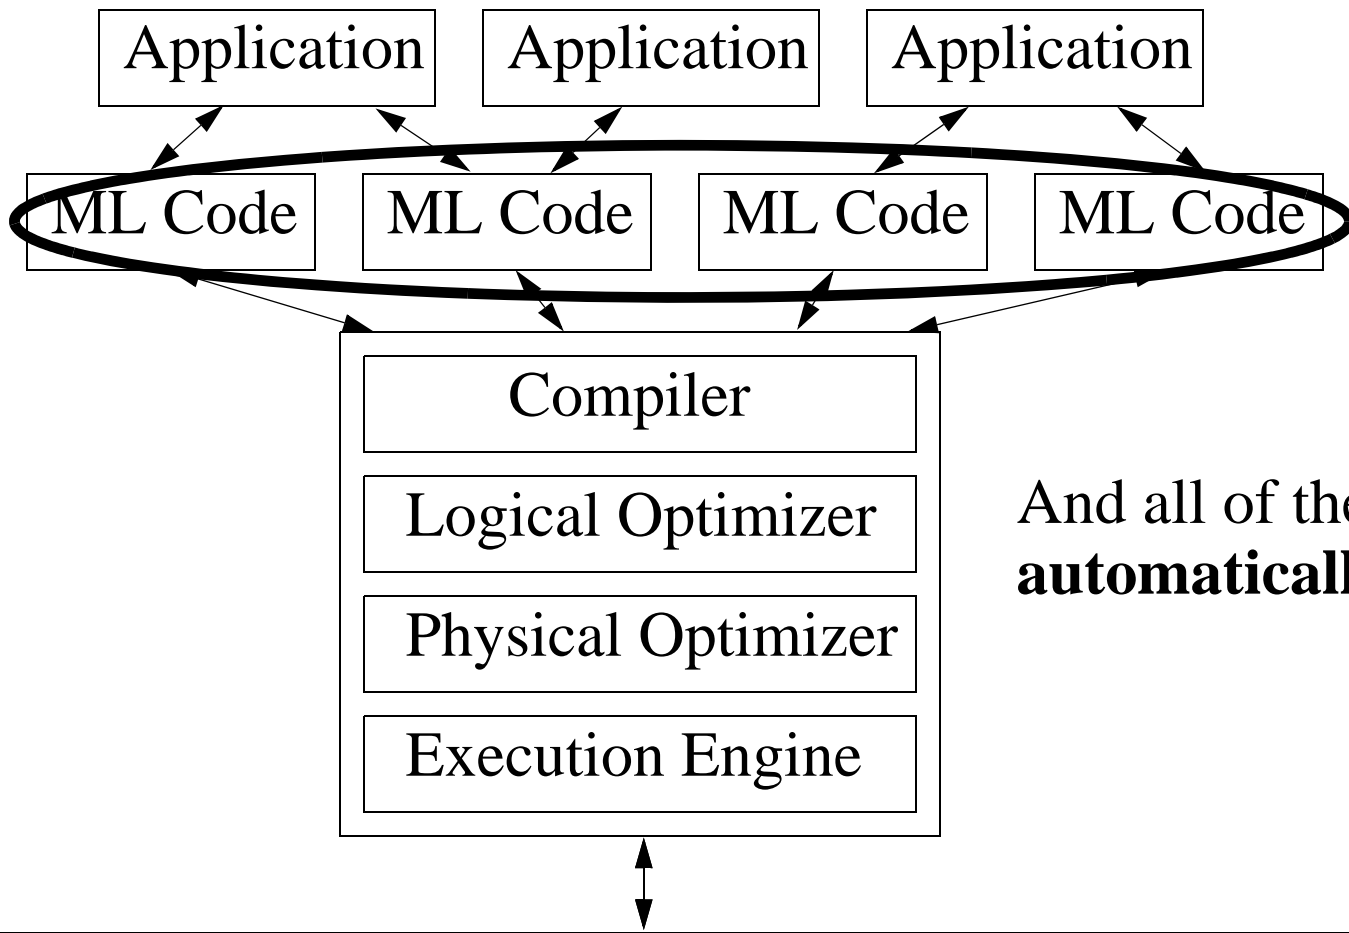
Machine Compiler in System Based on Data and ML



To add alternative execution strategies, make changes here...

Distributed compute cluster

Machine Compiler in System Based on Data and ML



And all of these codes **automatically** benefit

Distributed compute cluster

- Once written, code at top of stack can remain unchanged
 - #1 selling point of RDBMS tech for 30 years!

- Once written, code at top of stack can remain unchanged
 - #1 selling point of RDBMS tech for 30 years!
- Unfortunately, not accepted by ML community
 - ML people write great ML-on-GPU papers
 - They design platforms such as GraphLab (higher-level, MPI-like framework)

- Once written, code at top of stack can remain unchanged
 - #1 selling point of RDBMS tech for 30 years!
- Unfortunately, not accepted by ML community
 - ML people write great ML-on-GPU papers
 - They design platforms such as GraphLab (higher-level, MPI-like framework)
- Some in DB community have looked at declarative dataflow...
 - Spark SQL on Spark
 - Meteor on Stratosphere
 - Asterix from UC Irvine

- Once written, code at top of stack can remain unchanged
 - #1 selling point of RDBMS tech for 30 years!
- Unfortunately, not accepted by ML community
 - ML people write great ML-on-GPU papers
 - They design platforms such as GraphLab (higher-level, MPI-like framework)
- Some in DB community have looked at declarative dataflow...
 - Spark SQL on Spark
 - Meteor on Stratosphere
 - Asterix from UC Irvine
- But this is far from declarative ML
 - The codes don't look anything like math!

- Start with mathematical spec of learning algorithm...

1. $r \sim \text{Normal}(\mathbf{A}^{-1} \mathbf{X}^T \tilde{\mathbf{y}}, \sigma^2 \mathbf{A}^{-1})$

2. $\sigma^2 \sim \text{InvGamma}\left(\frac{(n-1) + p}{2}, \frac{(\tilde{\mathbf{y}} - \mathbf{X}r)^T (\tilde{\mathbf{y}} - \mathbf{X}r) + r^T \mathbf{D}^{-1} r}{2}\right)$

3. $\tau_j^{-2} \sim \text{InvGaussian}\left(\frac{\lambda \sigma}{r_j}, \lambda^2\right)$

— where $\mathbf{A} = \mathbf{X}^T \mathbf{X} + \mathbf{D}^{-1}$, $\mathbf{D}^{-1} = \text{diag}(\tau_1^{-2}, \tau_2^{-2}, \dots)$

This is math for the Bayesian Lasso, lifted from original paper

— Bayesian regression model with regularizing prior on regression coeffs

- Programmer writes code that looks just like the math...

```
data {
  n: range (responses); p: range (regressors);
  X: array[n, p] of real; y: array[n] of real;
  lam: real
}

var {
  sig: real;
  r, t: array[p] of real; yy, Z: array[n] of real;
}

A <- inv(X `* X + diag(t));
yy <- (y[i] - mean(y) | i in 1:n);
Z <- yy - X * r;

init {
  sig ~ InvGamma (1, 1);
  t ~ (InvGauss (1, lam) | j in 1:p);
}

r ~ Normal (A *' X * yy, sig * A);
sig ~ InvGamma(((n-1) + p)/2,
  (Z `* Z + (r * diag(t) `* r)) / 2);
for (j in 1:p) {
  t[j] ~ InvGauss (sqrt((lam * sig) / r[j]), lam);
```

We call our
language “BUDS”

- Write code that looks just like the math...

```
data {
  n: range (responses); p: range (regressors);
  X: array[n, p] of real; y: array[n] of real;
  lam: real
}
```

```
var {
  sig: real;
  r, t: array[p] of real; yy, Z: array[n] of real;
}
```

```
A <- inv(X `* X + diag(t)); ←  $\mathbf{A} = \mathbf{X}^T \mathbf{X} + \mathbf{D}^{-1}$ 
yy <- (y[i] - mean(y) | i in 1:n);
Z <- yy - X * r;
```

```
init {
  sig ~ InvGamma (1, 1);
  t ~ (InvGauss (1, lam) | j in 1:p);
}
```

```
r ~ Normal (A *' X * yy, sig * A); ←  $r \sim \text{Normal}(\mathbf{A}^{-1} \mathbf{X}^T \tilde{y}, \sigma^2 \mathbf{A}^{-1})$ 
sig ~ InvGamma(((n-1) + p)/2, ←  $\sigma^2 \sim \text{InvGamma}\left(\frac{(n-1) + p}{2}, \frac{(\tilde{y} - \mathbf{X}r)^T (\tilde{y} - \mathbf{X}r)^T + r^T \mathbf{D} r}{2}\right)$ 
  (Z `* Z + (r * diag(t) `* r)) / 2);
for (j in 1:p) {
  t[j] ~ InvGauss (sqrt((lam * sig) / r[j]), lam); ←  $\tau_j^{-2} \sim \text{InvGaussian}\left(\frac{\lambda \sigma}{r[j]}, \lambda^2\right)$ 
```


- And the system compiles and executes this for a huge data set
 - On hundreds or thousands of machines...
 - Or on a desktop with a GPU...
 - Or for whatever backend the system can target...

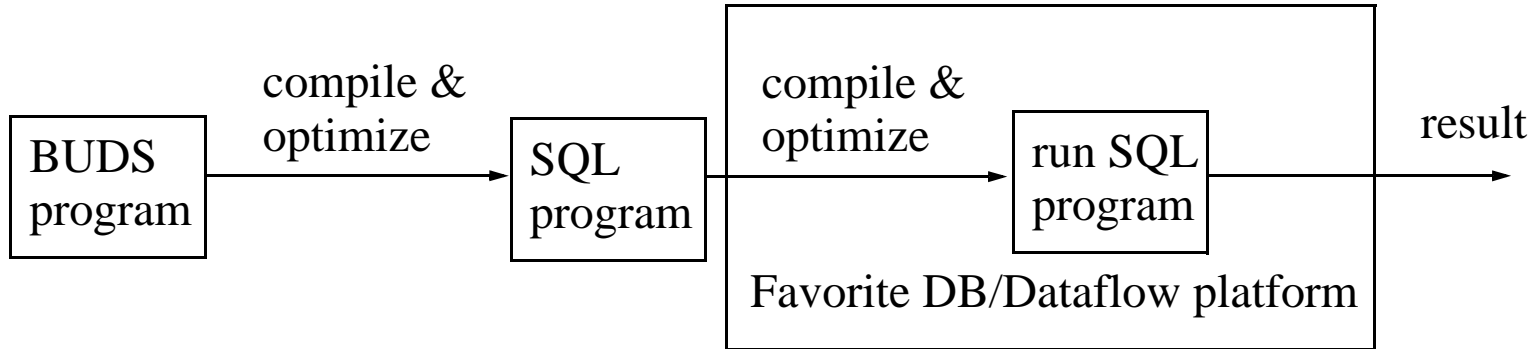
Also important

- We don't want to be like everyone and argue for a new DA stack
 - The world has too many dataflow platforms already

- SQL or SQL-like declarative language is the LCD for all systems
 - Including commercial databases
 - Free databases
 - And newfangled dataflow platforms

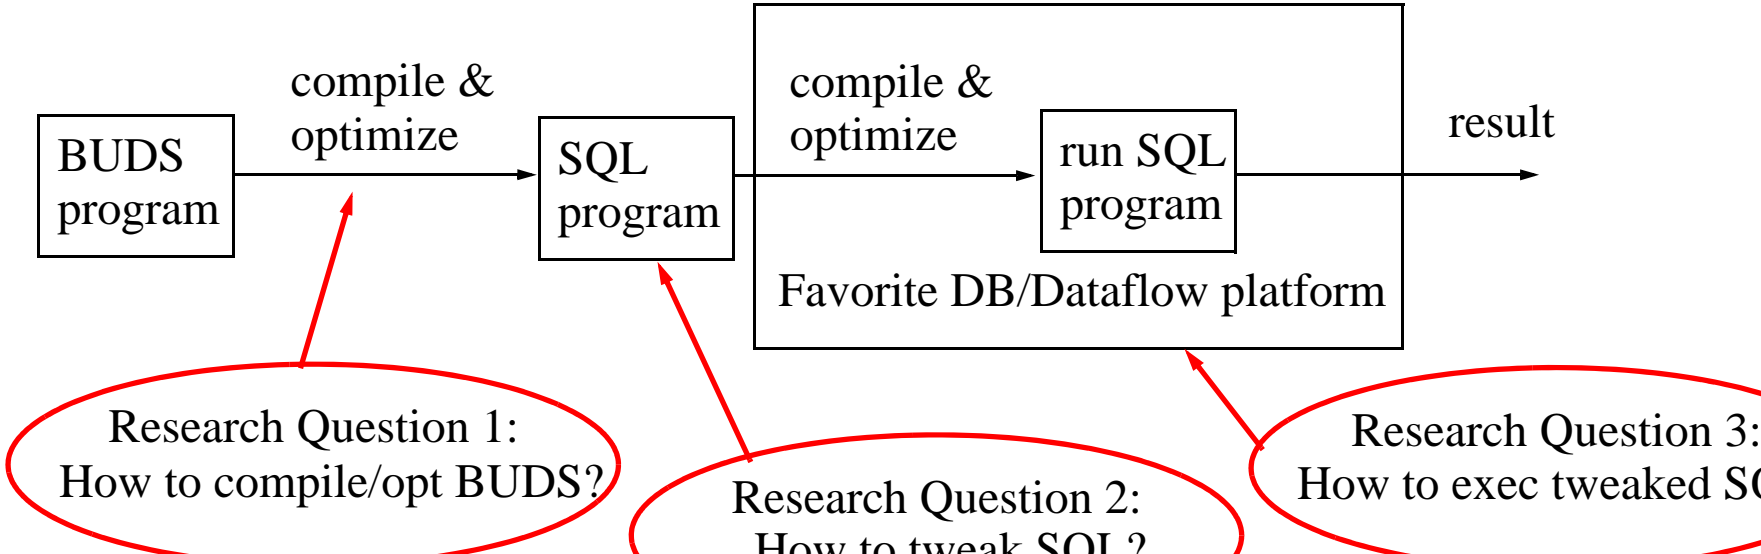
- SQL or SQL-like declarative language is the LCD for all systems
 - Including commercial databases
 - Free databases
 - And newfangled dataflow platforms

So here's the workflow we envision...



- SQL or SQL-like declarative language is the LCD for all systems
 - Including commercial databases
 - Free databases
 - And newfangled dataflow platforms

So here's the workflow we envision...



QUESTIONS: