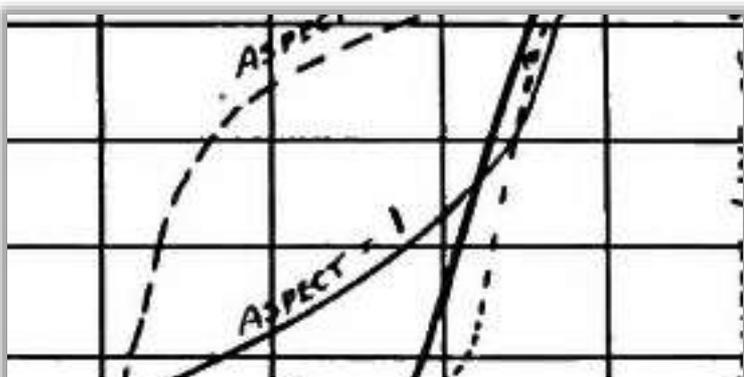




# **ZØ: An Optimizing Distributing Zero-Knowledge Compiler**

**Matt Fredrikson**  
University of Wisconsin

**Ben Livshits**  
Microsoft Research

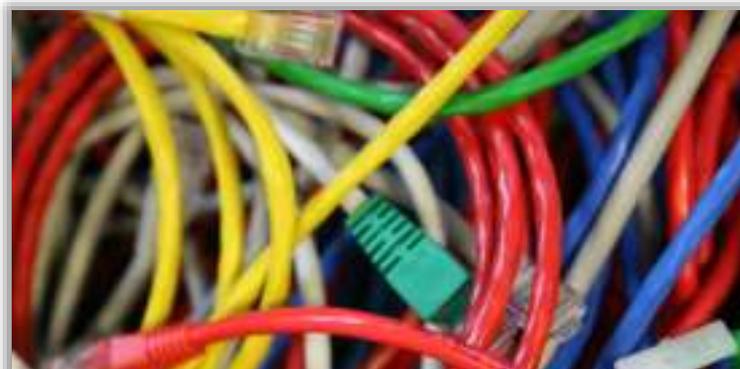


## Automatic Optimization

Cost modeling enables huge optimization opportunities

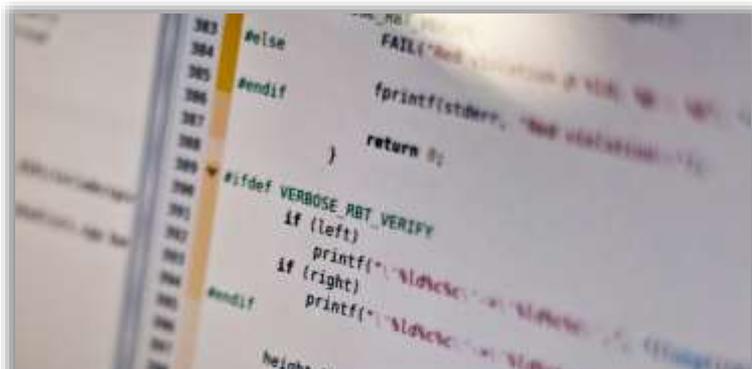
.30 .35 .40 .45

# This talk: at a glance



## Distributed Environments

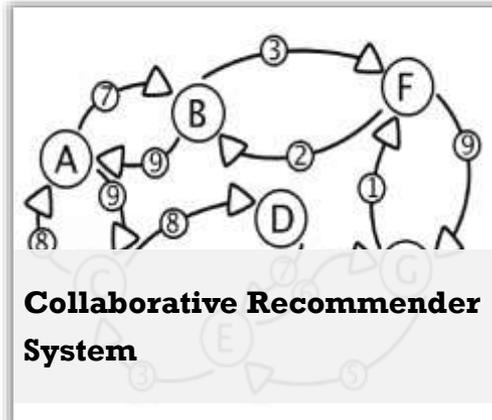
ZØ automatically places code on different computational tiers

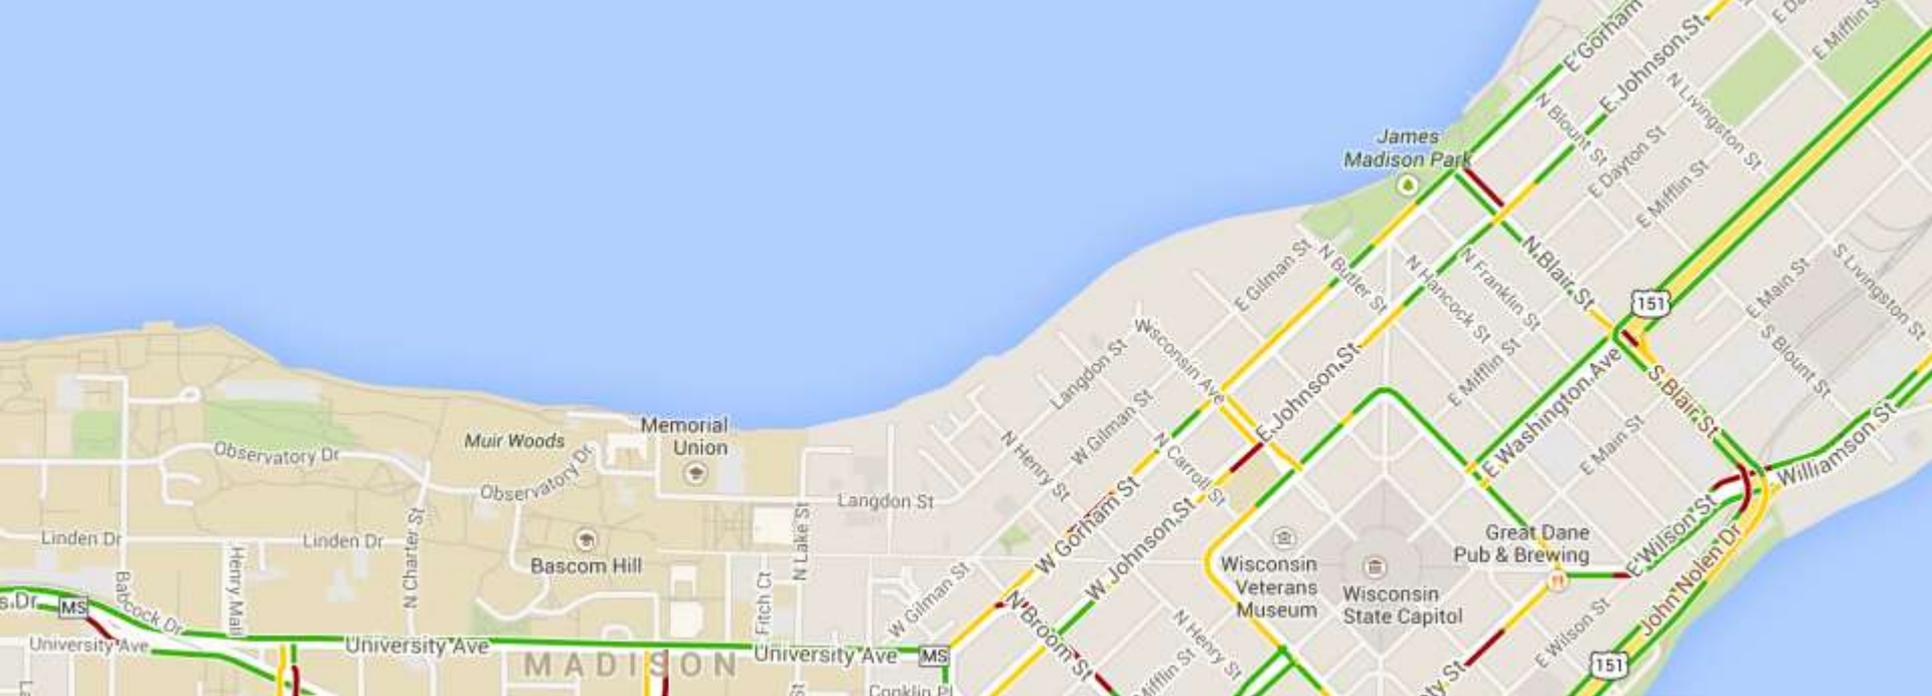


## “Zero-Knowledge for the Masses”

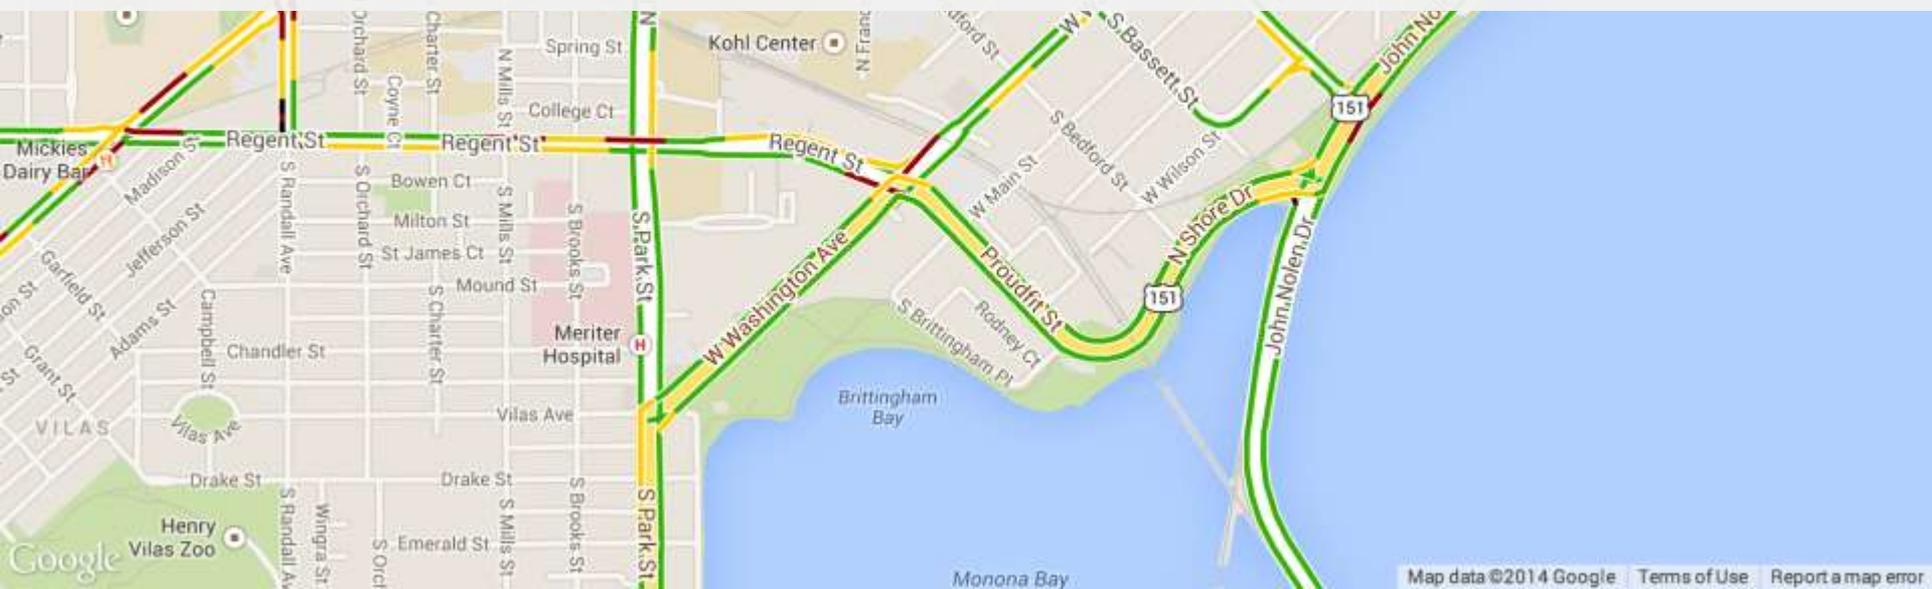
Users write ZK code in C#, as one part of a larger project

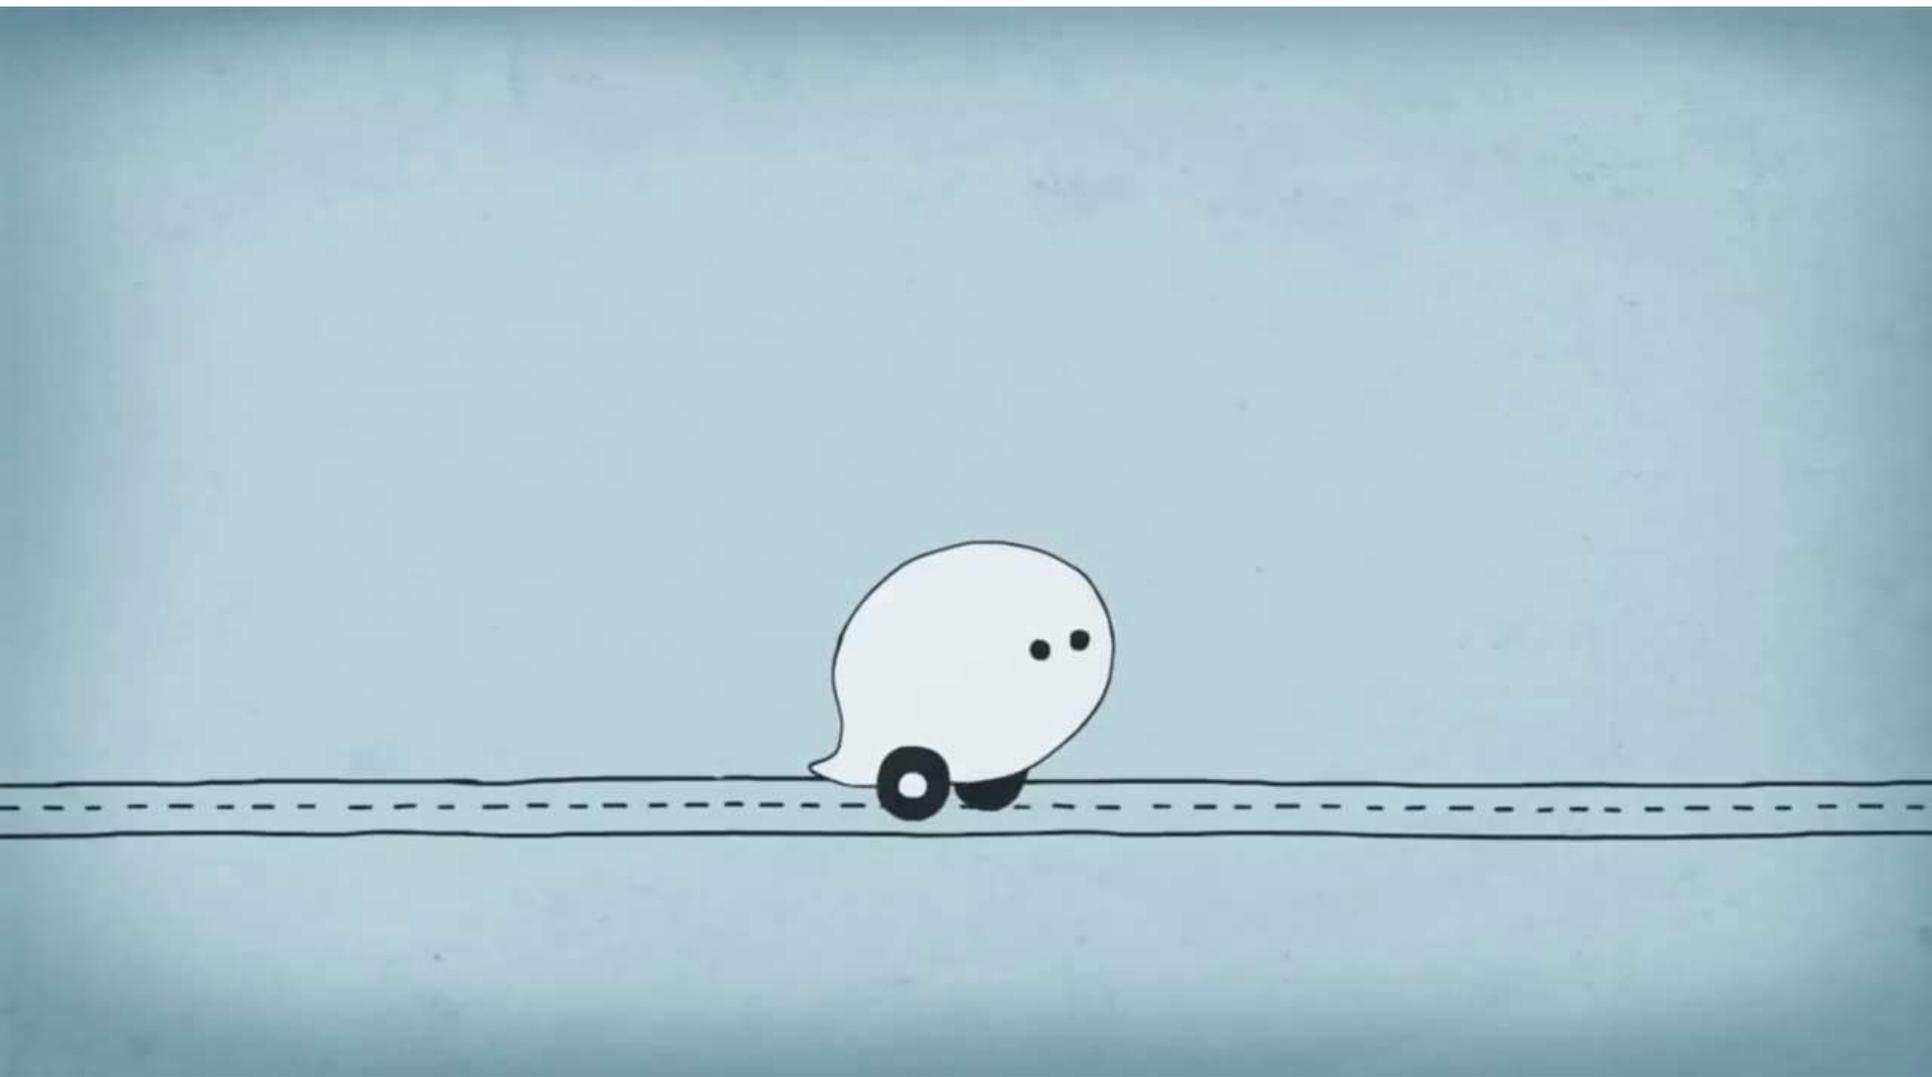
# This talk: at a glance

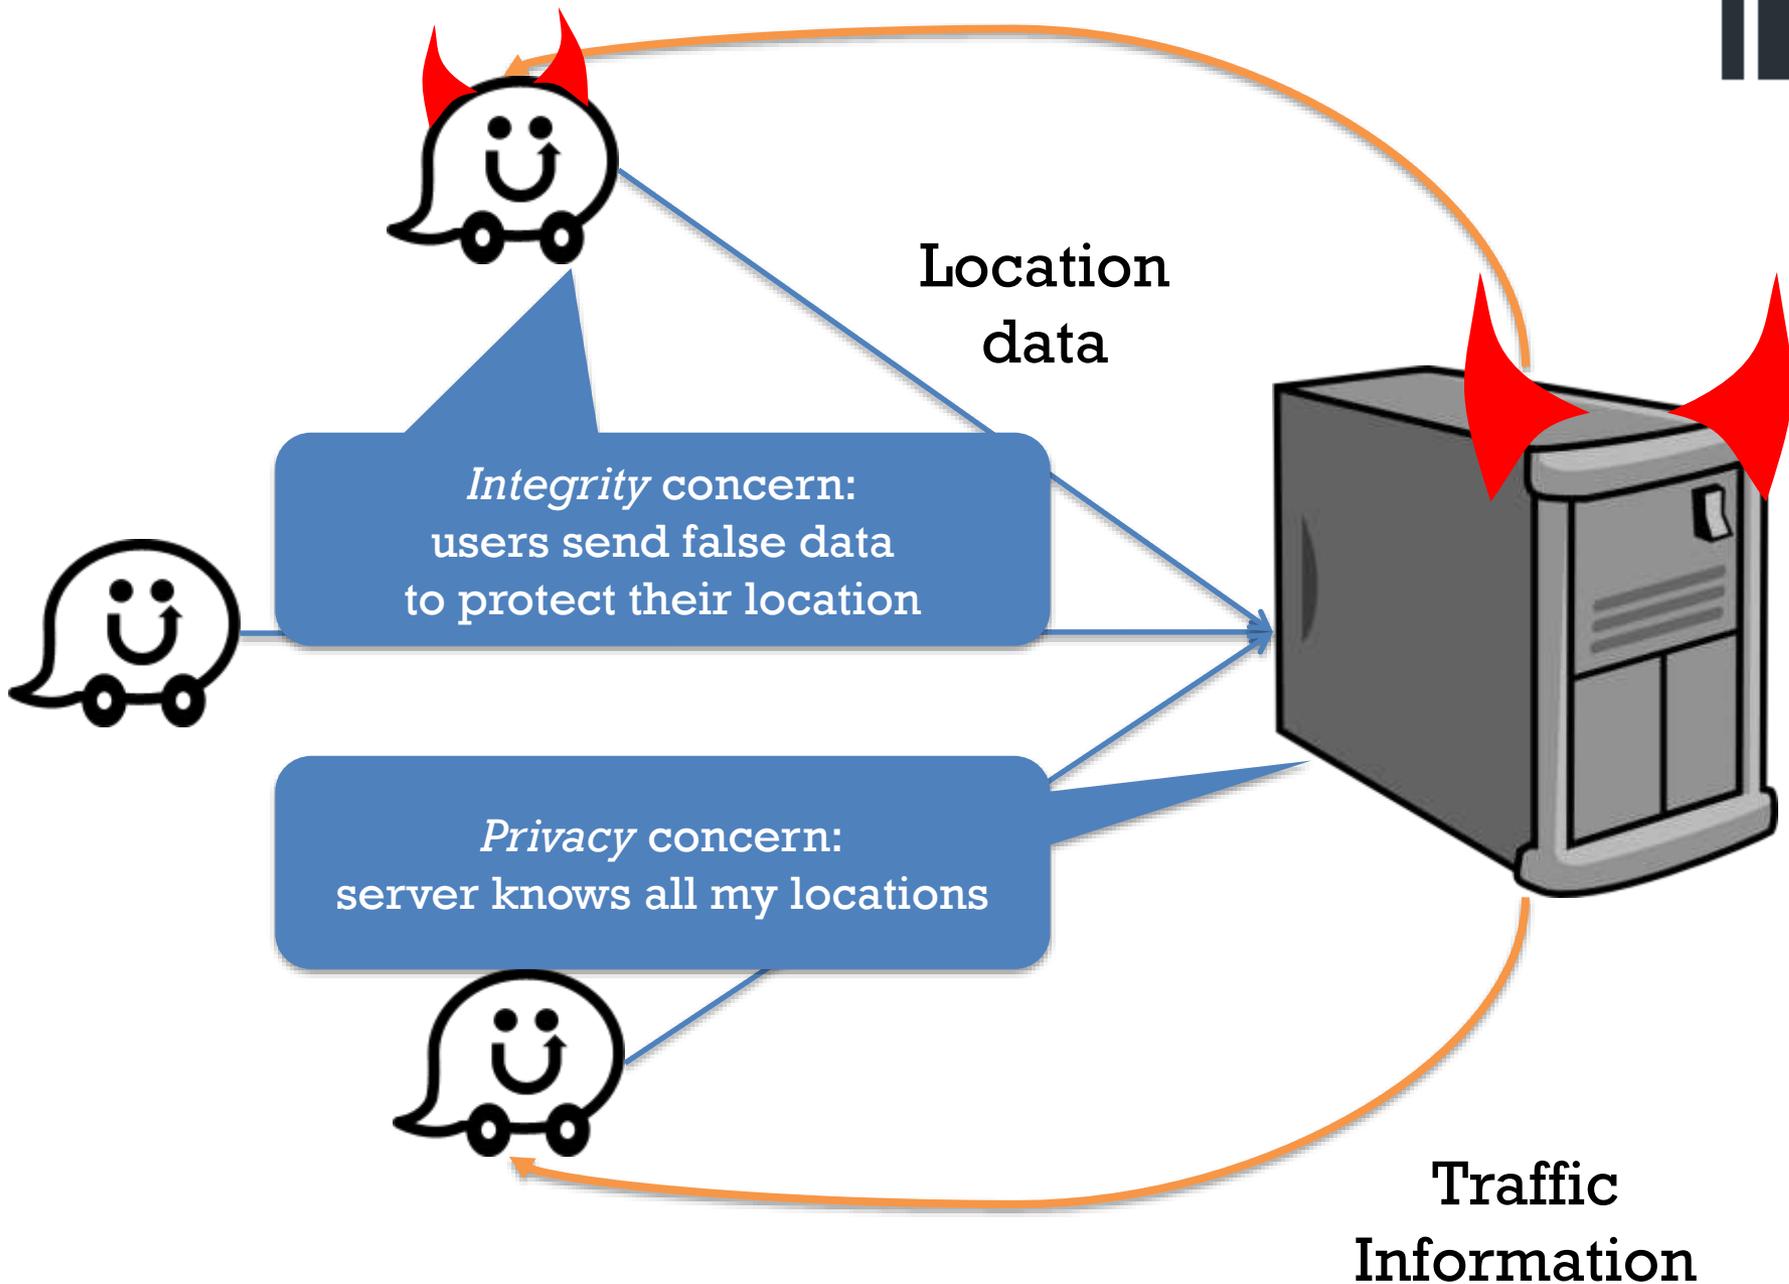


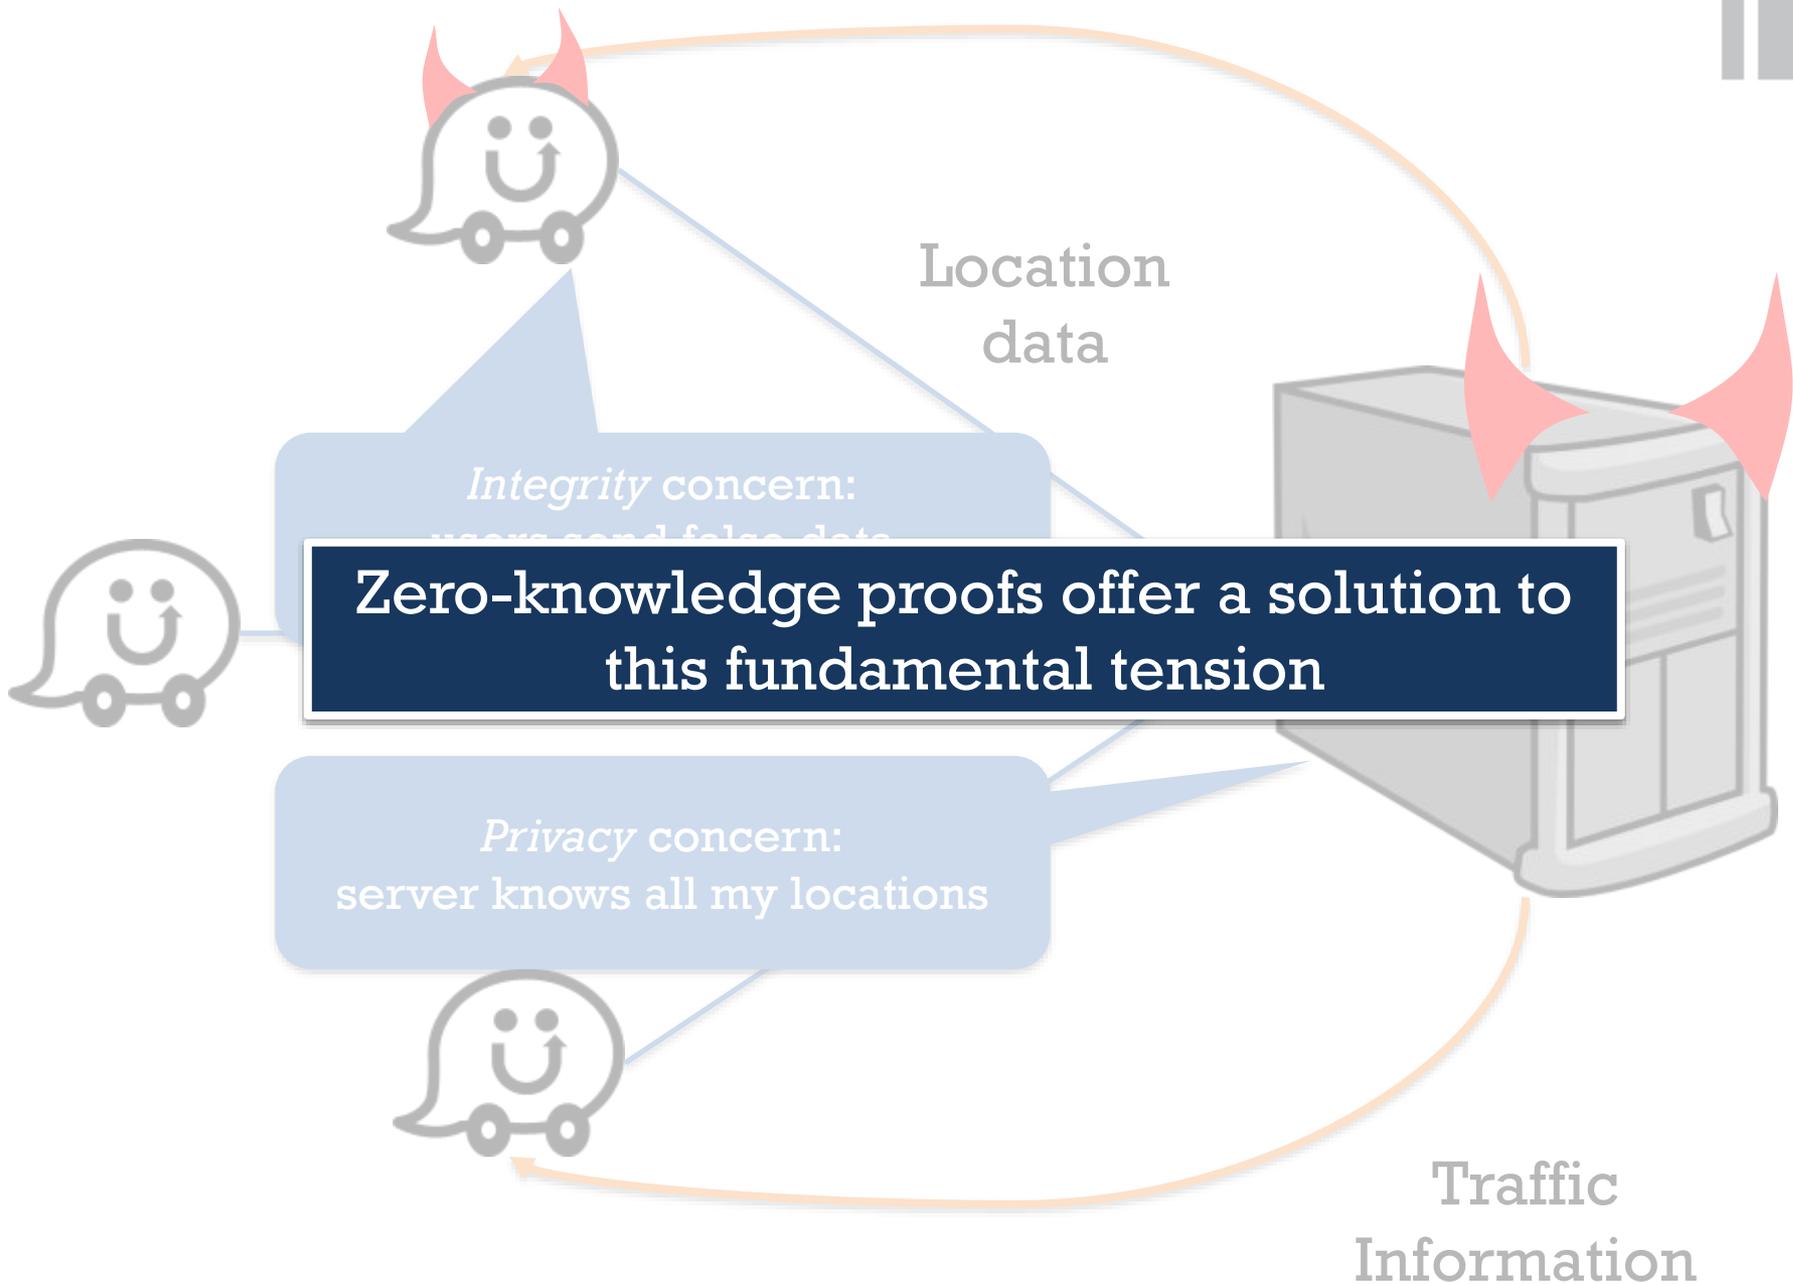


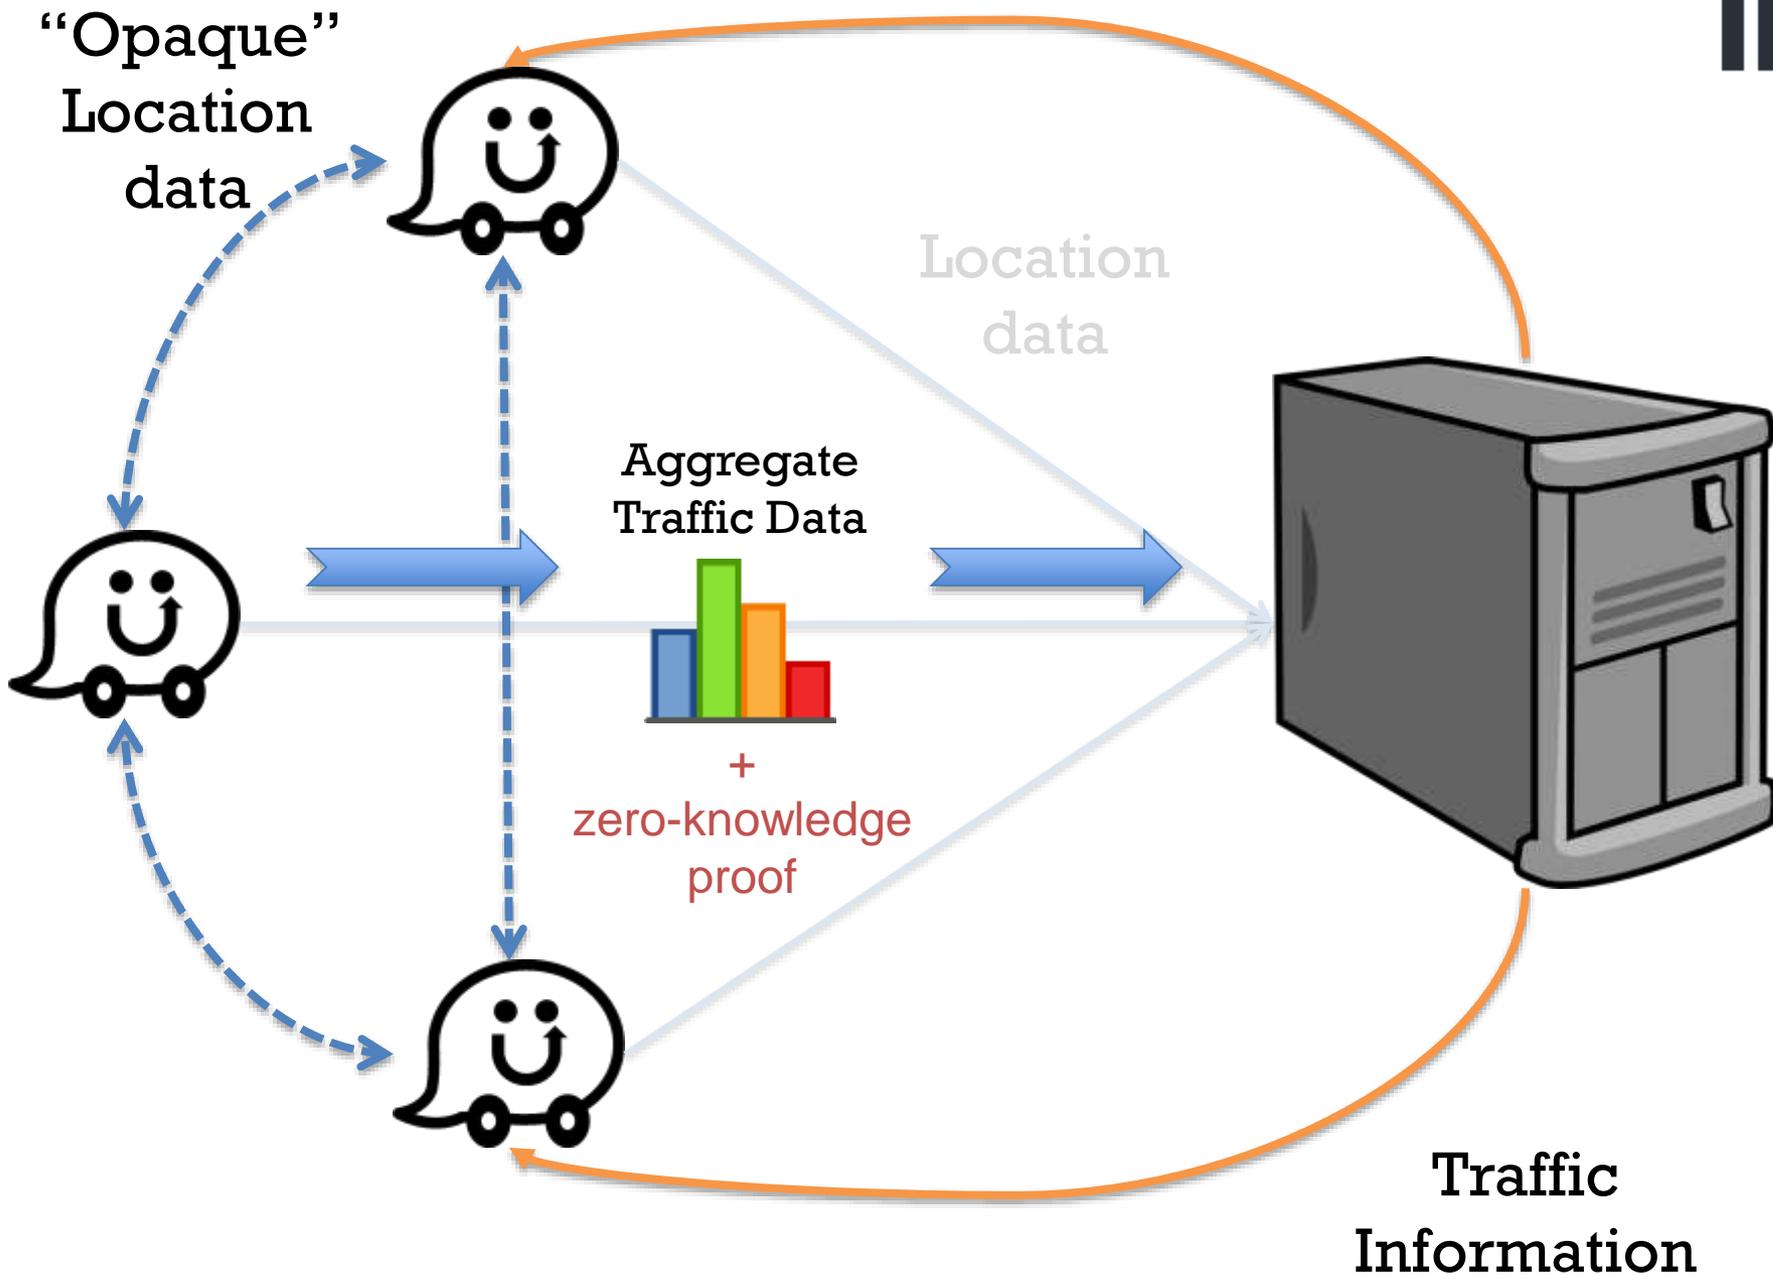
# Crowd-sourced traffic maps











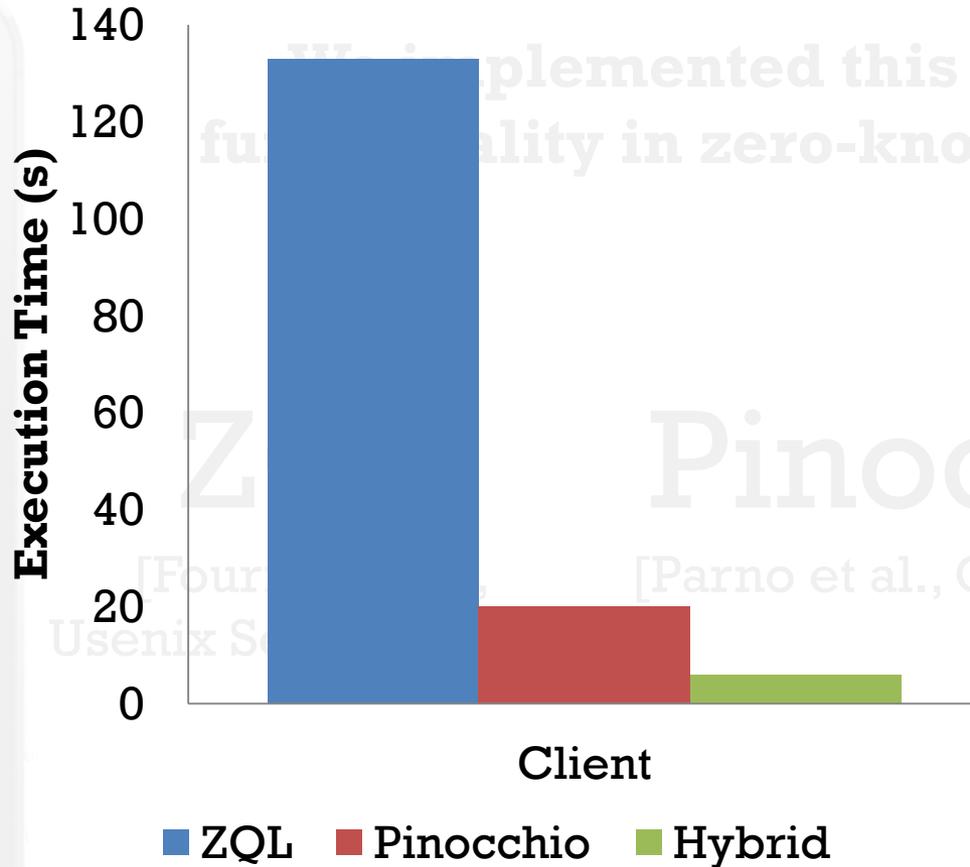
Partition roads into segments

Use *Shamir shares* on segment IDs



# Initial Experiments

## Client Time to Process a GPS Reading



Pinocchio

[Parno et al., Oakland 2013]

[Four  
Usenix S

implemented this core  
quality in zero-knowledge

# Why Such a Contrast?

**These zero-knowledge “back-ends” have significantly different execution models**

## ZQL

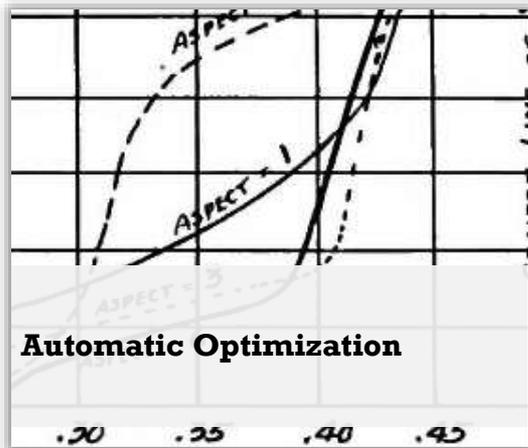
Compiles specialized language to F#, then CIL

## Pinocchio

Compiles C to a fixed circuit representation

# ZØ: An Optimizing Compiler for ZK

ZØ uses the best of both back-ends as appropriate for the application at hand



# ZØ: An Optimizing Compiler for ZK



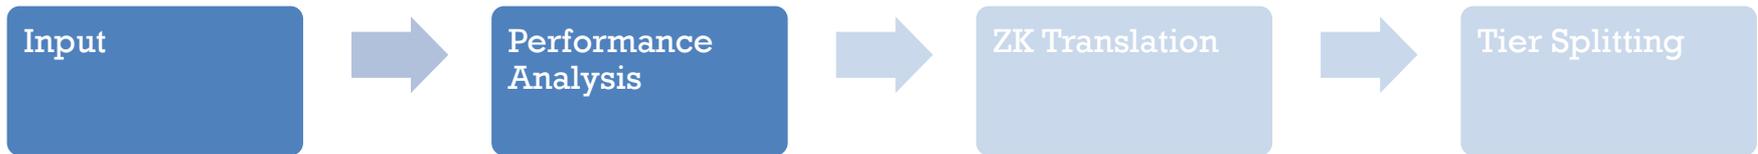
Users write code in C#

```
// Get the list of user names
var userNames = gpsPts.Select(loclist => loclist.Aggregate( , (acc, el) => el.User));
var squared = gpsQuads.Select(quad => (quad.nth<int>(3, 4) - quad.nth<int>(1, 4)) * (quad.nth<int>(3, 4) + quad.nth<int>(1, 4)));

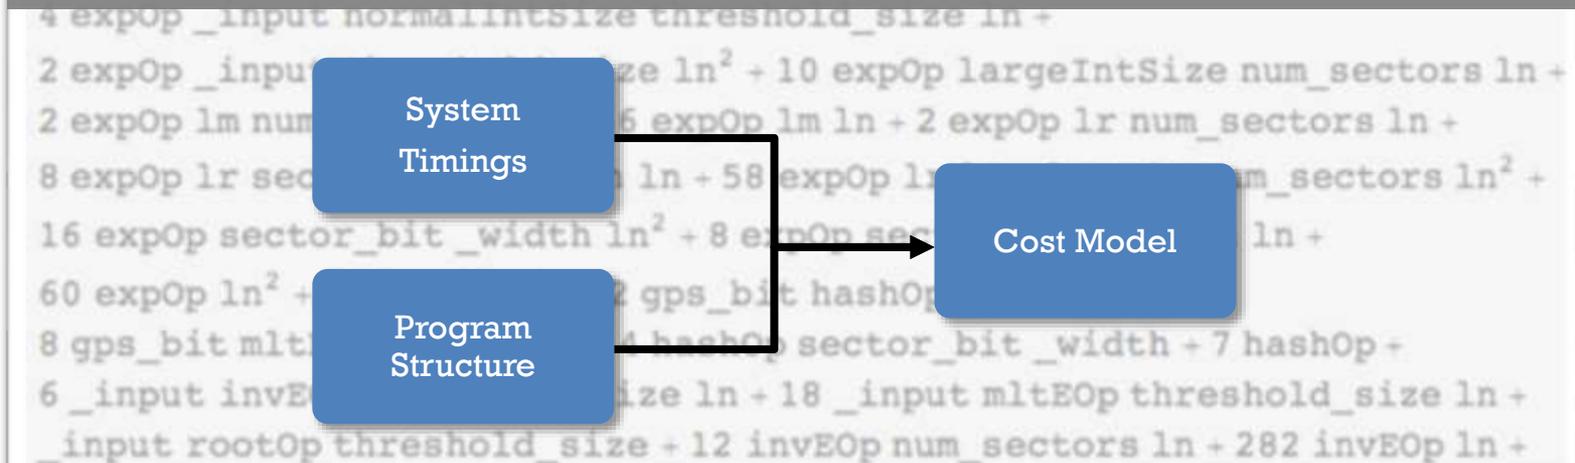
// Now do the zero-knowledge aggregation
// Get the sum of the squared incremental distances
ZeroKnowledgeBegin();
var distlisttup = squared.Select(sq => sqrtTable.First(row => row.nth<int>(1,2) == sq));
var distlist = distlisttup.Select(dltup => dltup.nth<int>(2, 2) * onevalue);
var distval = distlist.Aggregate(1, (acc, sum) => acc + sum);
ZeroKnowledgeEnd();

// Now wrap the result in a Distance type
// This is not going to be zero-knowledge
var dist = new List<int>() { distval };
var withUserNames = dist.Zip(userNames, (_dist, _user) => new { dist = _dist, user = _user });
var returnVal = new DistributedEnumerable<Distance>(withUserNames.Select(d => new Distance(d, user)));
```

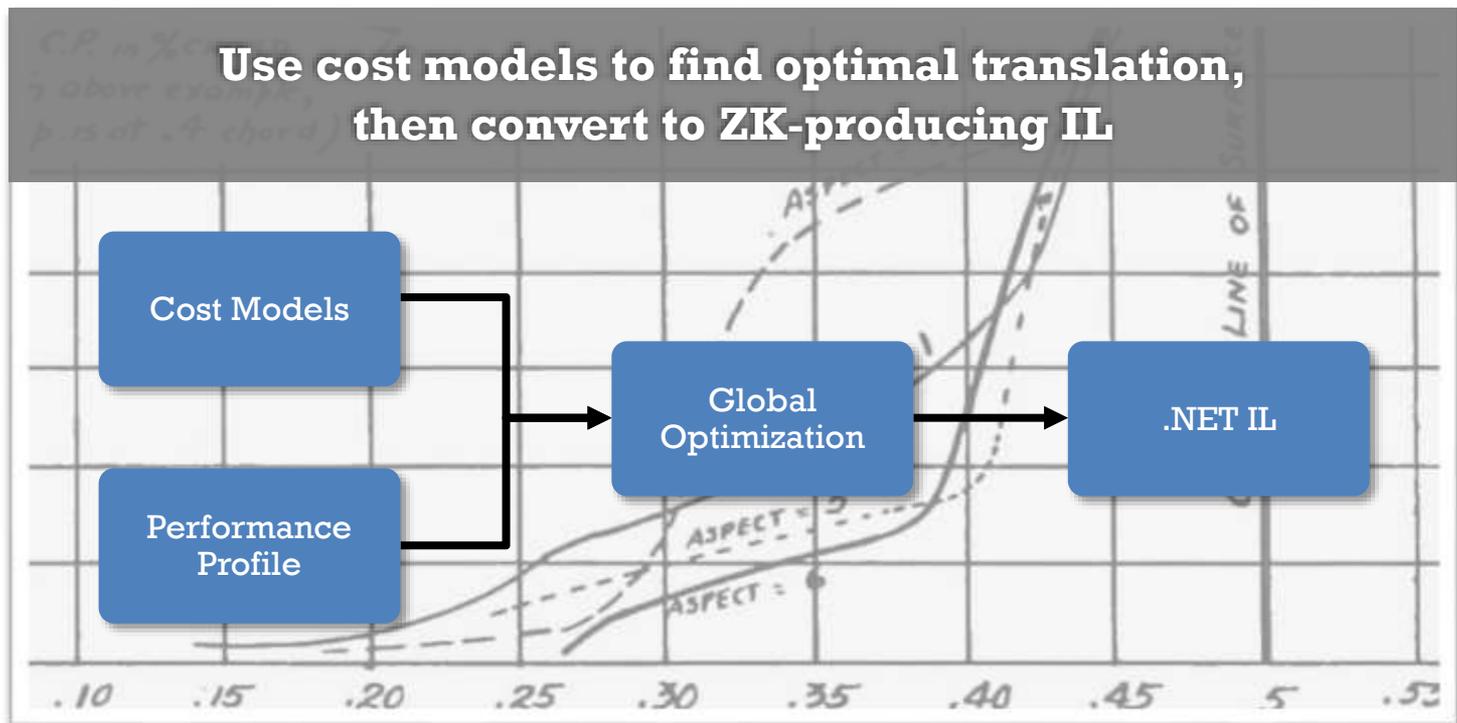
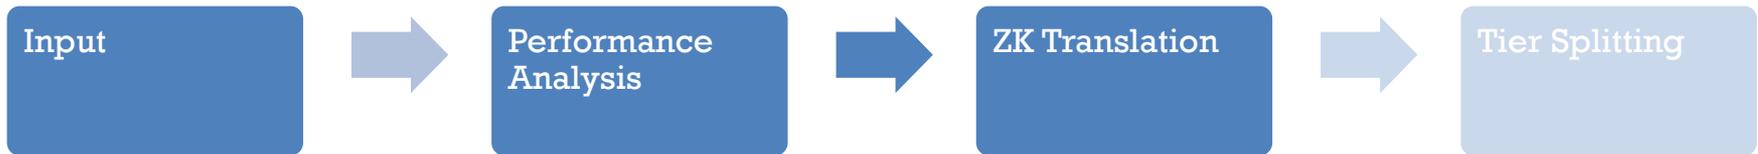
# ZØ: An Optimizing Compiler for ZK



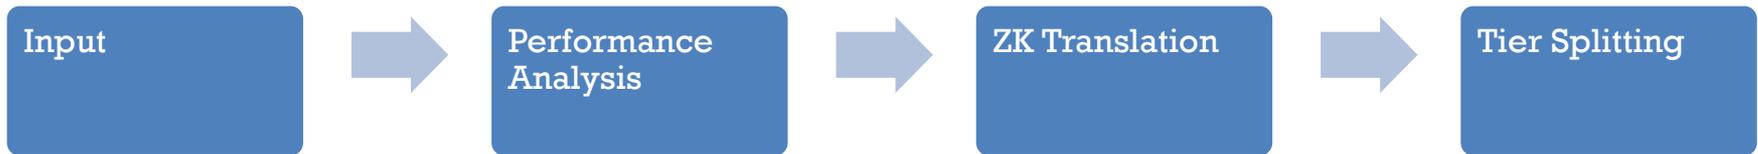
**Build detailed *cost models* that characterize how expensive C# will be when translated to zero-knowledge**



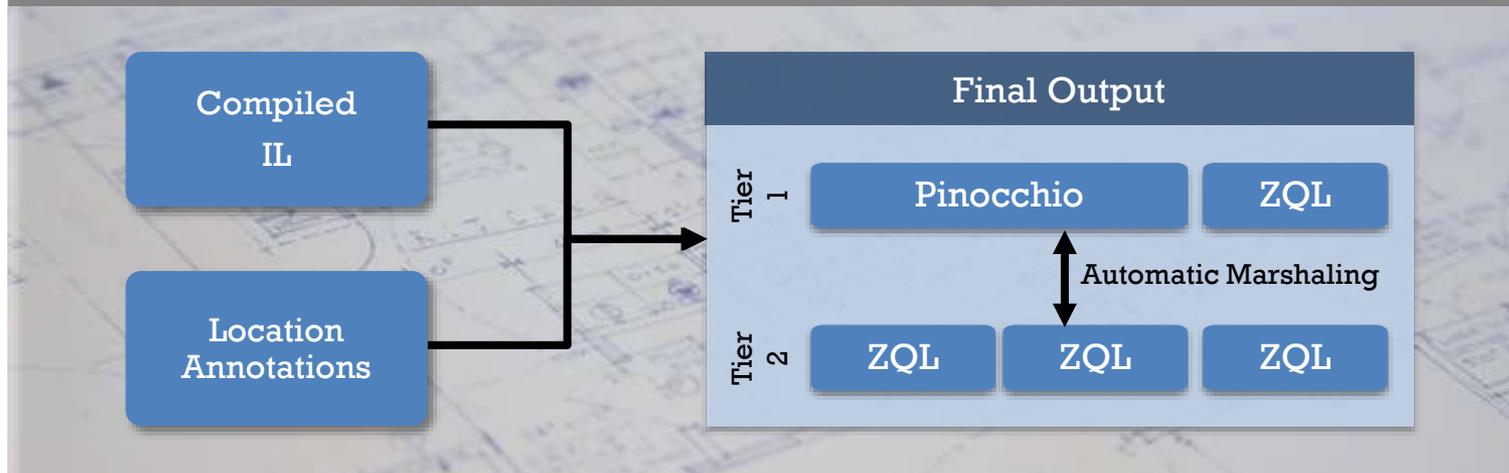
# ZØ: An Optimizing Compiler for ZK



# ZØ: An Optimizing Compiler for ZK



Use location annotations to split IL between tiers, insert automatic data transfer and synchronization



# ZERO-KNOWLEDGE IN C#

# Zero-Knowledge in C#

Location annotations drive tier-splitting

```
[ExecutionLocation(ExecutionLocationValue.Client)]  
private DistributedEnumerable<ShareValue> AggregateGpsReadings
```

```
(  
    [MaximumInputSize(1000)] DistributedEnumerable<ShareValue> shares  
)  
{
```

Specify ZK input sizes to help optimization

```
    var svalues = shares.Select(share => share.Value);
```

```
    ZeroKnowledgeBegin();
```

Programmers specify ZK regions

```
    int aggShare = svalues.Aggregate(0, (acc, share) => acc + share);  
    ZeroKnowledgeEnd();
```

```
    return new DistributedEnumerable<ShareValue>
```

ZK operations given by LINQ expressions

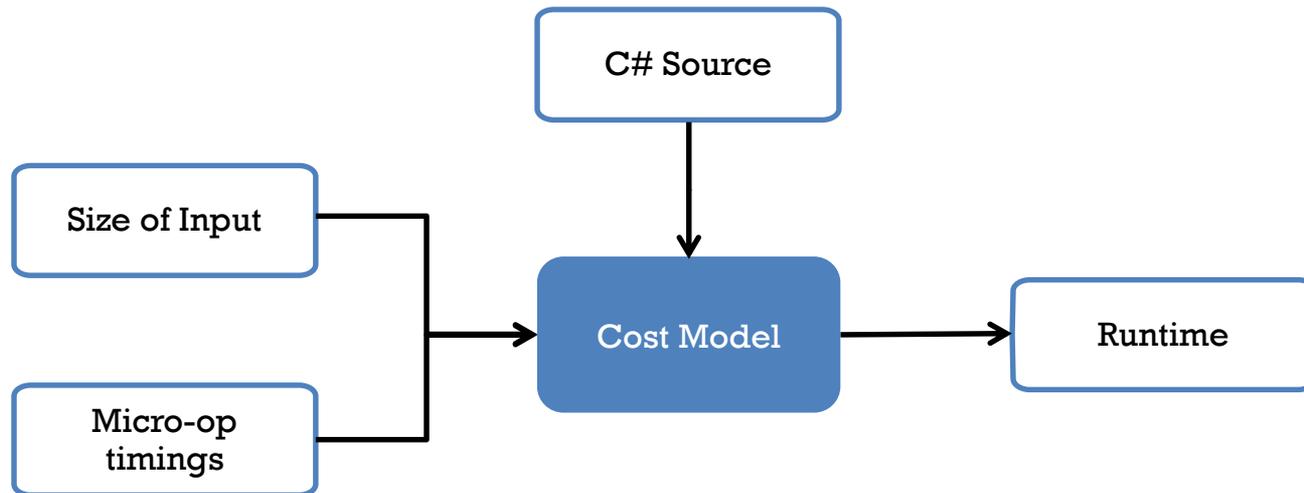
```
    (  
        new List<ShareValue>() { new ShareValue(aggShare, myEvalPoint) }  
    );
```

```
}
```

# **COST MODELING**

# Cost Models for Optimization

Cost models characterize the ZK runtime of C# code



$$\mathbf{F}(\text{inputListSize}) = \text{eqOp} * \text{inputListSize} + \text{addOp} + 12 * \text{expOp} + 3 * \text{extendOp} + 14 * \text{mltOp}$$

size of input

micro-op timings

# Building a Cost Model

Symbolic evaluation  
over polynomial domain

## ZQL

`map, fold, find`  
expressions: we can  
always bound the  
number of ops in each  
expression

Static circuit evaluation  
polynomials

## Pinocchio

Given a circuit, we can  
determine evaluation  
and proof generation  
time

# **TRANSLATION & TIER SPLITTING**

# Translating C# to Zero-Knowledge

## Cost Models

$f(\text{inputListSize}) = \text{eqOp} * \text{inputListSize} + \text{addOp} + \dots$

$f(\text{numPeers}) = \text{addOp} * \text{numPeers} + \text{multOp} + \dots$

$f(\text{numItems}) = \text{multOp} * \text{numItems} + \text{eqOp} + \dots$

## Performance Profile

Tier	Compute Cost	Transfer Cost
Mobile	2	3
Server	0.5	1
...	...	...

## Global Optimization

## .NET IL

Pinocchio

ZQL

ZQL

Pinocchio

ZQL

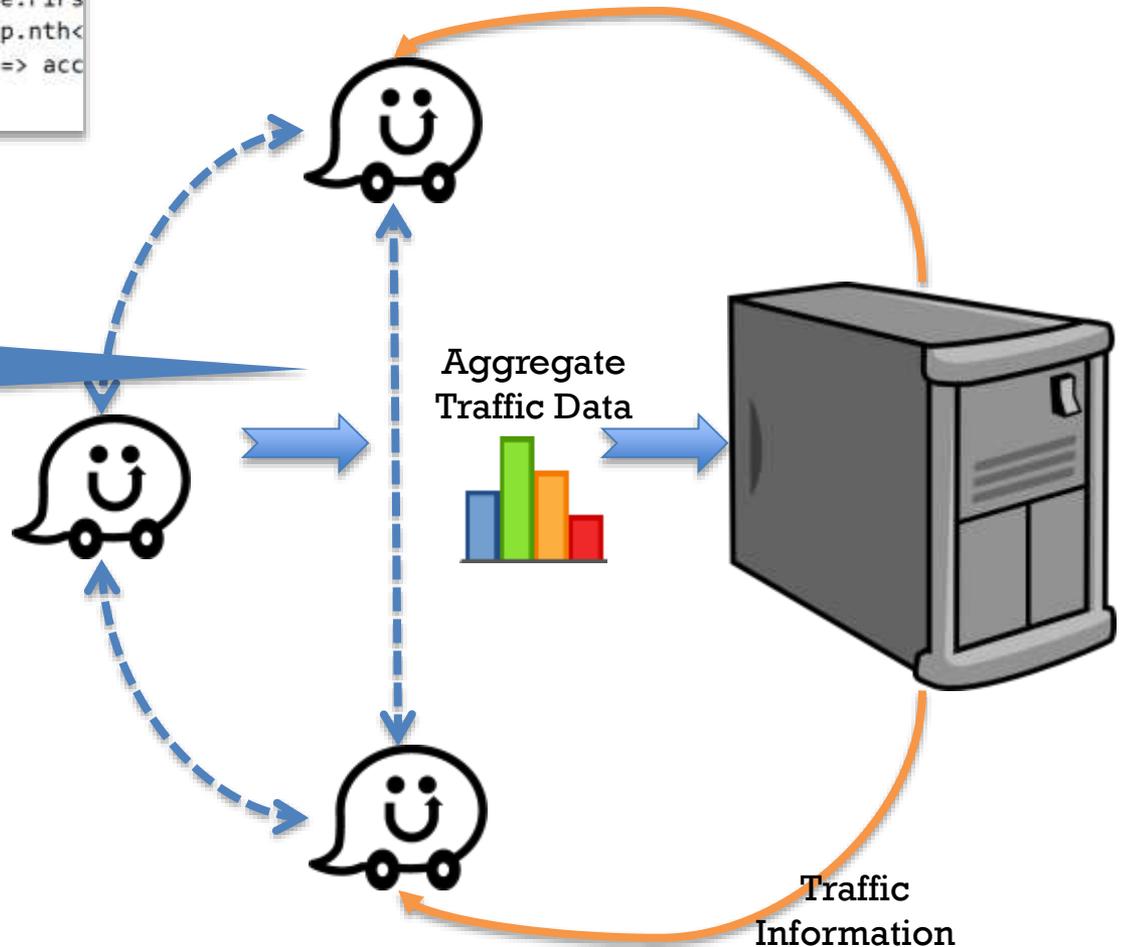
```

// Get the list of user names
var userNames = gpsPts.Select(loclist => loclist.Aggr
var userNames = gpsPts.Select(loclist => loclist.Aggr
var squared = gpsQuads.Select(quad => (quad.nth<int>()
var squared = gpsQuads.Select(quad => (quad.nth<int>()

// Now do the zero-knowledge aggregation
// Get the sum of the squared incremental distances
ZeroKnowledgeBegin();
var distlisttup = squared.Select(sq => sqrtTable.Firs
var distlisttup = squared.Select(sq => sqrtTable.Firs
var distlist = distlisttup.Select(ditup => ditup.nth<
var distlist = distlisttup.Select(ditup => ditup.nth<
ZeroKnowledgeEnd();
var distval = distlist.Aggregate(1, (acc, sum) => acc
var distval = distlist.Aggregate(1, (acc, sum) => acc
ZeroKnowledgeEnd();

```

Insert code for marshaling  
and synchronization



# ZØ: An Optimizing Distributing Zero-Knowledge Compiler

Matthew Fredrikson  
University of Wisconsin

Benjamin Livshits  
Microsoft Research

## Abstract

Traditionally, confidentiality and integrity have been two desirable design goals that are difficult to combine. *Zero-Knowledge Proofs of Knowledge* (ZKPK) offer a rigorous set of cryptographic mechanisms to balance these concerns. However, published uses of ZKPK have been difficult for regular developers to integrate into their code and, on top of that, have not been demonstrated to scale as required by most realistic applications.

This paper presents ZØ (pronounced “zee-not”), a compiler that consumes applications written in C# into code that automatically produces scalable zero-knowledge proofs of knowledge, while automatically splitting applications into distributed multi-tier code. ZØ builds detailed cost models and uses two existing zero-knowledge back-ends with varying performance characteristics to select the most efficient translation. Our case studies have been directly inspired by existing sophisticated widely-deployed commercial products that require both privacy and integrity. The performance delivered by ZØ is as much as 40× faster across six complex ap-

functionality to the client conflicts with a need for computational *integrity*: a malicious client can simply forge the results of a computation.

Traditionally, confidentiality and integrity have been two desirable design goals that are difficult to combine. *Zero-Knowledge Proofs of Knowledge* (ZKPK) offer a rigorous set of cryptographic mechanisms to balance these concerns, and recent theoretical developments suggest that they might translate well into practice. In the last several years, zero-knowledge approaches have received a fair bit of attention [23]. The premise of zero-knowledge computation is its promise of both privacy *and* integrity through the mechanism of cryptographic proofs. However, published uses of ZKPK [4, 5, 7, 8, 19, 36] have been difficult for regular developers to integrate into their code and, on top of that, have not been demonstrated to scale, as required by most realistic applications.

**Zero-knowledge example: pay as you drive insurance:** A frequently mentioned application and a good example of where zero-knowledge techniques excel is the practice of *mileage metering* to bill for car insurance.

# EVALUATION

# Experiments

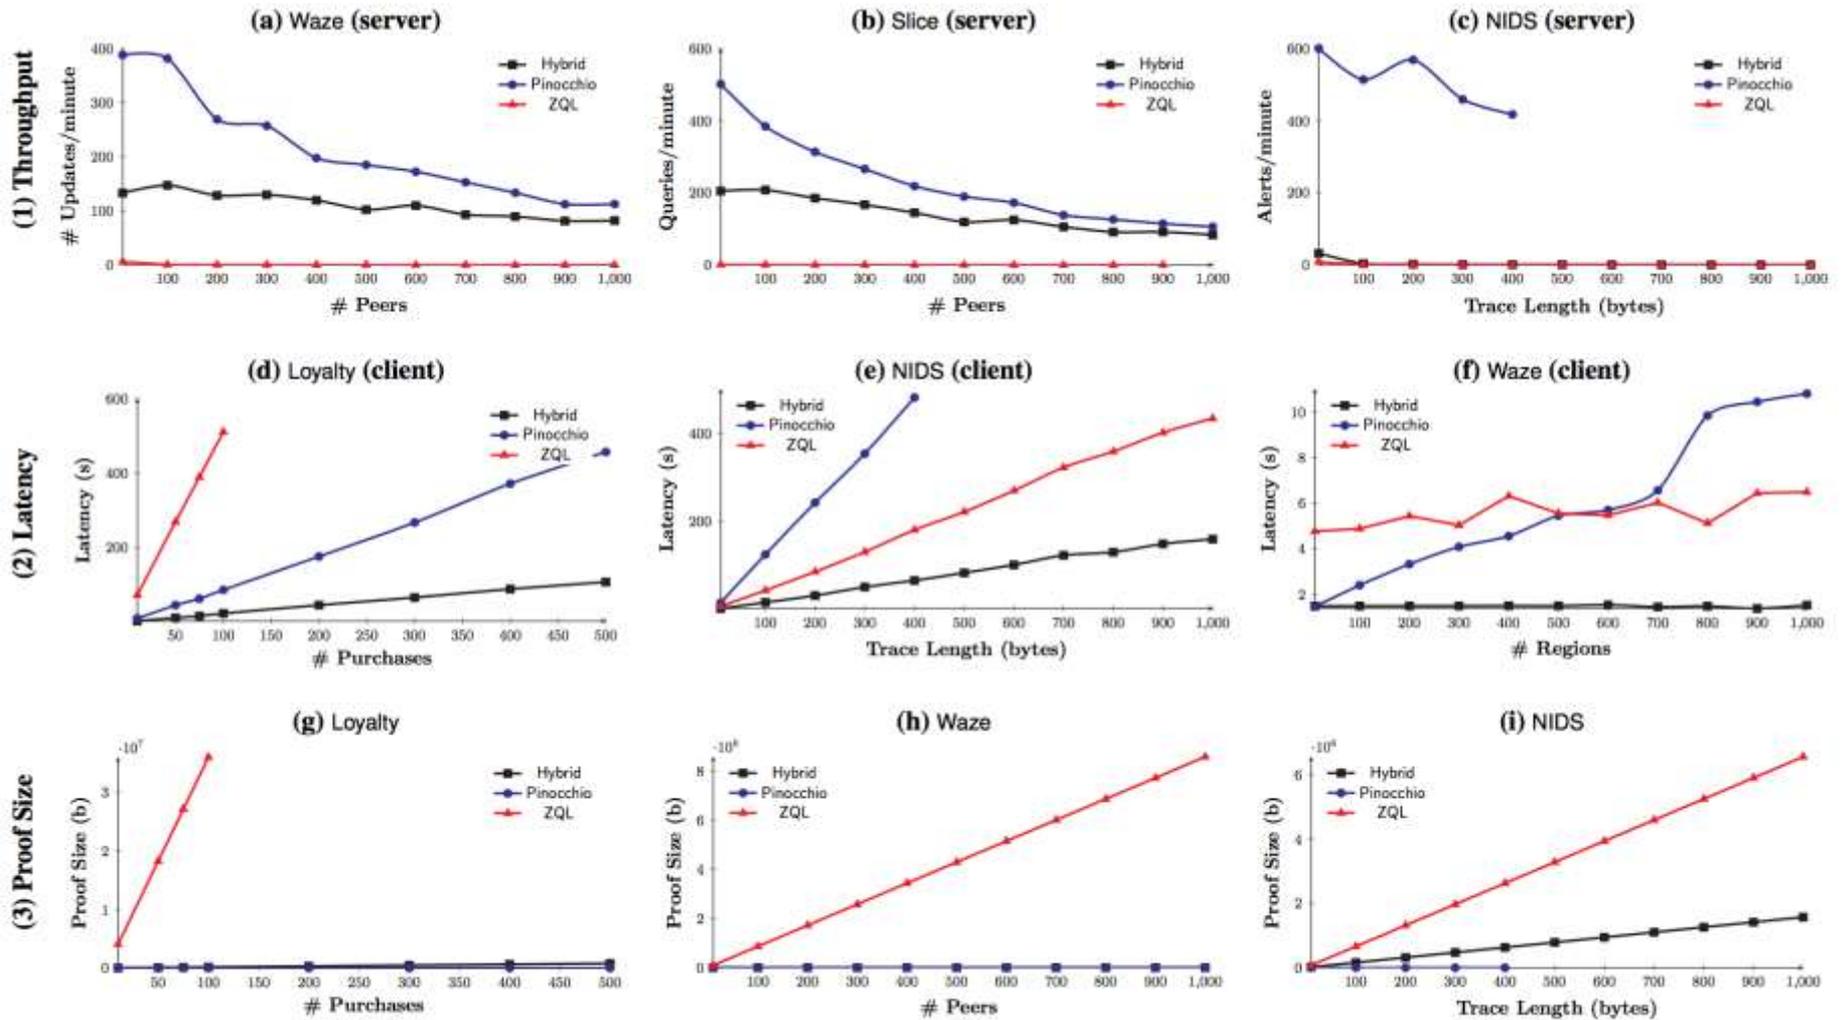
We ran each application in three configurations

## ZQL

## Pinocchio

## ZØ





**Figure 10:** (1) Throughput, (2) latency, and (3) proof size for a characteristic sample of application functionality.

## Loyalty Application, Client's Time to Process Transaction

600

ZQL times out on longer transactions

—■— ZØ  
—●— Pinocchio

ZØ's cost models identified expensive operation, used correct back-end

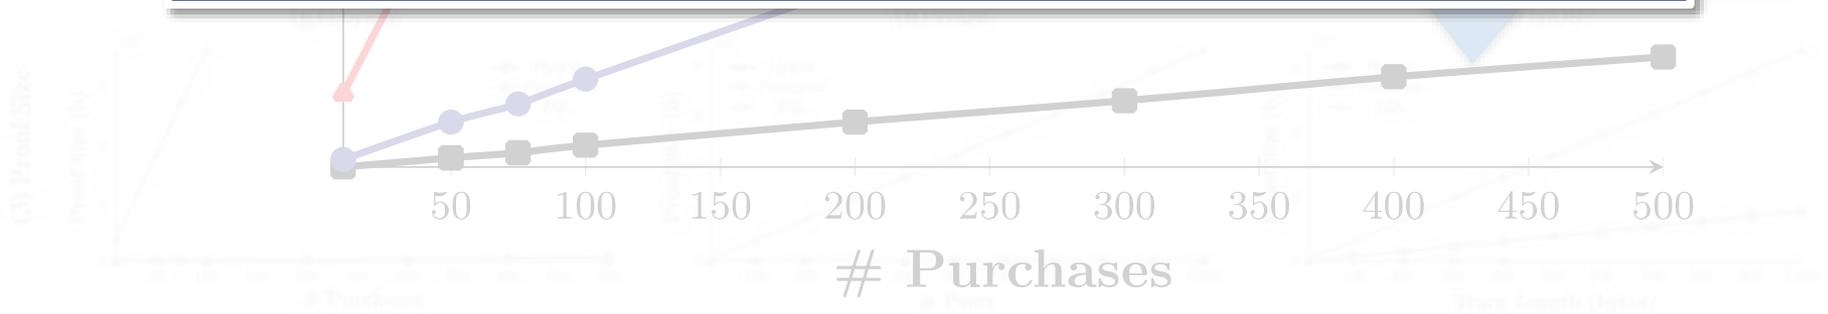


Figure 10: (1) Throughput, (2) latency, and (3) proof size for a characteristic sample of application functionality.

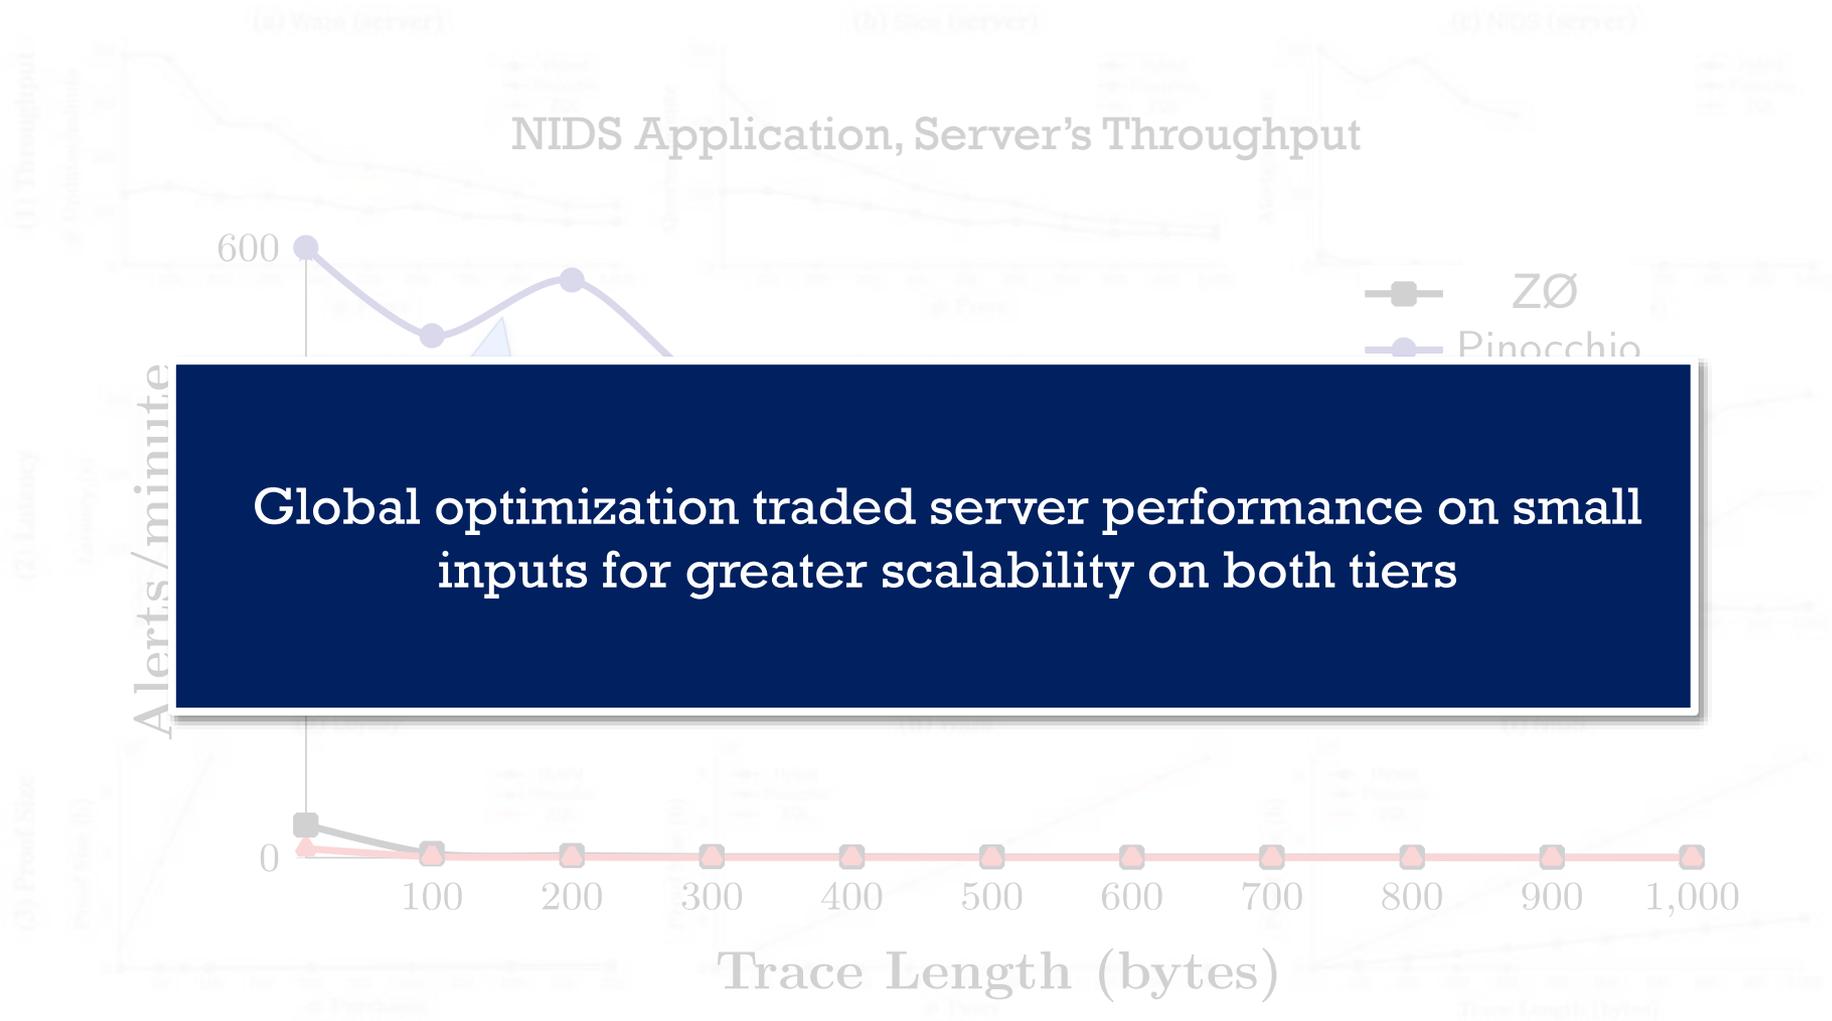


Figure 10: (1) Throughput, (2) latency, and (3) proof size for a characteristic sample of application functionality.

# Experiments

We ran each application in three configurations

ZQL

Pinocchio

ZØ

Scaling

Scales up to 10x larger data

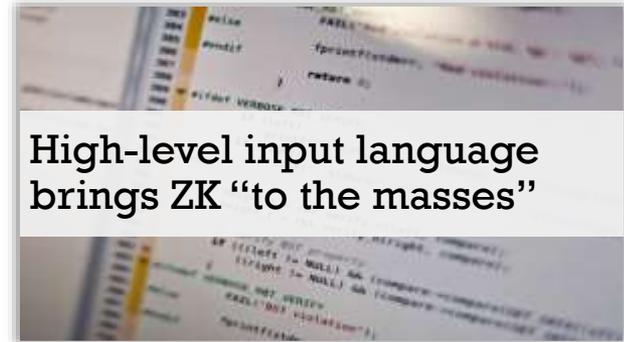
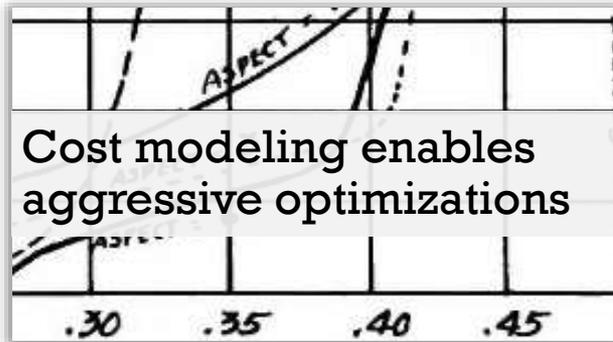
Performance

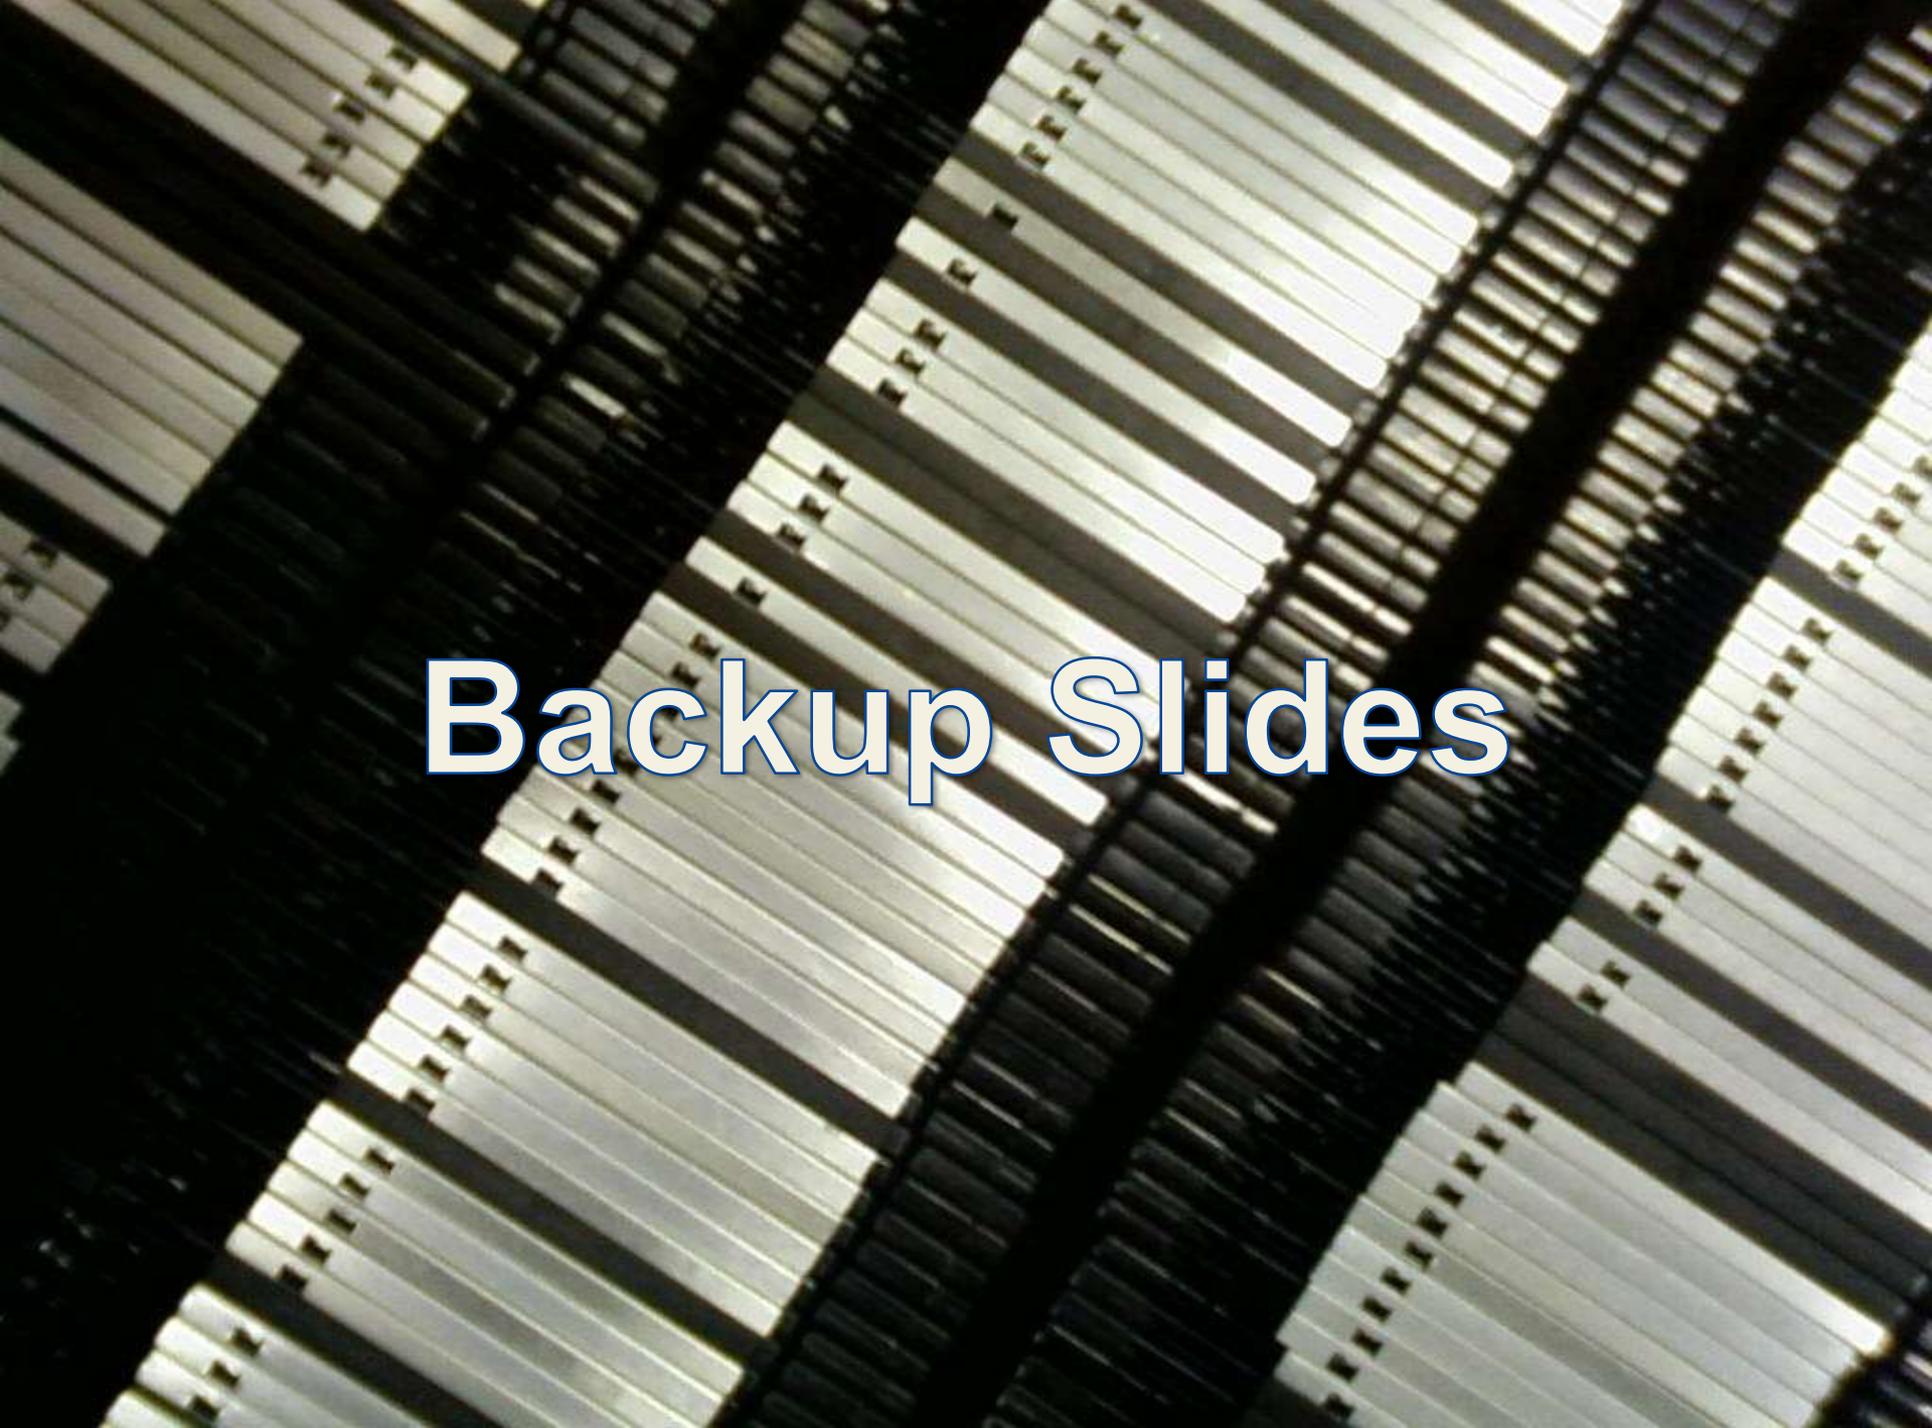
Up to 40x improvement in runtime

Proof Size

Up to 10-100x smaller than ZQL

# Conclusions



An aerial, high-angle photograph of a large stadium, likely the Allianz Arena in Munich. The stadium's roof and seating areas are composed of numerous rectangular panels, creating a complex, repeating geometric pattern of light and dark sections. The perspective is from directly above, looking down at the stadium's structure. The text "Backup Slides" is overlaid in the center of the image.

# Backup Slides

# This talk: at a glance



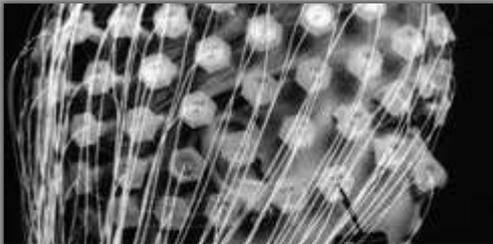
**Crowd-sourced traffic maps**



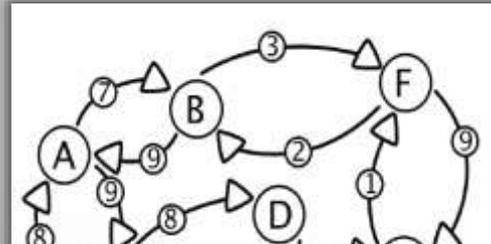
**Personal Fitness Rewards**



**Retail Loyalty Card**



**Human Subjects Studies**



**Collaborative Recommender System**



**Collaborative NIDS**

# Conclusions

- **ZØ is a new zero-knowledge compiler**
  - Detailed cost modeling enable aggressive optimizations
  - High-level language brings ZK “to the masses”
  - Automatic tier splitting simplifies distributed apps
- **Illustrated benefits with six interesting apps**
  - ZØ’s optimizations make these feasible

# Thanks!



Modern apps demand  
personal data

Often the need for data is  
legitimate

Pressure to address privacy  
concerns is widespread

In many applications, this  
creates a tension between  
*privacy* and *integrity*

# Zero-Knowledge: A Promising Solution

Batsvef

Prove that a computation was performed correctly without revealing inputs

The collage features several research papers:

- Efficient Proofs of Discrete Logarithm in Groups with Applications to Cryptography** by Andre Baerentzen, Jakob Bennecke, and Jan Camenisch.
- A Framework for Proving Composable Zero-Knowledge** by Jan Camenisch and Stephan Krenn.
- Pinocchio: Nearly Practical Verifiable Computation** by Bryan Parno, Jon Howell, Mariana Raykova, and Craig Gentry.

A large question mark is centered over the papers. Two blue arrows point from the question mark to the words "Privacy" (pointing right) and "Integrity" (pointing down).



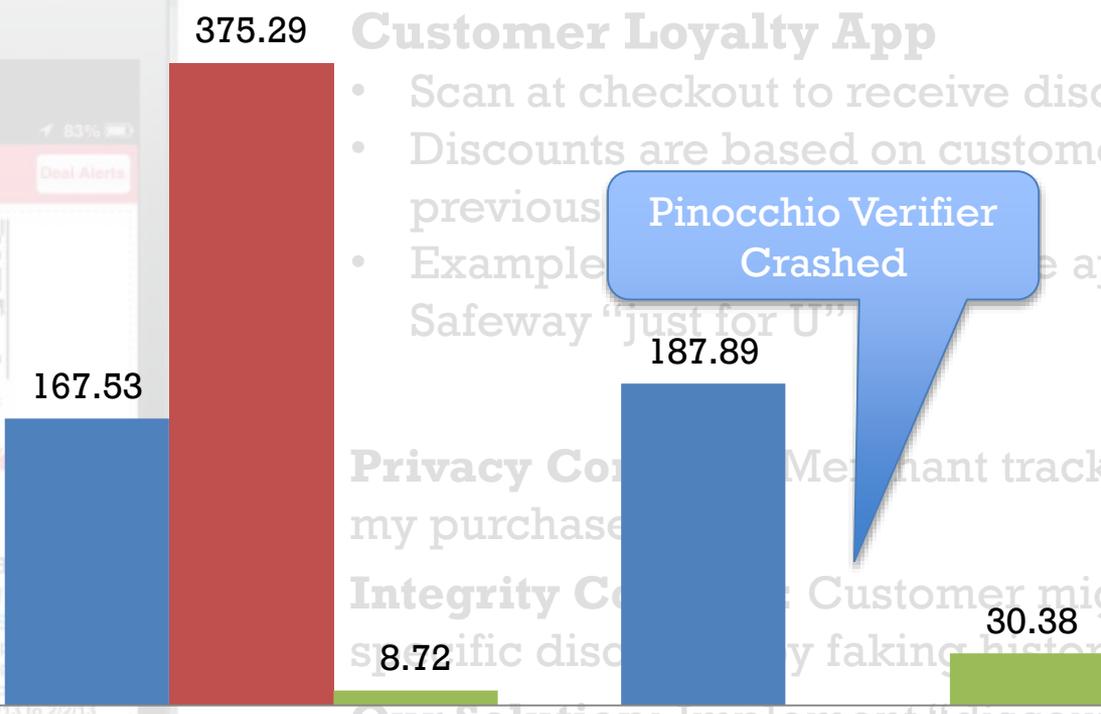
## Zero-Knowledge

- The map is broken into regions, and the desired statistic is the number of clients in each region at time  $t$ .
- At regular intervals, the server requests density stats from the clients.
- On receiving a request, each client:
  1. Takes a GPS reading
  2. Computes its map region
  3. Encodes its region as a vector, zero everywhere but the column for its region
  4. Creates shares of its vector, sends them to other clients
  5. On receiving the other clients' shares, each client sums all received shares and sends the result to the server
- On receiving the summed shares from the clients, the server reconstructs the sum to obtain the density map

## Time to Apply Discount

Execution Time (s)

400  
350  
300  
250  
200  
150  
100  
50  
0



### Customer Loyalty App

- Scan at checkout to receive discounts
- Discounts are based on customer's previous purchases
- Example: Safeway "just for U" app, Safeway "just for U"

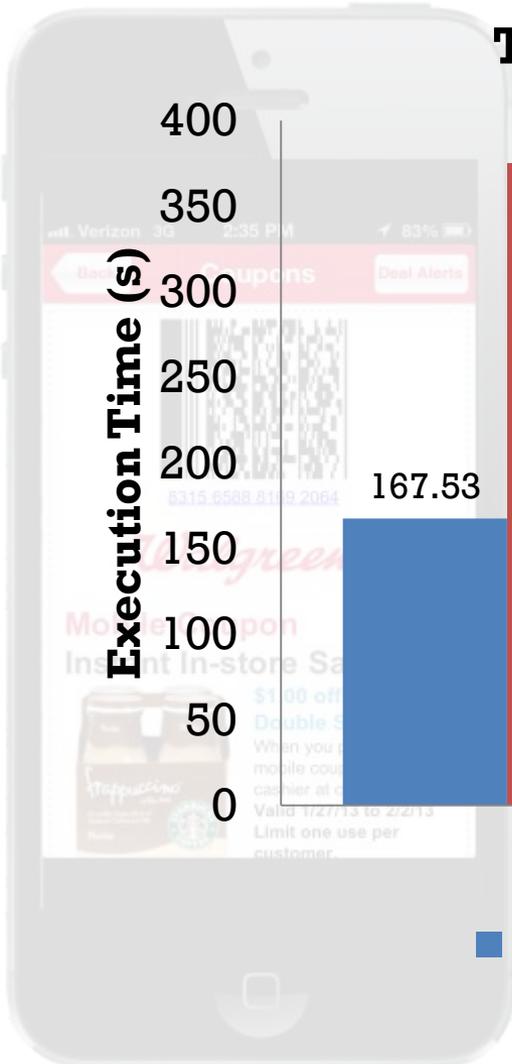
Pinocchio Verifier Crashed

**Privacy Concern:** Merchant tracks all of my purchase history

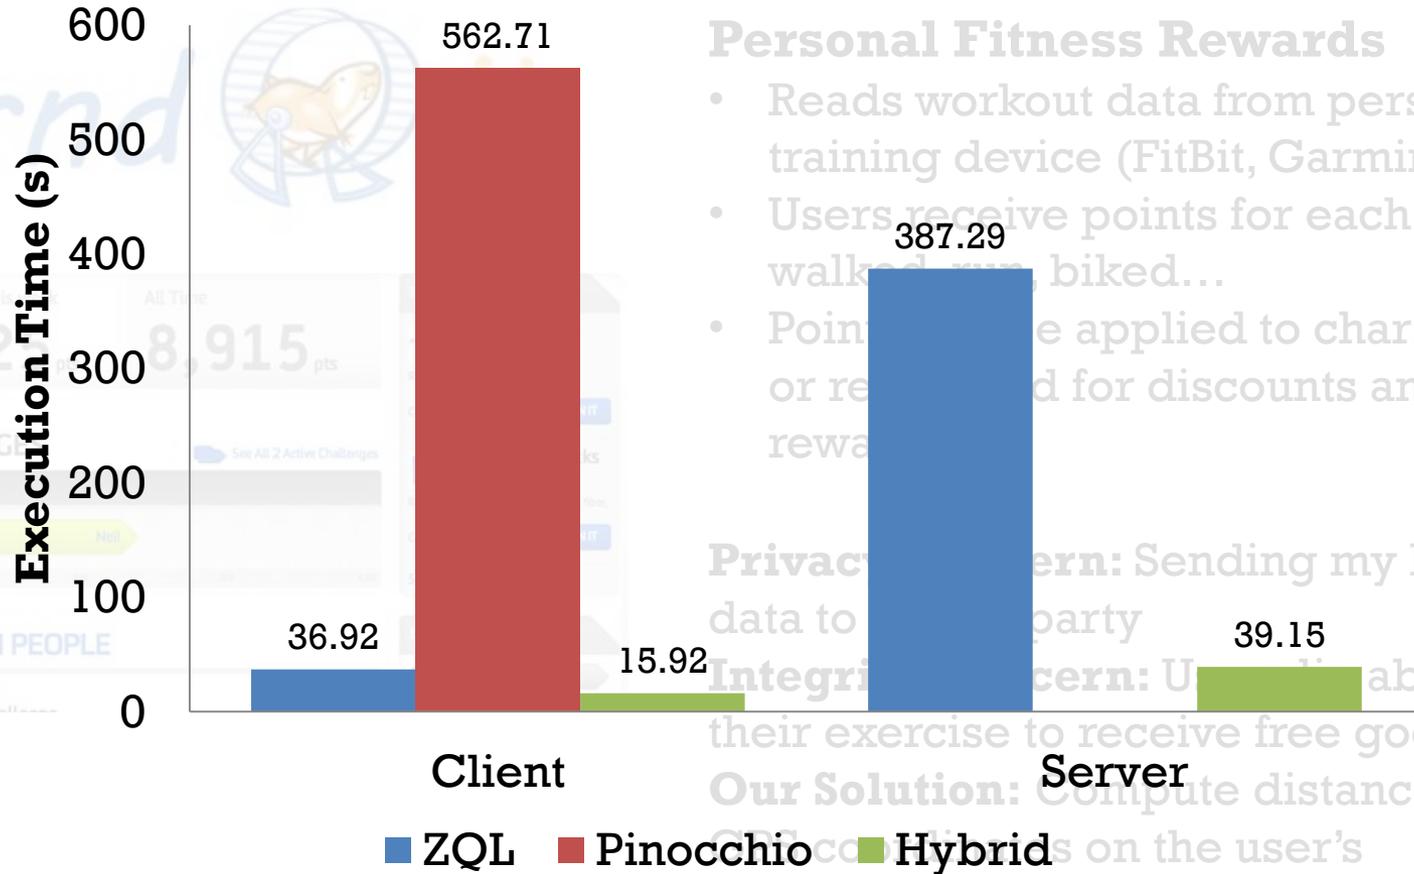
**Integrity Concern:** Customer might target specific discounts by faking purchase history

**Our Solution:** Implement "discounter" as transducer over purchase history, only send merchant.

Client Server  
 ■ ZQL ■ Pinocchio ■ Hybrid



## Time to Redeem Workout



### Personal Fitness Rewards

- Reads workout data from personal training device (FitBit, Garmin, ...)
- Users receive points for each mile walked, run, biked...
- Points can be applied to charities, or redeemed for discounts and rewards

**Privacy Concern:** Sending my location data to a third party

**Integrity Concern:** Users worried about their exercise to receive free goods

**Our Solution:** Compute distance from GPS coordinates on the user's computer, send final result to third party

# Cost Model Accuracy

Different stages of a single ZK computation

	ZQL			Pinocchio		
	Setup	Prover	Verif.	Keygen	Prover	Verif.
FitBit	0.01	1.81	0.10	0.39	0.20	0.00
Waze	0.11	0.29	0.25	0.04	0.02	0.00
Loyalty	0.03	0.35	0.11	0.31	0.20	0.00
Slice	0.06	0.41	0.32	0.05	0.03	0.00
<b>Average</b>	<b>0.05</b>	<b>0.72</b>	<b>0.20</b>	<b>0.20</b>	<b>0.11</b>	<b>0.00</b>

Absolute regression error (in seconds).

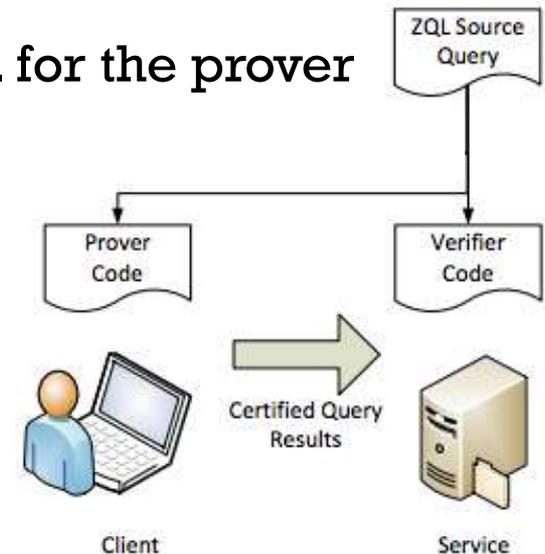
0.32 seconds on average  
(9%)

0.1 seconds on average  
(14%)

# ZQL

Target code is purely-functional, operates on F# lists

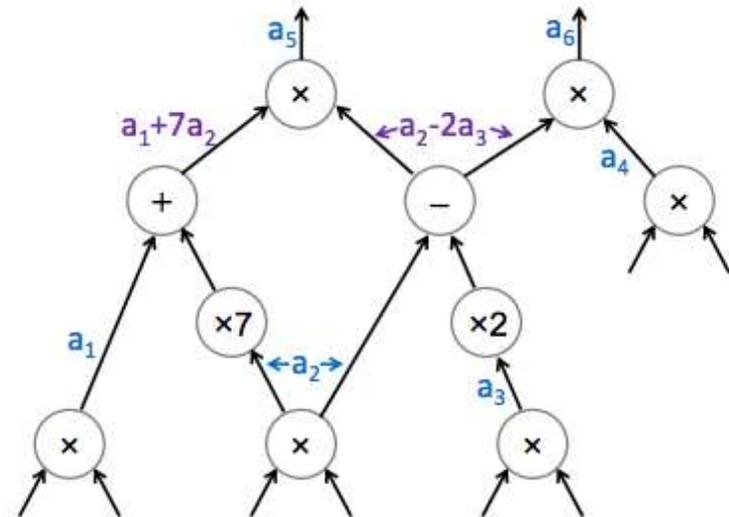
- Translated code mimics structure of original program, does additional cryptographic work for each primitive operation
- Relies heavily on a few primitive operations: `map`, `fold`, `find`
- Lambdas allowed only in limited contexts
- Translated code is highly parallelizable, esp. for the prover
- Runtime available for WP 7 and 8



# Pinocchio

Target code is a fixed-length arithmetic circuit

- Input language is C with static loops, constant dereferences, no recursion
- Everything is in-lined
- Values are broken into constituent bits, Boolean operations used
- Circuit evaluator/prover is optimized native code
- Requires polynomial interpolation and division
- No support for parallel execution



# Goals for ZØ

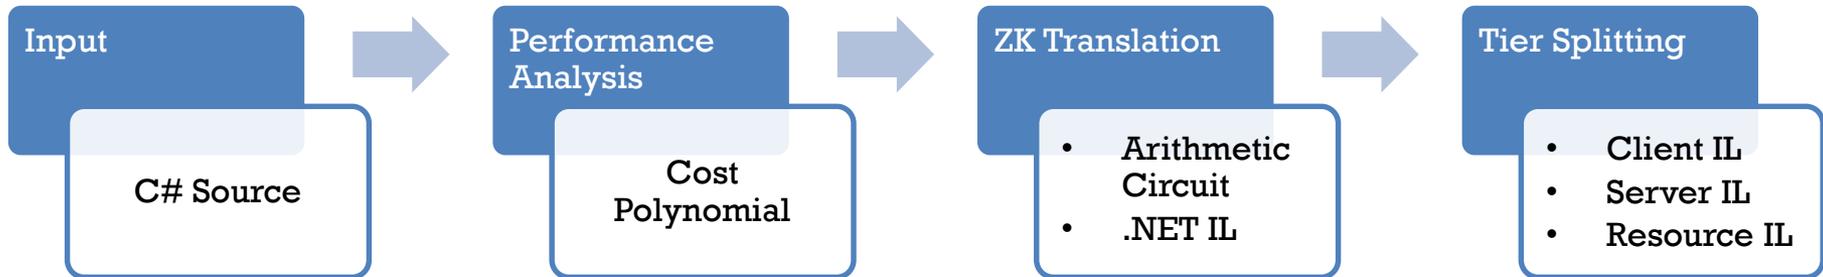
## Performance

- Neither back-end is one-size-fits-all
- Understanding performance requires specialized knowledge
- Bring zero-knowledge to “the masses”

## Usability

- Users should *never write their own crypto*
- Seamless integration with existing code
  - LINQ is our bridge to zero-knowledge
  - Can integrate ZK with large amounts of UI, Libraries, arbitrary logic
- Automates tier-splitting

# ZØ: An Optimizing Compiler for ZK



Implemented in C# and F#

- 9995 LoC
- Uses CCI for processing and analysis, operates on IL
- Uses Solver Foundation to resolve constraints

Still a work in progress

- Integrate cost model generator
- Tune cost model primitive coefficients

# Translation in Action

```
let y = fold (fun acc r ->
              (acc + r)) (toPrv 0) X
```

## ZQL

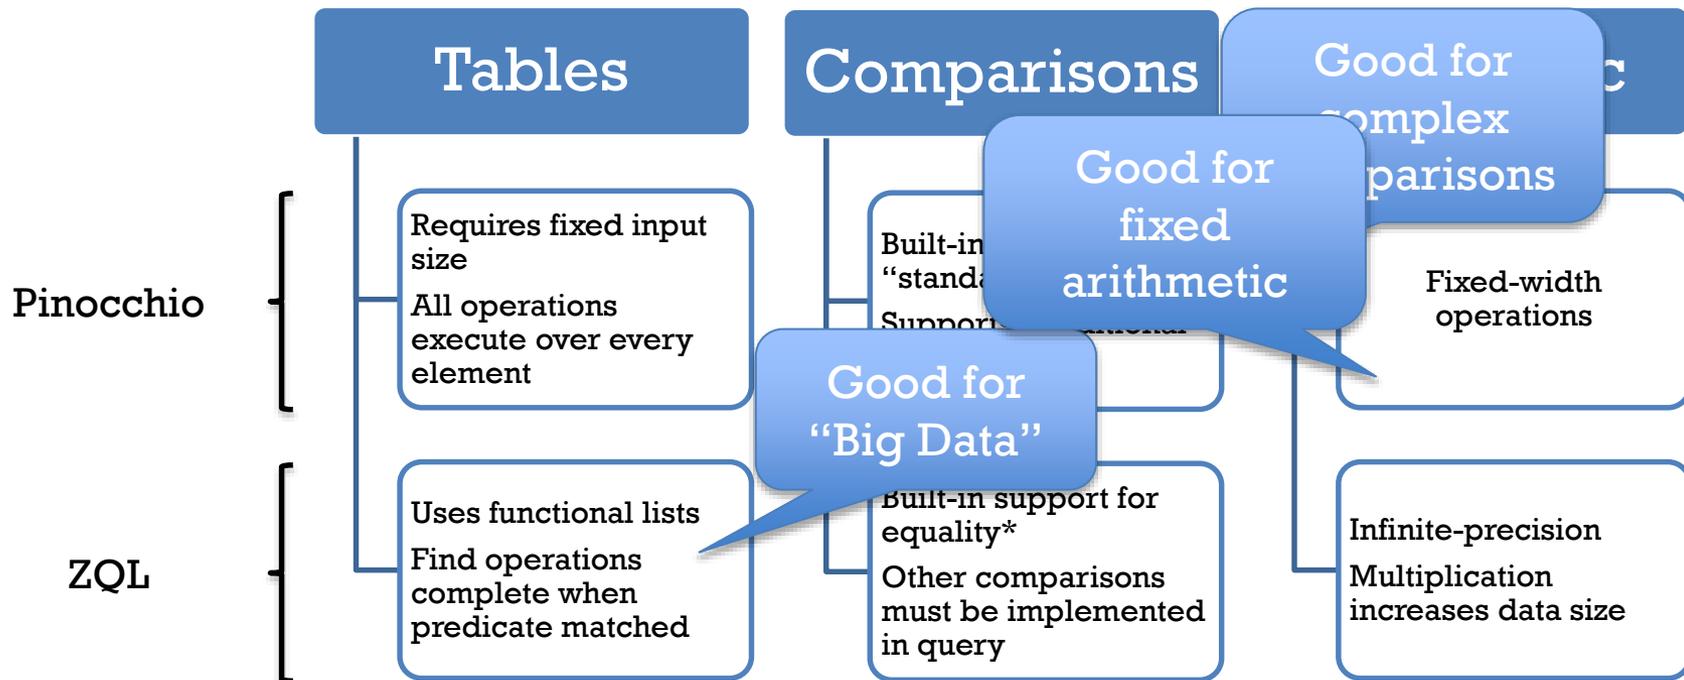
```
let y =
  fold
    (fun acc r ->
      let s_acc, ws_acc = acc
      let s_r, ws_r = r
      let s = add s_acc s_r
      let ws = add ws_acc ws_r
      (s, ws))
    _t2
    _f2
```

1 Multiplication  
100 Additions  
101 I/O Wires

## Pinocchio

```
total 203
input 0 # input
input 1 # input
input 2 # input
input 3 # input
input 4 # input
input 5 # input
input 6 # input
input 7 # input
```

# Performance Comparison



# ZQL Performance

Symbolically execute code generated by ZQL compiler

```
let find (pred : 'a -> 'b)
  (t: 'a table) =
  let size_table = t.size_table
  let size_columns = t.size_columns
  assign_env_coff size_table
  let _ = pred size_columns
  restore_env_coff ()
  size_columns
```

Tracks iterations

Executes nested

Accumulate cost  
of nested op.

```
let (_,c1,c2,c3,c4) = find (fun (reg,x1,x2,y1,y2) -> (reg = reg)) regionList
```

Cryptographic  
Overhead

$eqOp * regionListSize + addOp + 12 * expOp + 3 * extendOp + 14 * mltOp + \dots$

# Pinocchio Performance

Static polynomial based on circuit characteristics

Interpolation

```
mfredrik@mfredrik-PC /cygdrive/z/Desktop/pinoch-new/pinoch-new/code/ccompiler/in
put
$ ../src/vercomp.py z0.c --arith test.arith --cpparg Ibuild/ DBIT_WIDTH=32 DPARA
M=1
mul          : 6002
raw_mul     : 76002
split       : 2000

(info) Linted 30009 field ops from 25007 buses

mfredrik@mfredrik-PC /cygdrive/z/Desktop/pinoch-new/pinoch-new/code/ccompiler/in
put
$ ..
$ ..
```



$O(6002 \log^2 6002)(\text{add}+\text{mul}) + 6507 \text{ExpT} + 44034 \text{ExpB} + 50541 \text{ExpMulB} + \dots$

# Compiling to Zero-Knowledge

IL  
Code

Combinations of  
list-structured data

```

LinqExpr ::= LambdaLinqExpr | ZipLinqExpr
LambdaLinqExpr ::= Id.LambdaLinqId(Lambda)
LambdaLinqId ::= Select | Aggregate | First
ZipLinqExpr ::= Id.Zip(Id, NewAnonObj)
  
```

```

void fun11(struct QuintZ0LEN100 *param12, struct Quint *nthtarget0)
{
    *(nthtarget0) = param12->Enumerable[0];
    int itv1;
    for(itv1 = 0; itv1 < 100; itv1 += 1)
    {
        if(fun10(&(param12->Enumerable[itv1]))) *(nthtarget0) = param12->Enumerable[itv1];
    }
}

```

## 1. Inf

```

void fun13(struct Quint *param14, Int32 *param15)
{
    *(param15) = param14->fld0;
}

```

$con(expr) =$   
 $\{id.elt\}$

$\{id_1, id_2\}$

$\{id\}$  void fun16(Int32 \*param17)

$\{id\}$  {

$\{id.n\}$

```

    struct Quint *nthtarget0;
    nthtarget0 = &(nthtarget0alloc);
    (fun11(&(*regionlist), nthtarget0));
    fun13(nthtarget0, param17);
}

```

$con(id) = \{$

$\}$

$\overline{\cup \{\varphi\}}$

$= 1$

$\overline{d \Rightarrow \Gamma \cup \{\varphi\}}$

$\overline{id_n^f \geq id_n}$

$\overline{i = v}$

C-Basic  $\overline{\Gamma, id_1}$  void outsource(struct Input \*in, struct Output \*out)

{

C-Zip  $\overline{\Gamma, i}$

```

    checkresult = 1;
    regionlist = &(in->regionlist);
    myRegClaimed = &(in->myRegClaimed);
    Int32 myRegalloc;
    Int32 *myReg = &(myRegalloc);
    (fun16(myReg));
    out->out = *(myReg);
    out->checkresult = checkresult;
}

```

$\overline{= v}$

## 3. En

}

# LINQ -> ZQL

## 1. Mostly straightforward translation from LINQ to F#

```

let query = squared.Select
              (squared : (int * int), (sqrtTable : (int * int)[])) ->
              (arg1, arg2) -> (arg1 = sq))
var sqrts = squared.Select(squared, sqrtTable.First(row =>
                          row.Item1 == sq));
  
```

squared.Select  
 sqrtTable.First  
 (arg1, arg2) -> (arg1 = sq)

Caveat: ZQL queries cannot output structured data

## 2. Generate output check

Descend on the structure of the output

```

let _ =
  map2
    (fun lhs0 rhs0 in
      let (lhs2d0, rhs3d0) = (lhs0, rhs0)
      let rhsPrv5 = (toPrv rhs3d0) in check(lhs2d0 == rhsPrv5)
      let rhsPrv9 = (toPrv lhs7d1) in check(lhs6d1 == rhsPrv9)
      ())
    distlist tup expected
  
```

Fail proof checking when false

Pass result of operation to ZQL query

< Demo >

# Back to our example...



Z

Look up region in a large table of coordinates

P

Show that GPS coordinates match result

P

Encode region as a vector

ZP

Creates shares of vector

ZP

Sums other clients' shares

# Distributing Across Tiers

Core Principle: Rely on runtime whenever possible

Minimize the role of the compiler:

- 1.
- 2.

Only the main function can call code on multiple tiers

when tiers

whenever

Functions called from main  
Enumerable

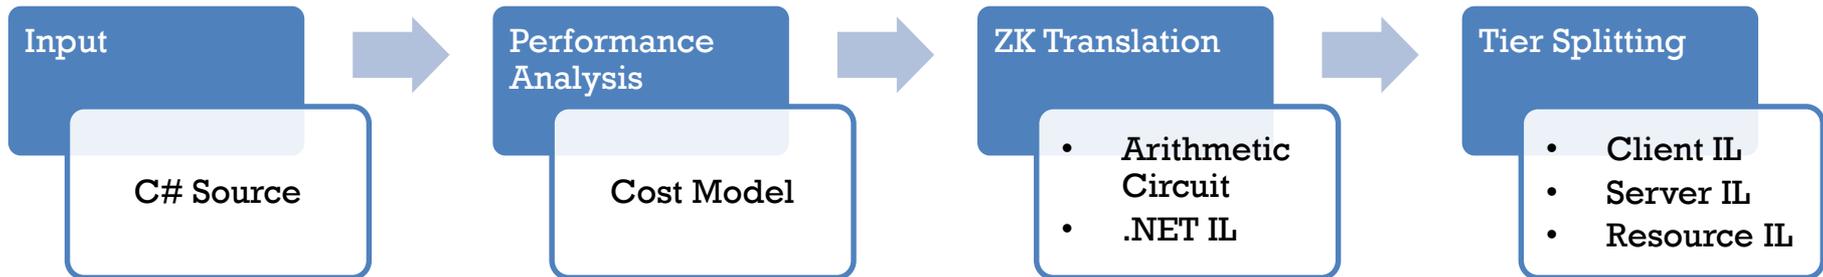
Each element inherits from  
"relocatable" type

```
gpsPts = GetGpsReadings();
shares = MakeSecretShares(gpsPts);
sumShares = AggregateGpsReadings(shares);
stats = InterpolateAndDecode(sumShares);
newDispMap = GenerateDisplayMap(stats);
RenderDisplayMap(newDispMap);
```

```
(shares, ...);
External PendingProjects(External);
Z0iti Shares.GetEnumerator();
while (Z0iti.MoveNext()) {
    Z0E1 = Z0iti.Current;
    Z0E1.Exfiltrate(Z0Wait, External);
}
Z0.Relocatable.WaitForExfiltration(sumShares, Z0Wait);
```

# ZØ: An Optimizing Compiler for ZK

ZØ uses the best of both back-ends as appropriate for the application at hand



# Translating C# To Zero-Knowledge

Specify ZK input sizes to help optimization

```
private int AggregateList([MaximumInputSize(50)] IEnumerable<int> list)
{
    ZeroKnowledgeBegin();
    int agg = list.Aggregate(0, (acc, x) => acc + x);
    ZeroKnowledgeEnd();

    return agg;
}
```

Programmers specify ZK regions

ZK operations given by LINQ expressions