

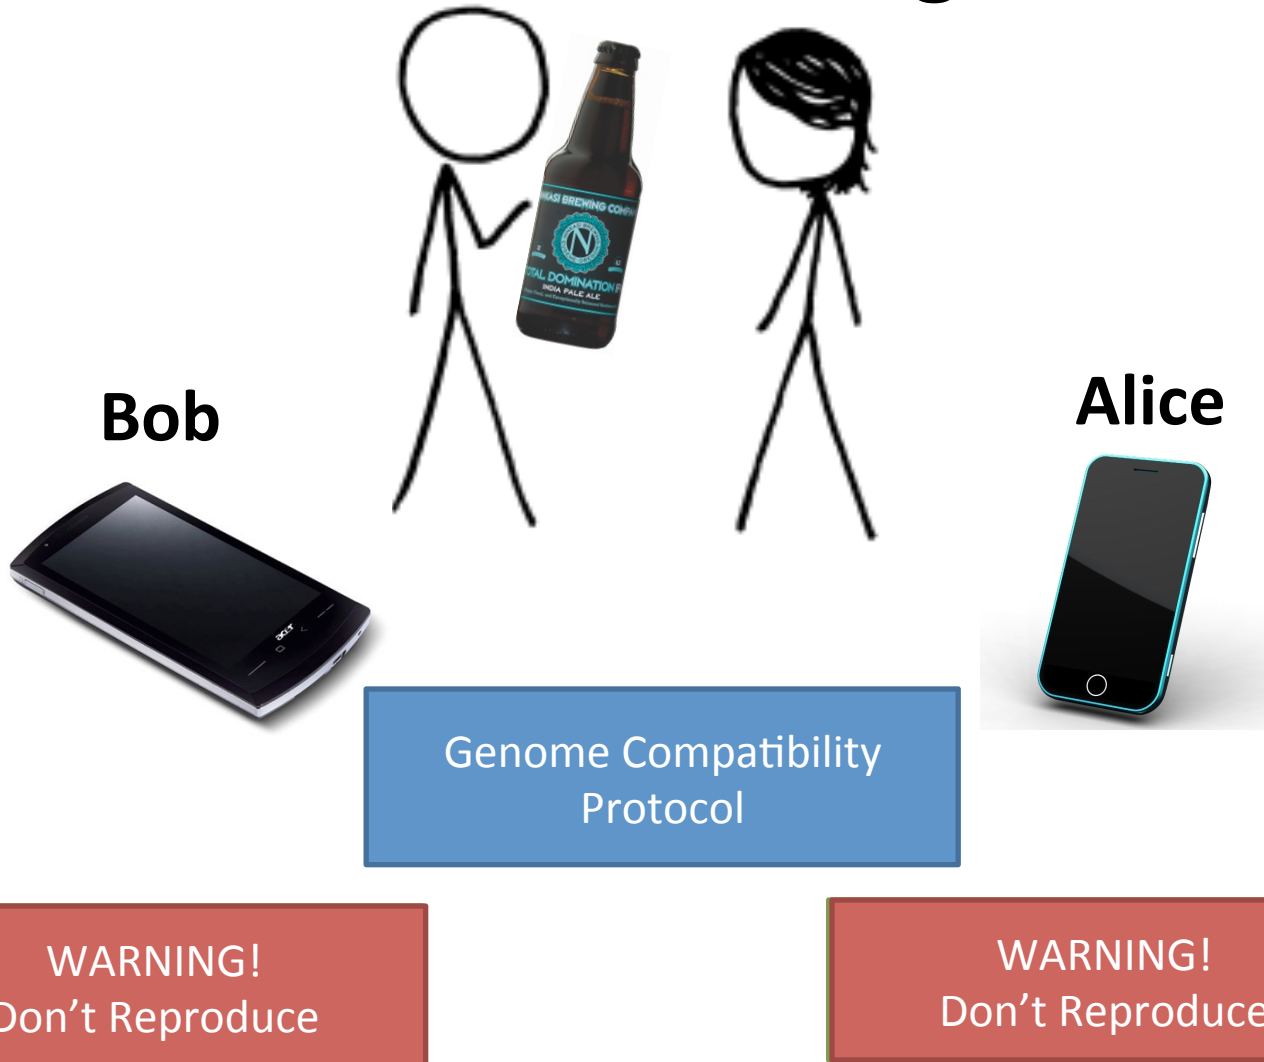


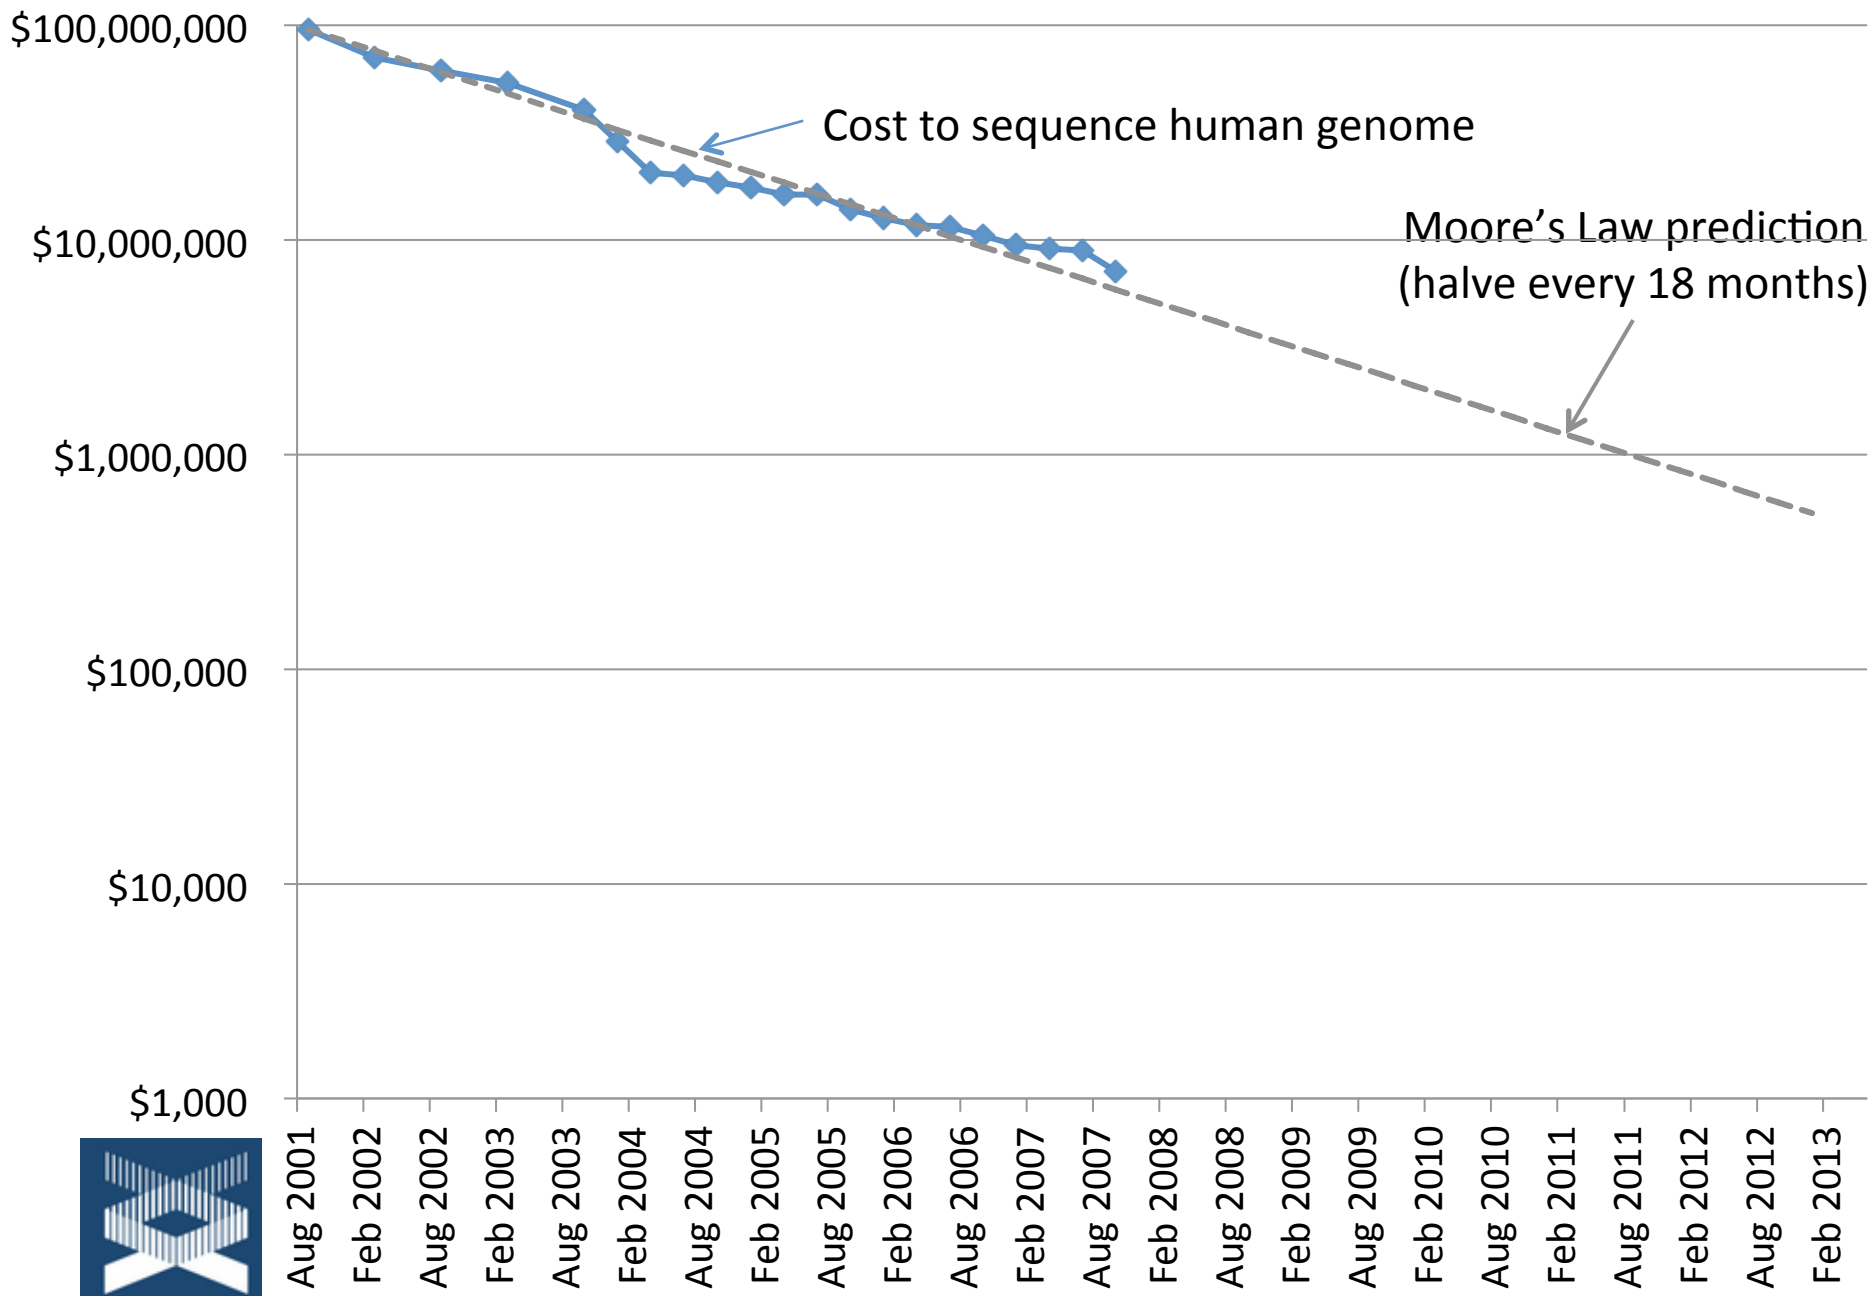
Scaling Secure Computation

Oregon Security Day
5 April 2013


David Evans
University of Virginia
<http://www.cs.virginia.edu/evans>
<http://MightBeEvil.com>

(De)Motivating Application: “Genetic Dating”





genome.gov/sequencingcosts



Ion torrent Personal Genome Machine



genome.gov/sequencingcosts

Year	reference	Technology	Sample	Average Reported Coverage depth (fold)	Reported sequencing consumables cost	Estimated cost per 40-fold coverage
	S4	Sanger (ABI)	JCV	7	\$10,000,000	\$57,000,000
	S5	Roche(454)	JDW	7	\$1,000,000	\$5,700,000
	S6	Illumina	NA18507	30	\$250,000	\$330,000
	S7	Helicos	SRQ	28	\$48,000	\$69,000
2009	this work	this work	NA07022	87	\$8,005	\$3,700
2009	this work	this work	NA19240	63	\$3,451	\$2,200
2009	this work	this work	NA20431	45	\$1,726	\$1,500


[Human Genome Sequencing Using Unchained Base Reads on Self-Assembling DNA Nanoarrays.](#) Radoje Drmanac, Andrew B. Sparks, Matthew J. Callow, Aaron L. Halpern, Norman L. Burns, Bahram G. Kermani, Paolo Carnevali, Igor Nazarenko, Geoffrey B. Nilsen, George Yeung, Fredrik Dahl, Andres Fernandez, Bryan Staker, Krishna P. Pant, Jonathan Baccash, Adam P. Borcharding, Anushka Brownley, Ryan Cedeno, Linsu Chen, Dan Chernikoff, Alex Cheung, Razvan Chirita, Benjamin Curson, Jessica C. Ebert, Coleen R. Hacker, Robert Hartlage, Brian Hauser, Steve Huang, Yuan Jiang, Vitali Karpinchyk, Mark Koenig, Calvin Kong, Tom Landers, Catherine Le, Jia Liu, Celeste E. McBride, Matt Morenzoni, Robert E. Morey, Karl Mutch, Helena Perazich, Kimberly Perry, Brock A. Peters, Joe Peterson, Charit L. Pethiyagoda, Kaliprasad Pothuraju, Claudia Richter, Abraham M. Rosenbaum, Shaunak Roy, Jay Shafto, Uladzislau Sharanovich, Karen W. Shannon, Conrad G. Sheppy, Michel Sun, Joseph V. Thakuria, Anne Tran, Dylan Vu, Alexander Wait Zaranek, Xiaodi Wu, Snezana Drmanac, Arnold R. Oliphant, William C. Banyai, Bruce Martin, Dennis G. Ballinger, George M. Church, Clifford A. Reid. ***Science***, **January 2010.**

Dystopia



Personalized Medicine

Secure Two-Party Computation



Can Alice and Bob compute a function of their private data, without exposing anything about their data besides the result?

Secure Function Evaluation

Alice (circuit generator)

Bob (circuit evaluator)

Picks $a \in \{0, 1\}^s$

Agree on
 $f(a, b) \rightarrow x$

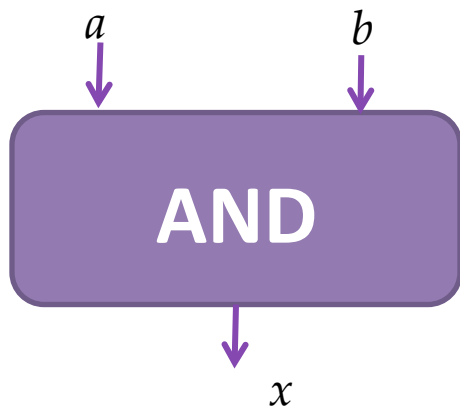
Picks $b \in \{0, 1\}^t$

Garbled Circuit Protocol

Outputs $x = f(a, b)$
without revealing a
to Bob or b to Alice.

Regular Logic


Inputs		Output
a	b	x
0	0	0
0	1	0
1	0	0
1	1	1



Computing with Meaningless Values?

Inputs		Output
a	b	x
a_0	b_0	x_0
a_0	b_1	x_0
a_1	b_0	x_0
a_1	b_1	x_1



a_i, b_i, x_i are **random** values, chosen by the **circuit generator** but **meaningless** to the **circuit evaluator**.



Computing with Garbled Tables

Inputs		Output
a	b	x
a_0	b_0	$Enc_{a_0,b_0}(x_0)$
a_0	b_1	$Enc_{a_0,b_1}(x_0)$
a_1	b_0	$Enc_{a_1,b_0}(x_0)$
a_1	b_1	$Enc_{a_1,b_1}(x_1)$

Bob can only decrypt
one of these!



Random
Permutation

Garbled Circuit Protocol

Alice (circuit generator)

Creates random keys: $a_0, a_1, b_0, b_1, x_0, x_1$

Garbled Gate
$Enc_{a_0, b_1}(x_0)$
$Enc_{a_1, b_1}(x_1)$
$Enc_{a_1, b_0}(x_0)$
$Enc_{a_0, b_0}(x_0)$

Bob (circuit evaluator)

Sends a_i to Bob based on her input value

a_0

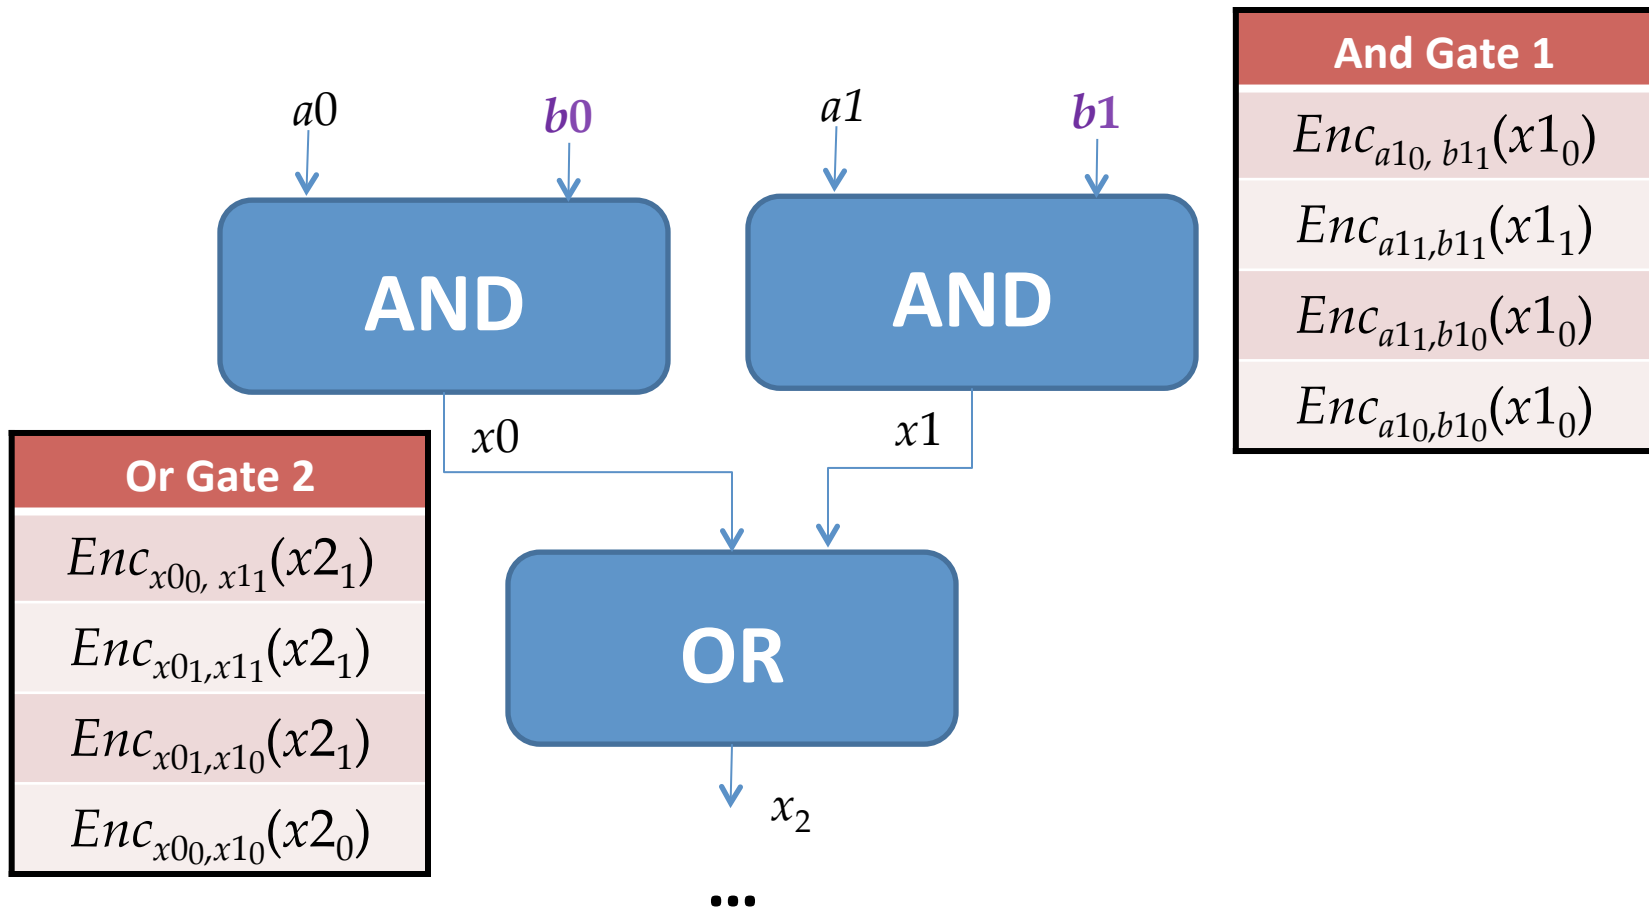
How does the Bob learn his own input wires?

Primitive: Oblivious Transfer




Oblivious: Alice doesn't learn which secret Bob obtains
Transfer: Bob learns one of Alice's secrets

Chaining Garbled Circuits



We can do *any* computation privately this way!

Building Computing Systems



$$Enc_{x_{00}, x_{11}}(x_{2_1})$$

$$Enc_{x_{01}, x_{11}}(x_{2_1})$$

$$Enc_{x_{01}, x_{10}}(x_{2_1})$$

$$Enc_{x_{00}, x_{10}}(x_{2_0})$$

Digital Electronic Circuits

Operate on **known data**

One-bit logical operation requires moving a few electrons a few nanometers
(hundreds of Billions per second)

Reuse is great!

Garbled Circuits

Operate on **encrypted wire labels**

One-bit logical operation requires performing (up to) 4 encryption operations: **very slow execution**

Reuse is not allowed for privacy: **huge circuits needed**

Fairplay

Alice

Bob

```
program Millionaires {
  type int = Int<4>; // 4-bit integer
  type AliceInput = int;
  type BobInput = int;
  type AliceOutput = Boolean;
  type BobOutput = Boolean;
  type Output = struct {
    AliceOutput alice, BobOutput bob};
  type Input = struct {
    AliceInput alice, BobInput bob};

  function Output out(Input inp) {
    out.alice = inp.alice > inp.bob;
    out.bob = inp.bob > inp.alice;
  }
}
```

SFDL Program

SFDL
Compiler

SFDL
Compiler

Circuit
(SHDL)

Garbled Tables
Generator

Garbled Tables
Evaluator

Dahlia Malkhi, Noam Nisan,
Benny Pinkas and Yaron Sella
[USENIX Sec 2004]

Faster Circuit Execution



Yan Huang

(UVa PhD ~~Student~~)

Graduate

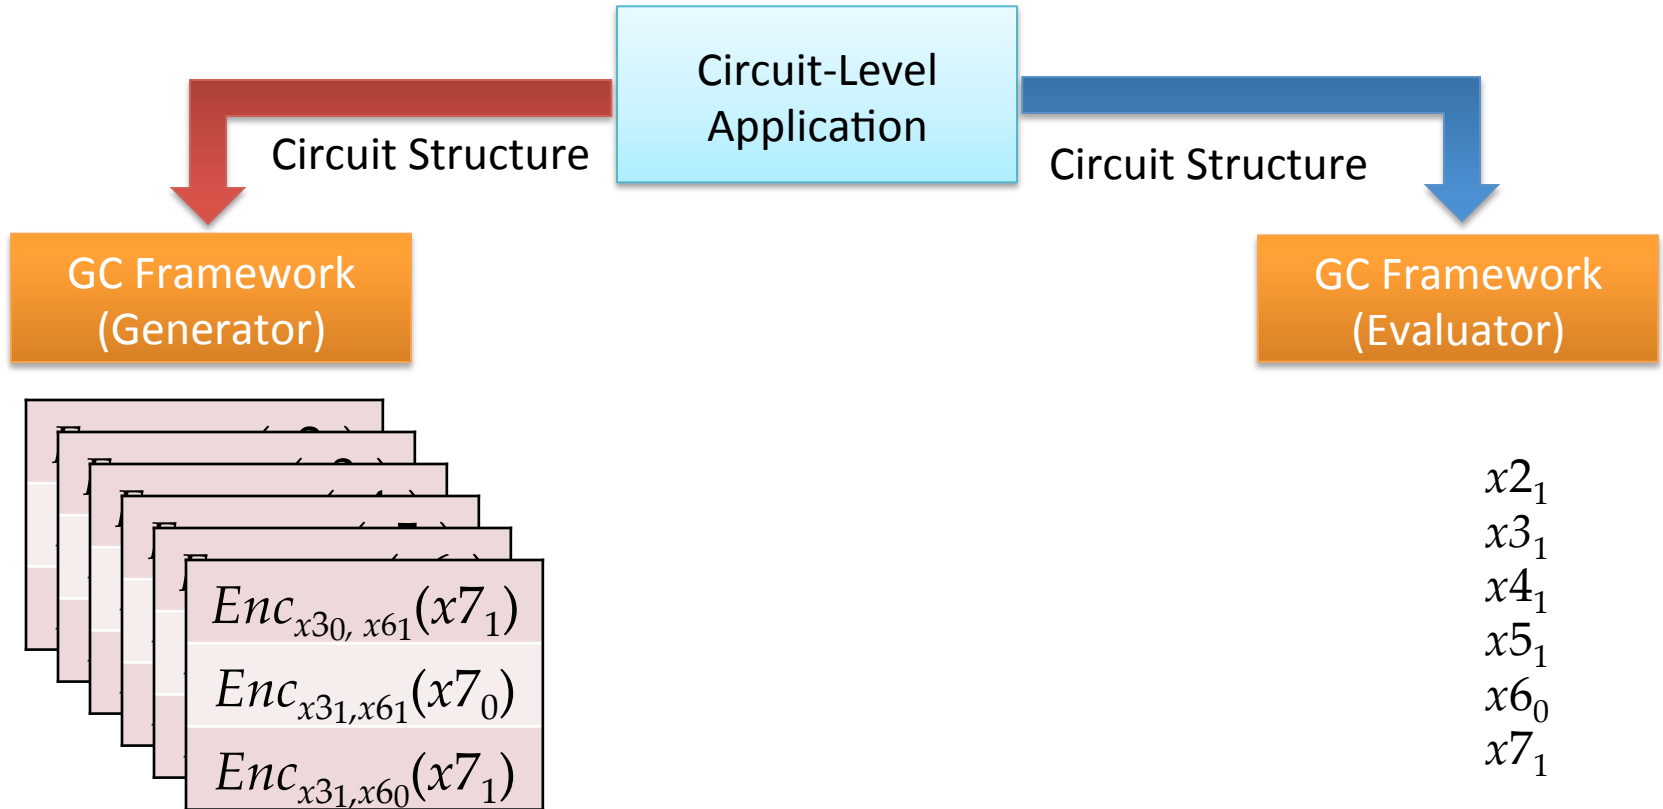
Pipelined Execution

Optimized Circuit Library

Partial Evaluation

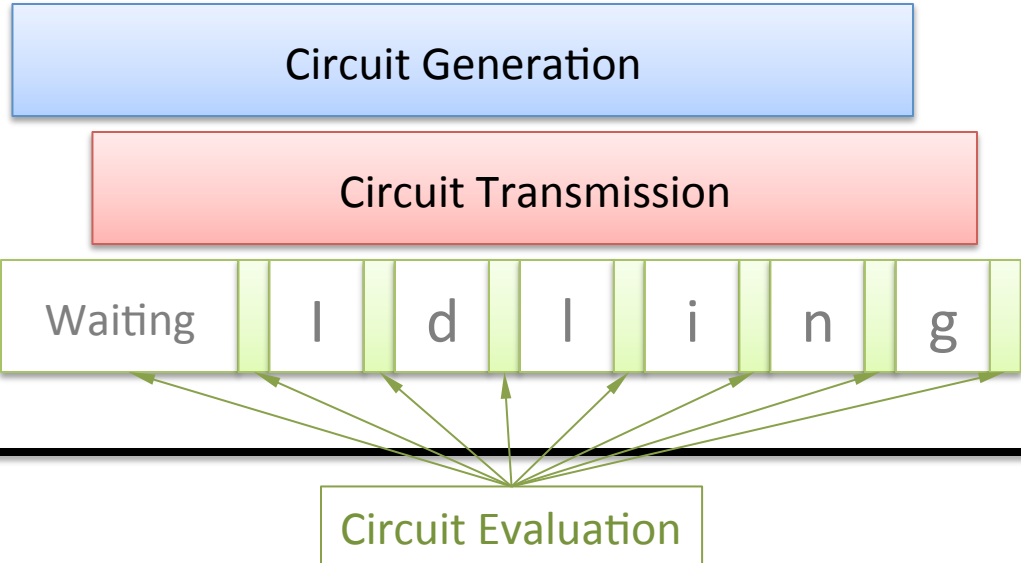
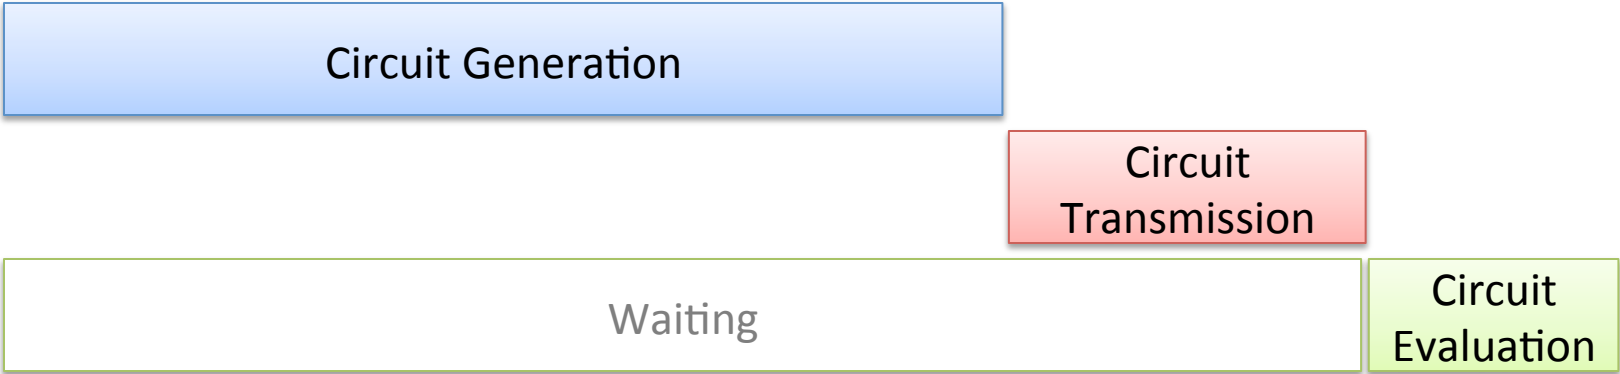
Yan Huang, David Evans, Jonathan Katz, and Lior Malka. *Faster Secure Two-Party Computation Using Garbled Circuits*. **USENIX Security 2011**.

Pipelined Execution



Saves memory: never need to keep whole circuit in memory

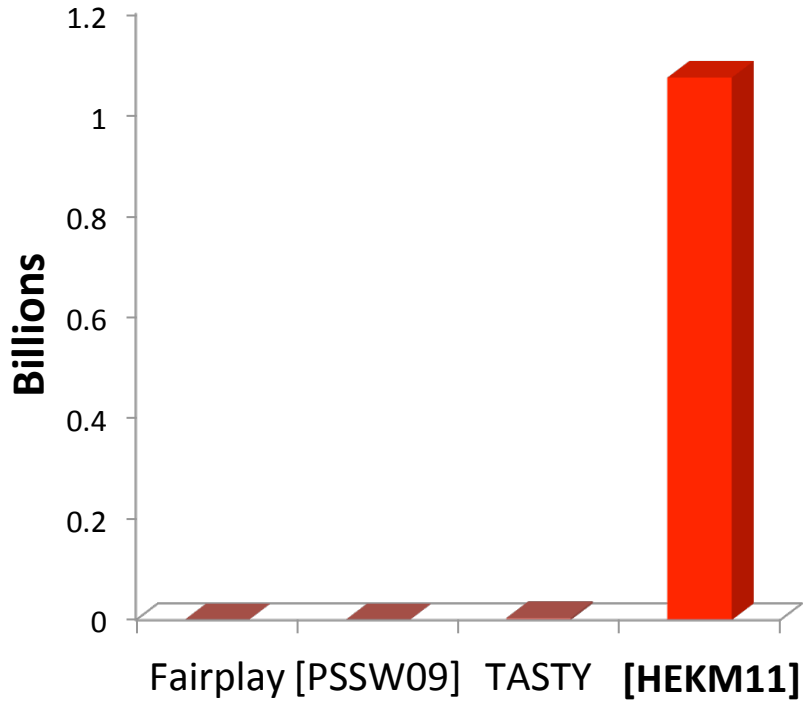
Pipelining



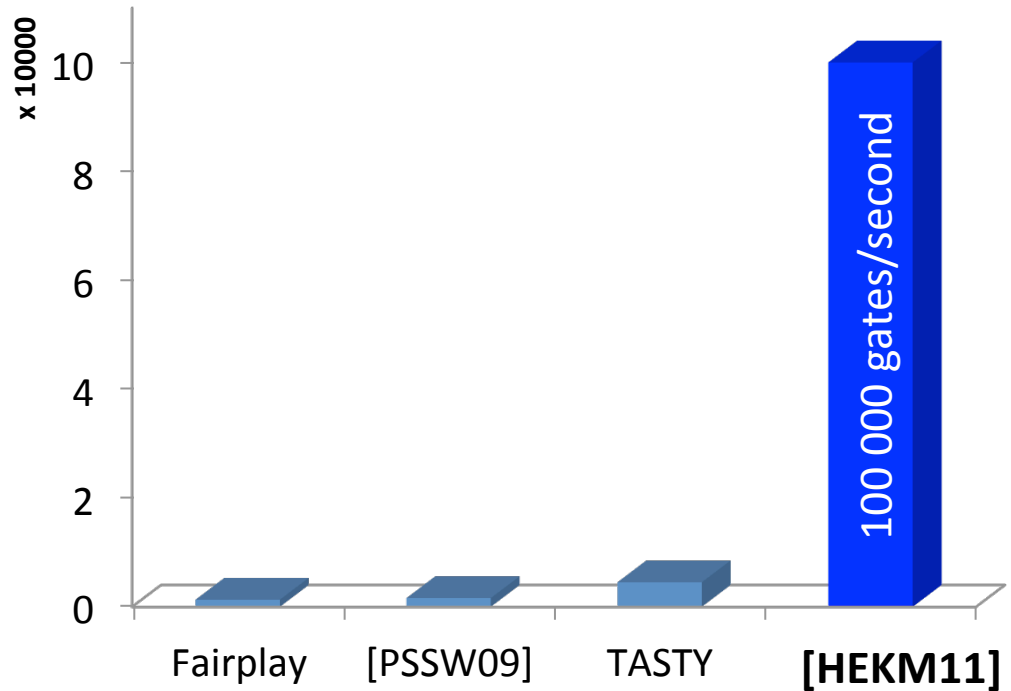
Saves **time**: reduces latency and improves throughput



Results



Scalability
(billions of gates)



Performance
(10,000x non-free gates per second)

Semi-Honest is Half-Way There

Privacy

Nothing is revealed other than the output

Correctness

The output of the protocol is indeed $f(x,y)$

Generator


Evaluator



As long as evaluator doesn't send result back, and a malicious-resistant OT is used, **privacy** for evaluator is guaranteed.

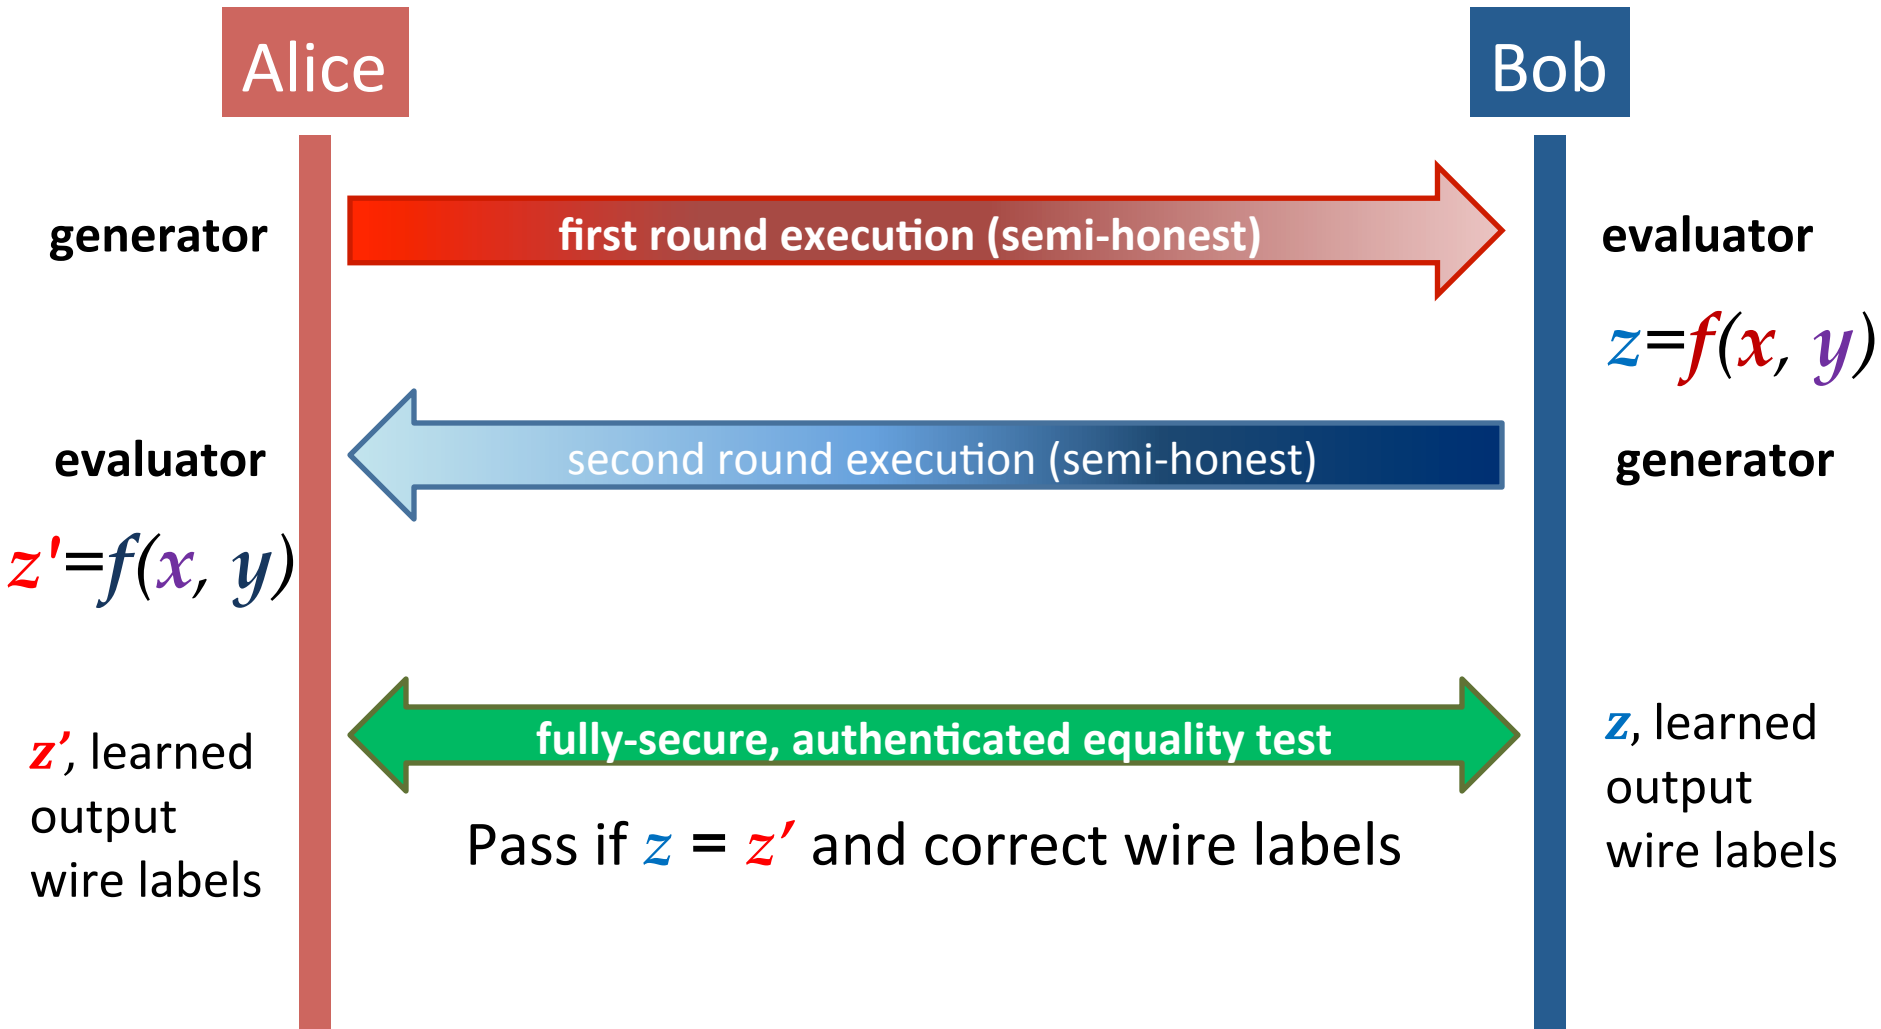
How can we get both correctness, and maintain privacy while giving both parties result?

Dual Execution Protocols



Yan Huang, Jonathan Katz, and David Evans. *Quid-Pro-Quo-tocols: Strengthening Semi-Honest Protocols with Dual Execution*. IEEE Security and Privacy (Oakland) 2012.

Dual Execution Protocol



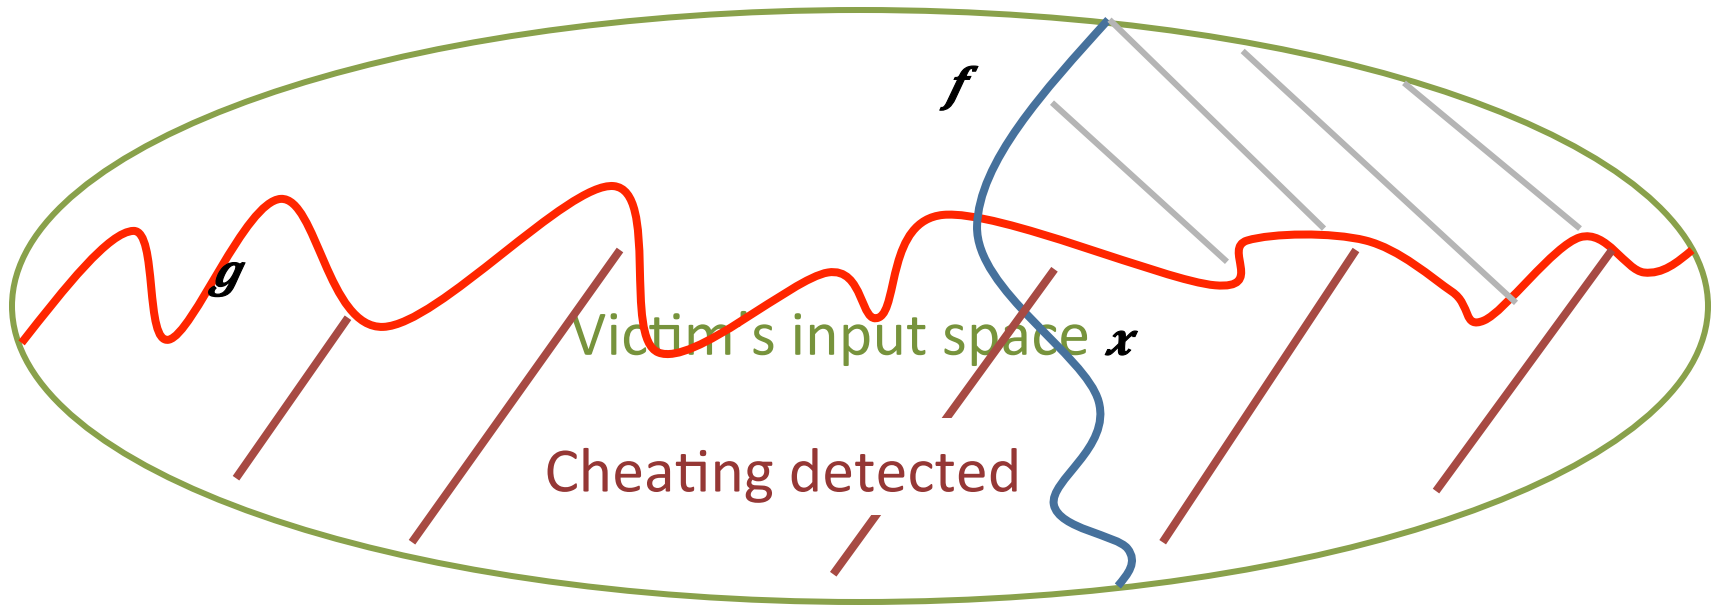
Security Properties

Correctness: guaranteed by authenticated, secure equality test

Privacy: Leaks **one** (extra) **bit** on average
adversarial circuit generator provides a circuit that fails on $\frac{1}{2}$ of inputs

Malicious generator can decrease likelihood of being caught, and increase information leaked when caught (but decreases average information leaked): at extreme, circuit fails on just one input

1-bit Leak

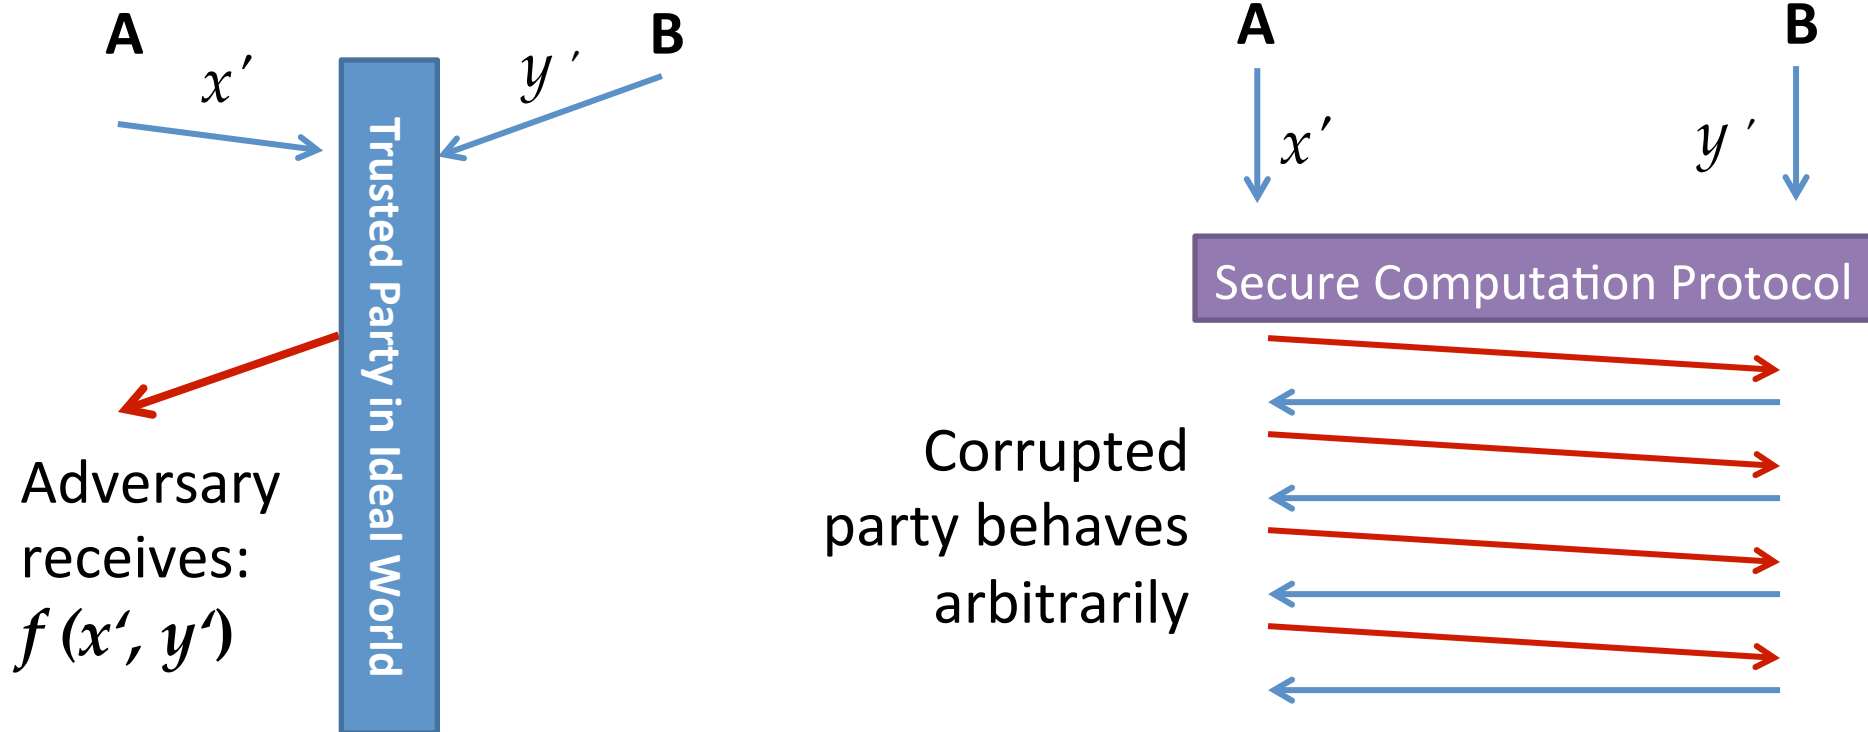


Proving Security: Malicious

Ideal World

Show equivalence

Real World




Standard Malicious Model: can't prove this for Dual Execution

Proof of Security: One-Bit Leakage

Ideal World

Controlled by
malicious \mathcal{A}



Can prove equivalence to this for Dual Execution protocols

Implementation

Alice

Bob

generator

first round execution (semi-honest)

evaluator

Recall: work to generate is ~3x work to evaluate!

$$z = f(x, y)$$

evaluator

second round execution (semi-honest)

generator

$$z' = f(x, y)$$


z' , learned
output
wire labels

fully-secure, authenticated equality test

Pass if $z = z'$ and correct wire labels

z , learned
output
wire labels

Performance



Circuits of arbitrary size can be done this way



Privacy-Preserving
Biometric Matching




Private AES
Encryption

Private
Personal
Genomics



Applications




Private Set Intersection

	Problem	Best Previous Result	Our Result	Speedup
NDSS 2012	Private Set Intersection (contact matching, common disease carrier)	Competitive with best custom protocols, scales to millions of 32-bit elements		
USENIX Security 2011	Hamming Distance (Face Recognition)	213s [SCiFI, 2010]	0.051s	4176
	Levenshtein Distance (genome, text comparison) – two 200-character inputs	534s [Jha+, 2008]	18.4s	29
	Smith-Waterman (genome alignment) – two 60-nucleotide sequences	[Not Implementable]	447s	-
	AES Encryption	3.3s [Henecka, 2010]	0.2s	16.5
	Fingerprint Matching (1024-entry database, 640x8bit vectors)	~83s [Barni, 2010]	18s	4.6
NDSS 2011				

Crazy Things in Typical Code

```
a[i] = x
```


Circuit for Array Update



$$a[i] = x$$

Easy (and Common) Case

```
for (i = 0; i < n; i++)  
    a[i] += 1
```



Circuit Structures

Design circuits to support typical **data structures** efficiently

Non-trivial access patterns, but **patterns nonetheless**

Main opportunities:

Locality and Batching



Samee Zahur

(UVa PhD Student)

Samee Zahur and David Evans. *Circuit Structures for Improving Efficiency of Security & Privacy Tools*. IEEE Security and Privacy (Oakland) 2013.

Locality: Stacks and Queues


```
if (x != 0)
  a[i] += 1
  if (a[i] > 10)
    i += 1
  a[i] = 5
```

Data-oblivious code
No branching allowed

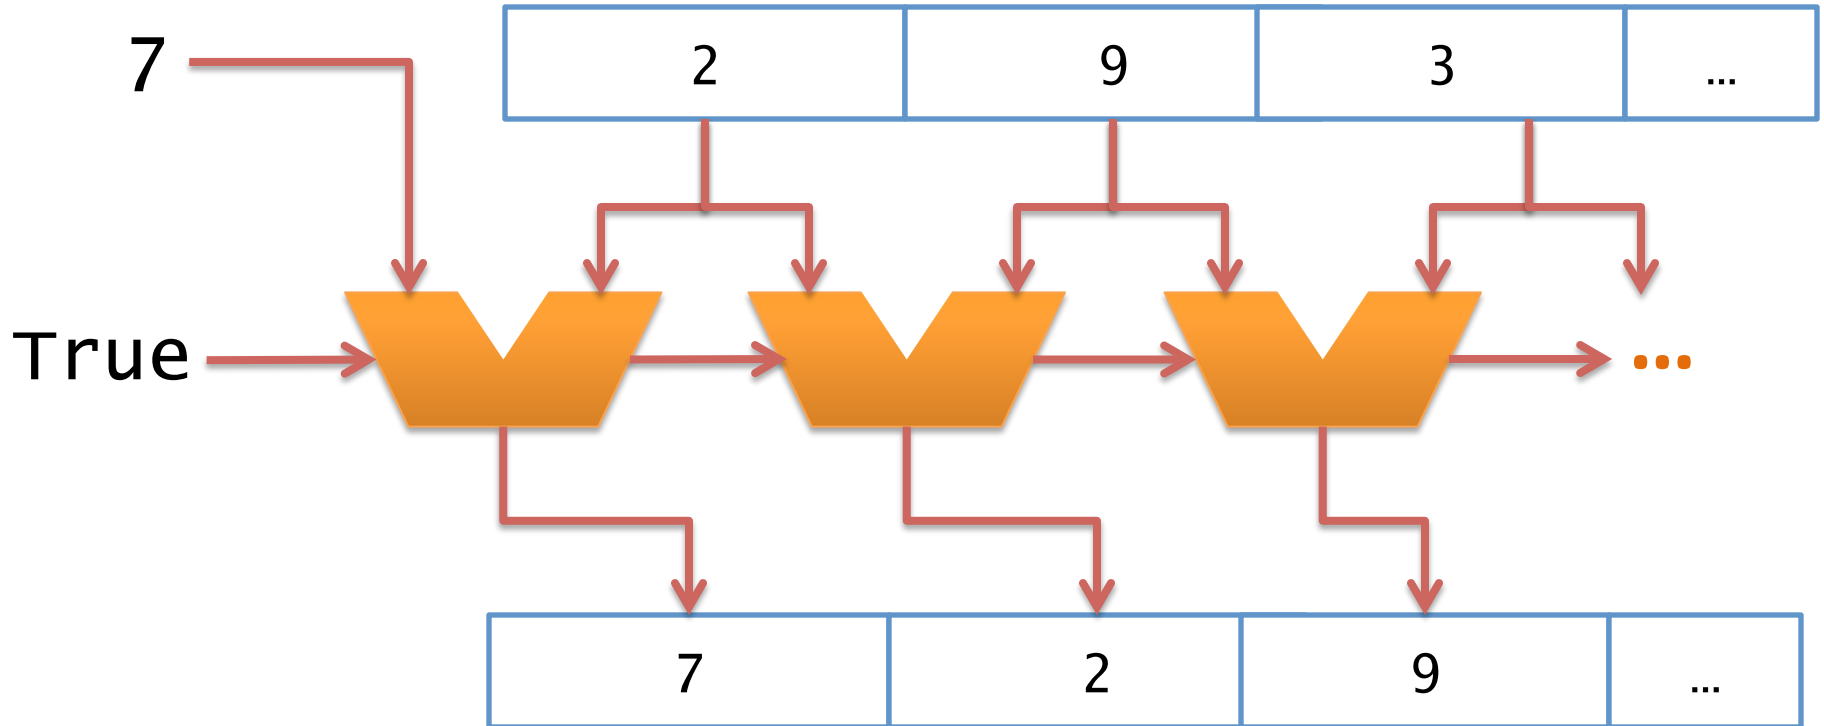


```
t := a.top() + 1
a.cond_update(x != 0, t)
a.cond_push(x != 0 && t > 10, *)
a.cond_update(x != 0, 5)
```

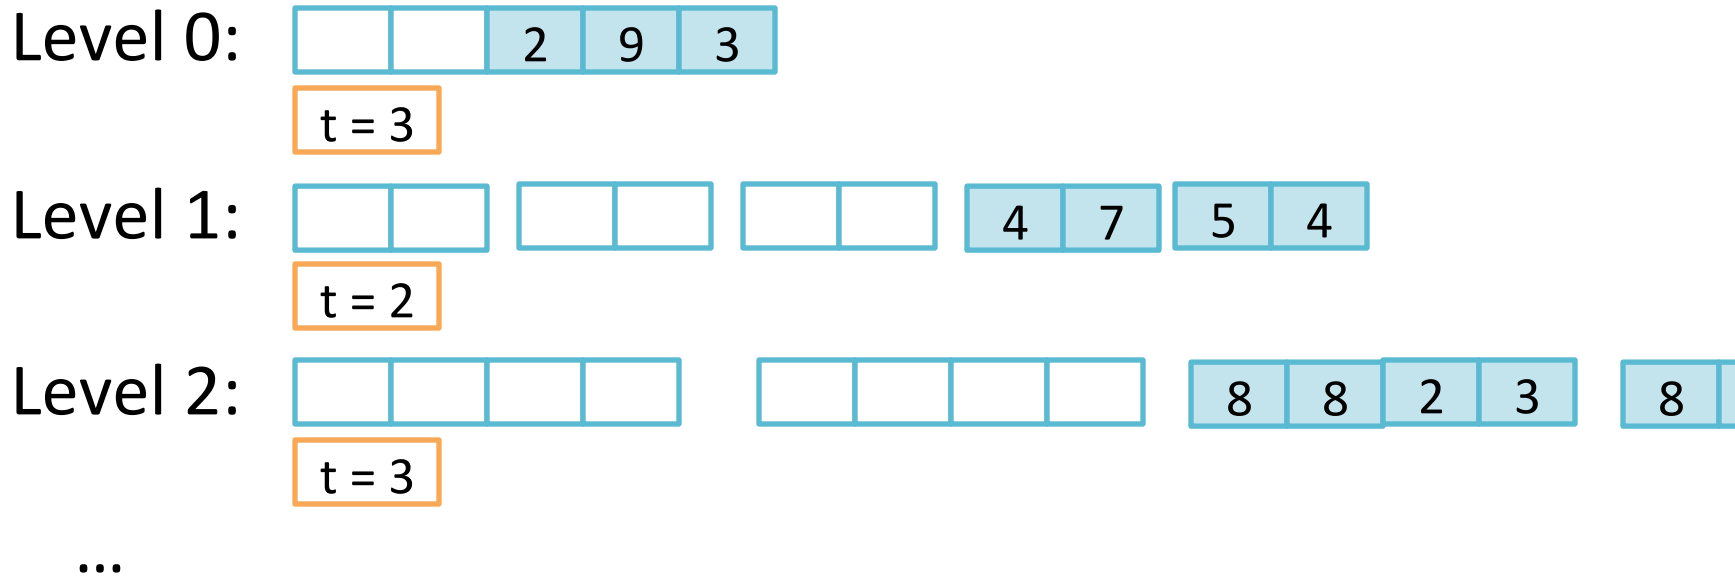
Naïve Conditional Push



Naïve Conditional Push



More Efficient Stack



Block size = 2^{level}

Each level has 5 blocks, at least 2 full and 2 empty

Level 0

Level 1

Level 2



t = 3



t = 2



t = 3

Conditional push (True, 7)



t = 4



t = 2



t = 3

Conditional push (True, 8)



t = 5



t = 2



t = 3

Shift



t = 3



t = 3



t = 3

Level 0

Level 1

Level



t = 3



t = 2



t = 3



t = 4


Amortized
 $\Theta(\log n)$ gates
 per operation



t = 5



t = 3



t = 3



t = 3



t = 3

Arbitrary Array Accesses (Associative Maps)

m[0] = 'A'
m[2] = 'U'
m[9] = 'M'
m[7] = 'R'
m[0] = 'D'
m[9] = 'Y'
m[9] = 'K'
m[7] = 'C'

0	2	7	9
'A'			
	'U'		
			'M'
		'R'	
'D'			
			'Y'
			'K'
		'C'	

Execution trace: **indexes and values are private values**

Batching Updates

m[0] = 'A'
m[2] = 'U'
m[9] = 'M'
m[7] = 'R'
m[0] = 'D'
m[9] = 'Y'
m[9] = 'K'
m[7] = 'C'



stable sort!

m[0] = 'A'
m[0] = 'D'
m[2] = 'U'
m[7] = 'R'
m[7] = 'C'
m[9] = 'M'
m[9] = 'Y'
m[9] = 'K'

Batching Updates

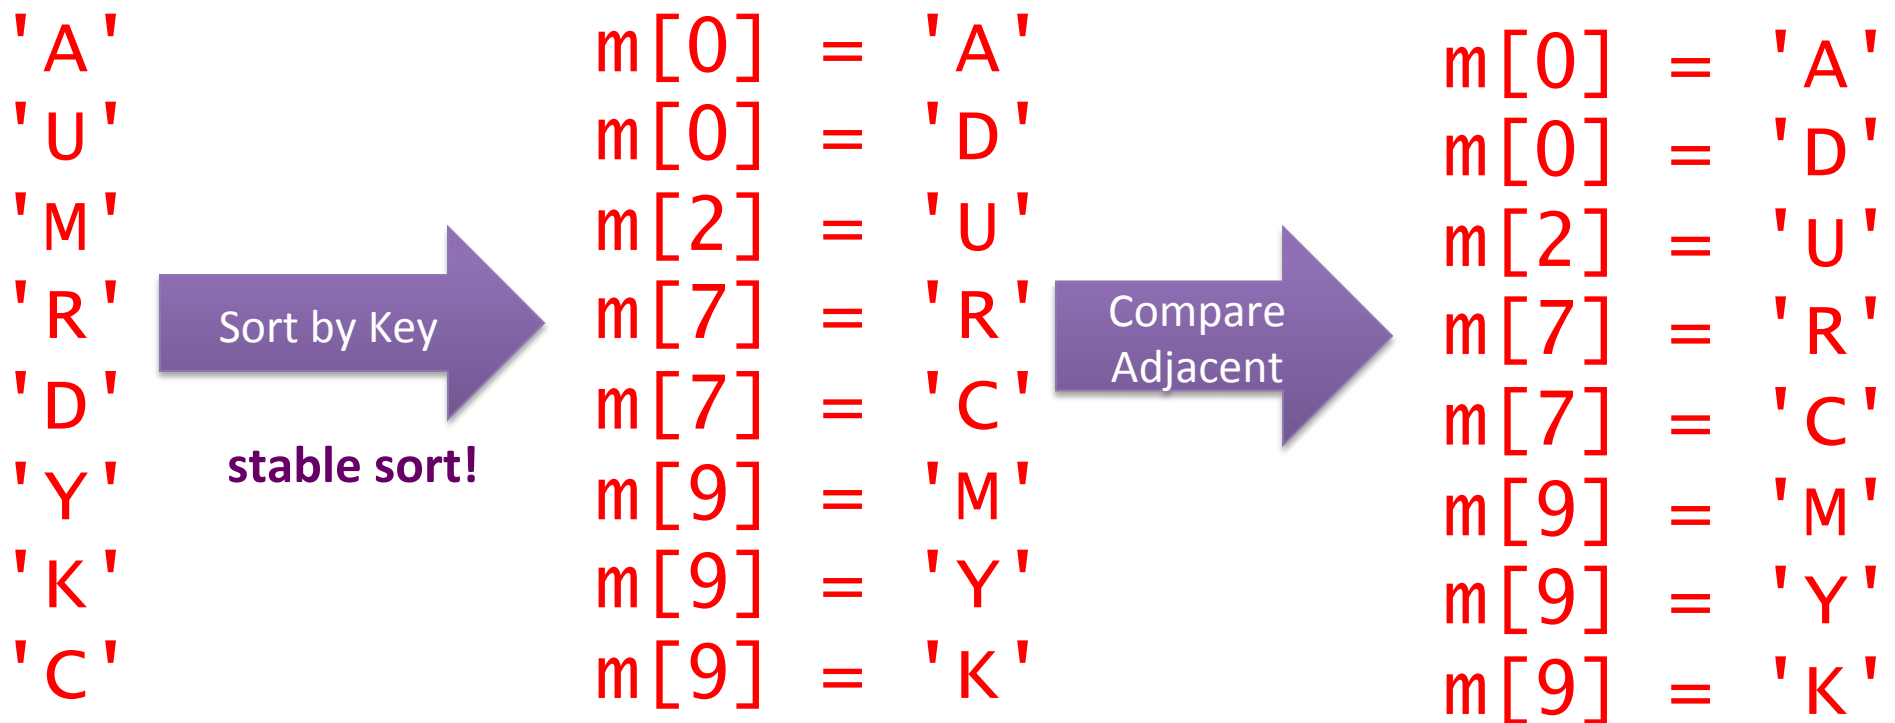
m[0] = 'A'
m[2] = 'U'
m[9] = 'M'
m[7] = 'R'
m[0] = 'D'
m[9] = 'Y'
m[9] = 'K'
m[7] = 'C'




stable sort!

m[0] = 'A'
m[0] = 'D'
m[2] = 'U'
m[7] = 'R'
m[7] = 'C'
m[9] = 'M'
m[9] = 'Y'
m[9] = 'K'


Batching Updates



Batching Updates



Batching Updates



Sort by Key

stable sort!

m[0] = 'A'
m[0] = 'D'
m[2] = 'U'
m[7] = 'R'
m[7] = 'C'
m[9] = 'M'
m[9] = 'Y'
m[9] = 'K'

Compare
Adjacent

m[0] = 'A'
m[0] = 'D'
m[2] = 'U'
m[7] = 'R'
m[7] = 'C'
m[9] = 'M'
m[9] = 'Y'
m[9] = 'K'

output
wires

m[0] = 'D'
m[2] = 'U'
m[7] = 'C'
m[9] = 'K'

Discarded

m[0] = 'A'
m[7] = 'R'
m[9] = 'M'
m[9] = 'Y'

Sort by
Liveness

Associative Map Cost

Oblivious Stable Sort

```
graph TD; A[Oblivious Stable Sort] --> B[Comparisons and Liveness Marking]; B --> C[Oblivious Sort Liveness/Key];
```

Comparisons and Liveness
Marking

Oblivious Sort
Liveness/Key

Circuit Size:

$\Theta(n \log n) \times$ comparison cost

$\Theta(n \log^2 n)$

Example Application: DBScan

Martin Ester, Hans-Peter Kriegel,
Jörg Sander, Xiaowei Xu. *KDD* 1996


Density-based clustering:
depth-first search to find dense clusters



Alice's Data



Bob's Data



Joint Clusters

$n \leftarrow |P|$

$c \leftarrow 0$

$s \leftarrow \text{emptyStack}$

$cluster \leftarrow [0, 0, \dots]$

for $i \leftarrow [1, n]$ **do**

if $cluster[i] \neq 0$ **then**

continue

$V \leftarrow \text{getNeighbors}(i, P, \text{minpts}, \text{radius})$

if $\text{count}(V) < \text{minpts}$ **then**

continue

$c \leftarrow c + 1$

▷ Start a new cluster

for $j \leftarrow [1, n]$ **do**

if $V[j] = \text{true} \wedge cluster[j] \neq 0$ **then**

$cluster[j] \leftarrow c$

$s.\text{push}(j)$



Conditional Push!

while $s \neq \emptyset$ **do**

$k \leftarrow s.\text{pop}()$

$V \leftarrow \text{getNeighbors}(k, P, \text{minpts}, \text{radius})$

if $\text{count}(V) < \text{minpts}$ **then**

continue

for $j \leftarrow [1, n]$ **do**

if $V[j] = \text{true} \wedge cluster[j] \neq 0$ **then**

$cluster[j] \leftarrow c$

$s.\text{push}(j)$

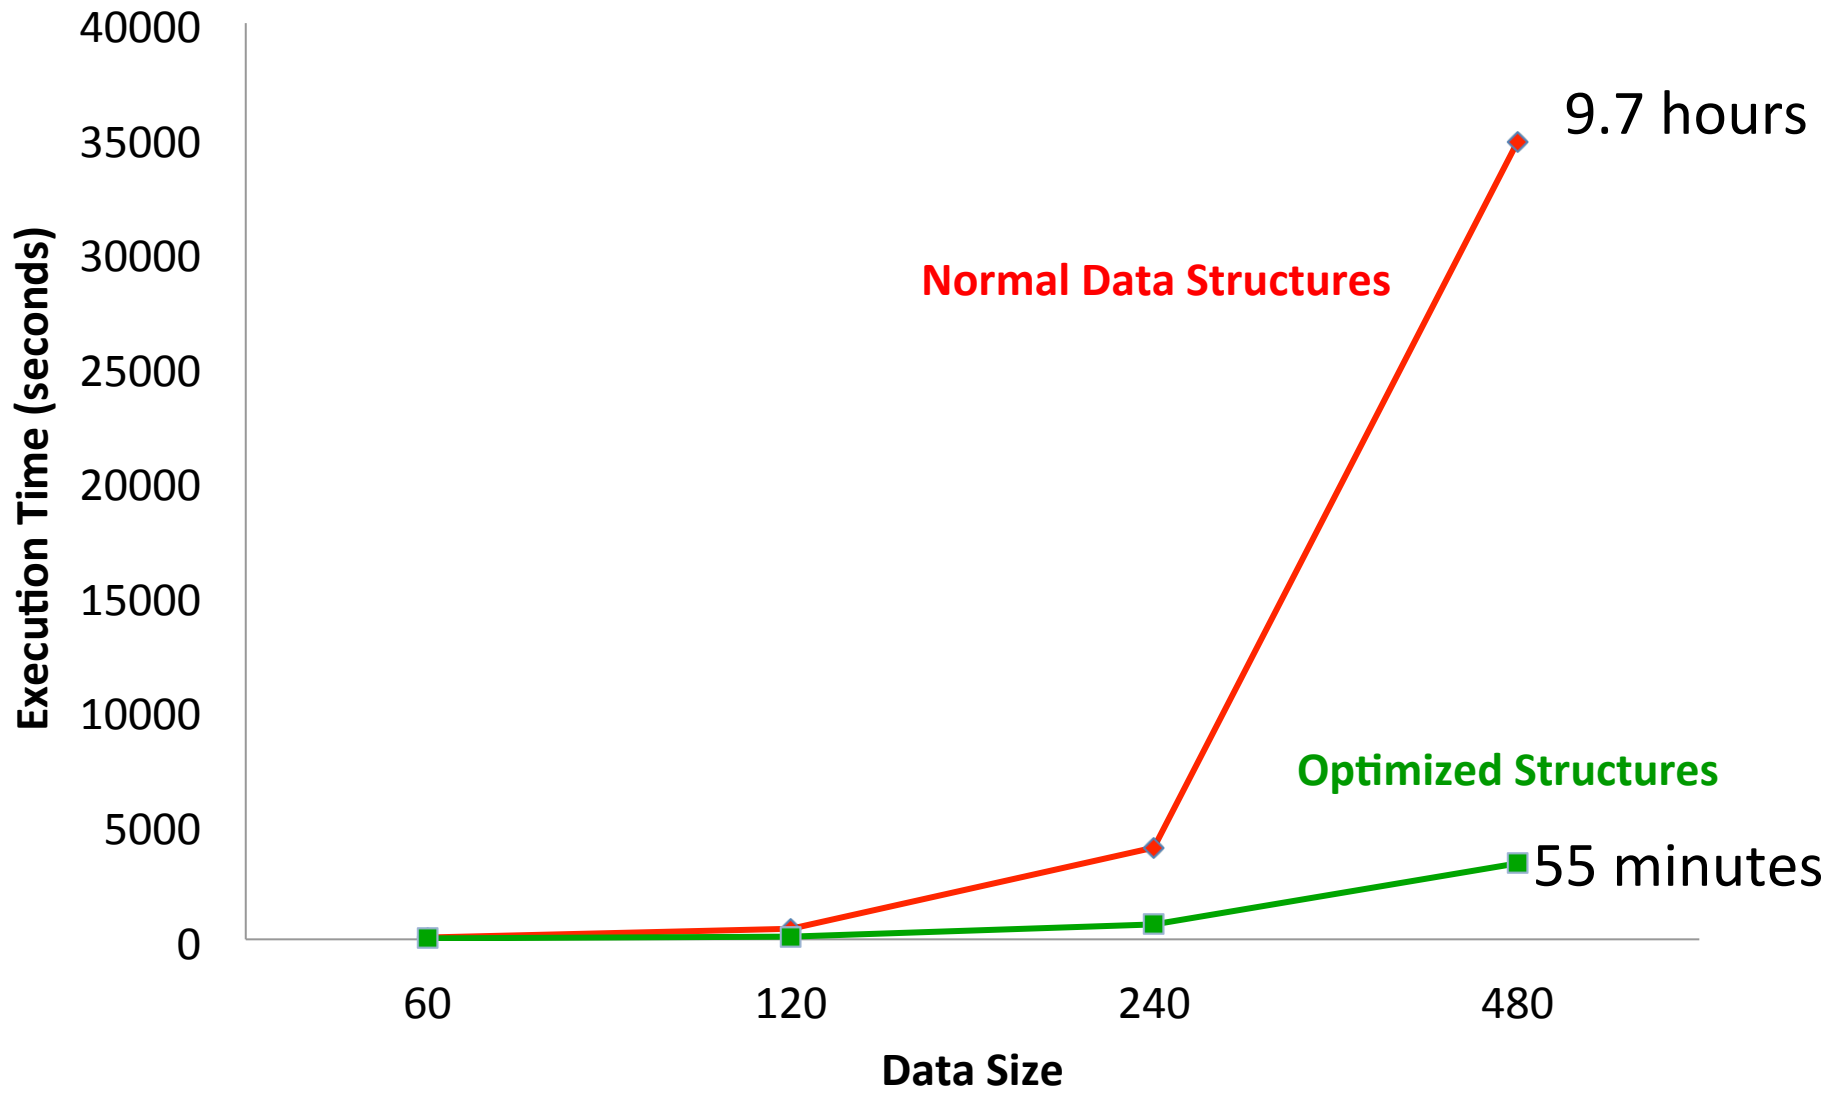



Array update!


Private Input: P – array of points (combines private points from both parties)


Public inputs: minpts , radius


Output: cluster number for each point











mightbeevil.com

Questions?