

Trapdoor simulation
of quantum algorithms

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

Joint work with:

Tung Chou

Technische Universiteit Eindhoven

Algorithms in CS courses

“WHAT is your algorithm?”

Trapdoor simulation
of quantum algorithms

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

Joint work with:

Tung Chou

Technische Universiteit Eindhoven

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

Trapdoor simulation
of quantum algorithms

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

Joint work with:

Tung Chou

Technische Universiteit Eindhoven

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

Trapdoor simulation
of quantum algorithms

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

Joint work with:

Tung Chou

Technische Universiteit Eindhoven

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

“It sorts the input array in place.
Here’s a proof.”

Trapdoor simulation
of quantum algorithms

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

Joint work with:

Tung Chou

Technische Universiteit Eindhoven

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

“It sorts the input array in place.
Here’s a proof.”

“WHAT is its run time?”

Trapdoor simulation
of quantum algorithms

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

Joint work with:

Tung Chou

Technische Universiteit Eindhoven

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

“It sorts the input array in place.
Here’s a proof.”

“WHAT is its run time?”

“ $O(n \lg n)$ comparisons;
and $\Theta(n \lg n)$ comparisons
for most inputs. Here’s a proof.”

Trapdoor simulation
of quantum algorithms

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

Joint work with:

Tung Chou

Technische Universiteit Eindhoven

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

“It sorts the input array in place.
Here’s a proof.”

“WHAT is its run time?”

“ $O(n \lg n)$ comparisons;
and $\Theta(n \lg n)$ comparisons
for most inputs. Here’s a proof.”

“You may pass.”

or simulation

um algorithms

. Bernstein

ty of Illinois at Chicago &

he Universiteit Eindhoven

ork with:

ou

he Universiteit Eindhoven

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

“It sorts the input array in place.
Here’s a proof.”

“WHAT is its run time?”

“ $O(n \lg n)$ comparisons;
and $\Theta(n \lg n)$ comparisons
for most inputs. Here’s a proof.”

“You may pass.”

Algorithms

Critical c

How har

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

“It sorts the input array in place.
Here’s a proof.”

“WHAT is its run time?”

“ $O(n \lg n)$ comparisons;
and $\Theta(n \lg n)$ comparisons
for most inputs. Here’s a proof.”

“You may pass.”

Algorithms for hard

Critical question for

How hard is ECDL

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

“It sorts the input array in place.
Here’s a proof.”

“WHAT is its run time?”

“ $O(n \lg n)$ comparisons;
and $\Theta(n \lg n)$ comparisons
for most inputs. Here’s a proof.”

“You may pass.”

Algorithms for hard problems

Critical question for ECC security

How hard is ECDLP?

ago &
hoven

hoven

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

“It sorts the input array in place.
Here’s a proof.”

“WHAT is its run time?”

“ $O(n \lg n)$ comparisons;
and $\Theta(n \lg n)$ comparisons
for most inputs. Here’s a proof.”

“You may pass.”

Algorithms for hard problems

Critical question for ECC security:

How hard is ECDLP?

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

“It sorts the input array in place.
Here’s a proof.”

“WHAT is its run time?”

“ $O(n \lg n)$ comparisons;
and $\Theta(n \lg n)$ comparisons
for most inputs. Here’s a proof.”

“You may pass.”

Algorithms for hard problems

Critical question for ECC security:
How hard is ECDLP?

Standard estimate for “strong”
ECC groups of prime order ℓ :
Latest “negating” variants of
“distinguished point” rho methods
break an average ECDLP instance
using $\approx 0.886\sqrt{\ell}$ additions.

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

“It sorts the input array in place.
Here’s a proof.”

“WHAT is its run time?”

“ $O(n \lg n)$ comparisons;
and $\Theta(n \lg n)$ comparisons
for most inputs. Here’s a proof.”

“You may pass.”

Algorithms for hard problems

Critical question for ECC security:
How hard is ECDLP?

Standard estimate for “strong”
ECC groups of prime order ℓ :
Latest “negating” variants of
“distinguished point” rho methods
break an average ECDLP instance
using $\approx 0.886\sqrt{\ell}$ additions.

Is this proven? No!

Is this provable? Maybe not!

Algorithms in CS courses

“WHAT is your algorithm?”

“Heapsort. Here’s the code.”

“WHAT does it accomplish?”

“It sorts the input array in place.
Here’s a proof.”

“WHAT is its run time?”

“ $O(n \lg n)$ comparisons;
and $\Theta(n \lg n)$ comparisons
for most inputs. Here’s a proof.”

“You may pass.”

Algorithms for hard problems

Critical question for ECC security:
How hard is ECDLP?

Standard estimate for “strong”
ECC groups of prime order ℓ :
Latest “negating” variants of
“distinguished point” rho methods
break an average ECDLP instance
using $\approx 0.886\sqrt{\ell}$ additions.

Is this proven? No!

Is this provable? Maybe not!

So why do we think it’s true?

Algorithms in CS courses

“What is your algorithm?”

“I don’t know. Here’s the code.”

“What does it accomplish?”

“It sorts the input array in place. I have a proof.”

“What is its run time?”

“ $O(n)$ comparisons;

$O(\lg n)$ comparisons

for n inputs. Here’s a proof.”

“How can it possibly pass.”

Algorithms for hard problems

Critical question for ECC security:
How hard is ECDLP?

Standard estimate for “strong”
ECC groups of prime order ℓ :
Latest “negating” variants of
“distinguished point” rho methods
break an average ECDLP instance
using $\approx 0.886\sqrt{\ell}$ additions.

Is this proven? No!

Is this provable? Maybe not!

So why do we think it’s true?

2000 Ga
inadequa
of a neg

courses

algorithm?"

the code."

accomplish?"

array in place.

time?"

isons;

parisons

here's a proof."

Algorithms for hard problems

Critical question for ECC security:
How hard is ECDLP?

Standard estimate for "strong"
ECC groups of prime order ℓ :
Latest "negating" variants of
"distinguished point" rho methods
break an average ECDLP instance
using $\approx 0.886\sqrt{\ell}$ additions.

Is this proven? No!

Is this provable? Maybe not!

So why do we think it's true?

2000 Gallant–Lam
inadequately speci
of a negating rho

Algorithms for hard problems

Critical question for ECC security:
How hard is ECDLP?

Standard estimate for “strong”
ECC groups of prime order ℓ :
Latest “negating” variants of
“distinguished point” rho methods
break an average ECDLP instance
using $\approx 0.886\sqrt{\ell}$ additions.

Is this proven? No!

Is this provable? Maybe not!

So why do we think it's true?

2000 Gallant–Lambert–VanS
inadequately specified stater
of a negating rho algorithm.

Algorithms for hard problems

Critical question for ECC security:
How hard is ECDLP?

Standard estimate for “strong”
ECC groups of prime order ℓ :
Latest “negating” variants of
“distinguished point” rho methods
break an average ECDLP instance
using $\approx 0.886\sqrt{\ell}$ additions.

Is this proven? No!

Is this provable? Maybe not!

So why do we think it's true?

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

Algorithms for hard problems

Critical question for ECC security:
How hard is ECDLP?

Standard estimate for “strong”
ECC groups of prime order ℓ :
Latest “negating” variants of
“distinguished point” rho methods
break an average ECDLP instance
using $\approx 0.886\sqrt{\ell}$ additions.

Is this proven? No!

Is this provable? Maybe not!

So why do we think it's true?

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

2010 Bos–Kleinjung–Lenstra:
a plausible interpretation of
that algorithm is *non-functional*.

Algorithms for hard problems

Critical question for ECC security:
How hard is ECDLP?

Standard estimate for “strong”
ECC groups of prime order ℓ :
Latest “negating” variants of
“distinguished point” rho methods
break an average ECDLP instance
using $\approx 0.886\sqrt{\ell}$ additions.

Is this proven? No!

Is this provable? Maybe not!

So why do we think it's true?

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

2010 Bos–Kleinjung–Lenstra:
a plausible interpretation of
that algorithm is *non-functional*.

See 2011 Bernstein–Lange–
Schwabe for more history
and better algorithms.

Algorithms for hard problems

Critical question for ECC security:
How hard is ECDLP?

Standard estimate for “strong”
ECC groups of prime order ℓ :
Latest “negating” variants of
“distinguished point” rho methods
break an average ECDLP instance
using $\approx 0.886\sqrt{\ell}$ additions.

Is this proven? No!

Is this provable? Maybe not!

So why do we think it's true?

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

2010 Bos–Kleinjung–Lenstra:
a plausible interpretation of
that algorithm is *non-functional*.

See 2011 Bernstein–Lange–
Schwabe for more history
and better algorithms.

Why do we believe that
the latest algorithms work
at the claimed speeds?

Experiments!

Algorithms for hard problems

Open question for ECC security:
What is ECDLP?

Best estimate for “strong”
groups of prime order ℓ :
“negating” variants of
“washed point” rho methods
on average ECDLP instance
costs $0.886\sqrt{\ell}$ additions.

Proven? No!

Provable? Maybe not!

How do we think it's true?

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

2010 Bos–Kleinjung–Lenstra:
a plausible interpretation of
that algorithm is *non-functional*.

See 2011 Bernstein–Lange–
Schwabe for more history
and better algorithms.

Why do we believe that
the latest algorithms work
at the claimed speeds?

Experiments!

Similar to
we don't
best fact

and problems

for ECC security:
CDLP?

for “strong”

of order ℓ :

variants of

“strong” rho methods

for ECDLP instance

under conditions.

o!

Maybe not!

Think it's true?

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

2010 Bos–Kleinjung–Lenstra:
a plausible interpretation of
that algorithm is *non-functional*.

See 2011 Bernstein–Lange–
Schwabe for more history
and better algorithms.

Why do we believe that
the latest algorithms work
at the claimed speeds?

Experiments!

Similar story for R
we don't have pro
best factoring algo

s
curity:
ng"
:
of
ethods
stance
!
e?

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

2010 Bos–Kleinjung–Lenstra:
a plausible interpretation of
that algorithm is *non-functional*.

See 2011 Bernstein–Lange–
Schwabe for more history
and better algorithms.

Why do we believe that
the latest algorithms work
at the claimed speeds?

Experiments!

Similar story for RSA security
we don't have proofs for the
best factoring algorithms.

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

2010 Bos–Kleinjung–Lenstra:
a plausible interpretation of
that algorithm is *non-functional*.

See 2011 Bernstein–Lange–
Schwabe for more history
and better algorithms.

Why do we believe that
the latest algorithms work
at the claimed speeds?

Experiments!

Similar story for RSA security:
we don't have proofs for the
best factoring algorithms.

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

2010 Bos–Kleinjung–Lenstra:
a plausible interpretation of
that algorithm is *non-functional*.

See 2011 Bernstein–Lange–
Schwabe for more history
and better algorithms.

Why do we believe that
the latest algorithms work
at the claimed speeds?

Experiments!

Similar story for RSA security:
we don't have proofs for the
best factoring algorithms.

Code-based cryptography:
we don't have proofs for the
best decoding algorithms.

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

2010 Bos–Kleinjung–Lenstra:
a plausible interpretation of
that algorithm is *non-functional*.

See 2011 Bernstein–Lange–
Schwabe for more history
and better algorithms.

Why do we believe that
the latest algorithms work
at the claimed speeds?

Experiments!

Similar story for RSA security:
we don't have proofs for the
best factoring algorithms.

Code-based cryptography:
we don't have proofs for the
best decoding algorithms.

Lattice-based cryptography:
we don't have proofs for the
best lattice algorithms.

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

2010 Bos–Kleinjung–Lenstra:
a plausible interpretation of
that algorithm is *non-functional*.

See 2011 Bernstein–Lange–
Schwabe for more history
and better algorithms.

Why do we believe that
the latest algorithms work
at the claimed speeds?

Experiments!

Similar story for RSA security:
we don't have proofs for the
best factoring algorithms.

Code-based cryptography:
we don't have proofs for the
best decoding algorithms.

Lattice-based cryptography:
we don't have proofs for the
best lattice algorithms.

MQ-based cryptography:
we don't have proofs for the
best system-solving algorithms.

2000 Gallant–Lambert–Vanstone:
inadequately specified statement
of a negating rho algorithm.

2010 Bos–Kleinjung–Lenstra:
a plausible interpretation of
that algorithm is *non-functional*.

See 2011 Bernstein–Lange–
Schwabe for more history
and better algorithms.

Why do we believe that
the latest algorithms work
at the claimed speeds?

Experiments!

Similar story for RSA security:
we don't have proofs for the
best factoring algorithms.

Code-based cryptography:
we don't have proofs for the
best decoding algorithms.

Lattice-based cryptography:
we don't have proofs for the
best lattice algorithms.

MQ-based cryptography:
we don't have proofs for the
best system-solving algorithms.

Confidence relies on experiments.

Illant–Lambert–Vanstone:
vaguely specified statement
regarding rho algorithm.

Shor–Kleinjung–Lenstra:
multiple interpretation of
algorithm is *non-functional*.

1 Bernstein–Lange–
for more history
of factor algorithms.

we believe that
fast algorithms work
at claimed speeds?

Comments!

Similar story for RSA security:
we don't have proofs for the
best factoring algorithms.

Code-based cryptography:
we don't have proofs for the
best decoding algorithms.

Lattice-based cryptography:
we don't have proofs for the
best lattice algorithms.

MQ-based cryptography:
we don't have proofs for the
best system-solving algorithms.

Confidence relies on experiments.

Where's

Quantum
is moving
into algo

Example
exponent

Bernstei

Don't ex
for the b

to attack

How do
in analys

Quantum

Shor–Vanstone:
verified statement
algorithm.

Heath–Lenstra:
verification of
non-functional.

Shor–Lange–
history
algorithms.

Are there
algorithms work
needed?

Similar story for RSA security:
we don't have proofs for the
best factoring algorithms.

Code-based cryptography:
we don't have proofs for the
best decoding algorithms.

Lattice-based cryptography:
we don't have proofs for the
best lattice algorithms.

MQ-based cryptography:
we don't have proofs for the
best system-solving algorithms.

Confidence relies on experiments.

Where's my quantum

Quantum-algorithms
is moving beyond
into algorithms with

Example: subset-sum
exponent ≈ 0.241

Bernstein–Jefferys–

Don't expect proofs
for the best quantum
to attack post-quantum

How do we obtain
in analysis of these

Quantum experiments

stone:

ment

n:

onal.

Similar story for RSA security:

we don't have proofs for the best factoring algorithms.

Code-based cryptography:

we don't have proofs for the best decoding algorithms.

Lattice-based cryptography:

we don't have proofs for the best lattice algorithms.

MQ-based cryptography:

we don't have proofs for the best system-solving algorithms.

Confidence relies on experiments.

Where's my quantum computer?

Quantum-algorithm design is moving beyond textbook solutions into algorithms without proofs.

Example: subset-sum

exponent ≈ 0.241 from 2013

Bernstein–Jeffery–Lange–Me

Don't expect proofs or provability

for the best quantum algorithm

to attack post-quantum crypt

How do we obtain confidence

in analysis of these algorithms?

Quantum experiments are hard

Similar story for RSA security:
we don't have proofs for the
best factoring algorithms.

Code-based cryptography:
we don't have proofs for the
best decoding algorithms.

Lattice-based cryptography:
we don't have proofs for the
best lattice algorithms.

MQ-based cryptography:
we don't have proofs for the
best system-solving algorithms.

Confidence relies on experiments.

Where's my quantum computer?

Quantum-algorithm design
is moving beyond textbook stage
into algorithms without proofs.

Example: subset-sum
exponent ≈ 0.241 from 2013
Bernstein–Jeffery–Lange–Meurer.

Don't expect proofs or provability
for the best quantum algorithms
to attack post-quantum crypto.

How do we obtain confidence
in analysis of these algorithms?

Quantum experiments are hard.

story for RSA security:

we have proofs for the
factoring algorithms.

lattice-based cryptography:

we have proofs for the
decoding algorithms.

code-based cryptography:

we have proofs for the
decoding algorithms.

isogeny-based cryptography:

we have proofs for the
isogeny-solving algorithms.

confidence relies on experiments.

Where's my quantum computer?

Quantum-algorithm design
is moving beyond textbook stage
into algorithms without proofs.

Example: subset-sum

exponent ≈ 0.241 from 2013

Bernstein–Jeffery–Lange–Meurer.

Don't expect proofs or provability
for the best quantum algorithms
to attack post-quantum crypto.

How do we obtain confidence
in analysis of these algorithms?

Quantum experiments are hard.

Where's

Analogy
a 2^{80} NP

SA security:

Proofs for the
algorithms.

Cryptography:

Proofs for the
algorithms.

Cryptography:

Proofs for the
algorithms.

Cryptography:

Proofs for the
new algorithms.

Quantum experiments.

Where's my quantum computer?

Quantum-algorithm design
is moving beyond textbook stage
into algorithms without proofs.

Example: subset-sum

exponent ≈ 0.241 from 2013

Bernstein–Jeffery–Lange–Meurer.

Don't expect proofs or provability
for the best quantum algorithms
to attack post-quantum crypto.

How do we obtain confidence
in analysis of these algorithms?

Quantum experiments are hard.

Where's my big crypto?

Analogy: Public key
a 2^{80} NFS RSA-1024

Where's my quantum computer?

Quantum-algorithm design
is moving beyond textbook stage
into algorithms without proofs.

Example: subset-sum
exponent ≈ 0.241 from 2013
Bernstein–Jeffery–Lange–Meurer.

Don't expect proofs or provability
for the best quantum algorithms
to attack post-quantum crypto.

How do we obtain confidence
in analysis of these algorithms?
Quantum experiments are hard.

Where's my big computer?

Analogy: Public hasn't carried
a 2^{80} NFS RSA-1024 experi

Where's my quantum computer?

Quantum-algorithm design is moving beyond textbook stage into algorithms without proofs.

Example: subset-sum
exponent ≈ 0.241 from 2013
Bernstein–Jeffery–Lange–Meurer.

Don't expect proofs or provability for the best quantum algorithms to attack post-quantum crypto.

How do we obtain confidence in analysis of these algorithms?
Quantum experiments are hard.

Where's my big computer?

Analogy: Public hasn't carried out a 2^{80} NFS RSA-1024 experiment.

Where's my quantum computer?

Quantum-algorithm design is moving beyond textbook stage into algorithms without proofs.

Example: subset-sum

exponent ≈ 0.241 from 2013

Bernstein–Jeffery–Lange–Meurer.

Don't expect proofs or provability for the best quantum algorithms to attack post-quantum crypto.

How do we obtain confidence in analysis of these algorithms?

Quantum experiments are hard.

Where's my big computer?

Analogy: Public hasn't carried out a 2^{80} NFS RSA-1024 experiment.

But public has carried out 2^{50} , 2^{60} , 2^{70} NFS experiments.

Hopefully not too much extrapolation error for 2^{80} .

Where's my quantum computer?

Quantum-algorithm design is moving beyond textbook stage into algorithms without proofs.

Example: subset-sum
exponent ≈ 0.241 from 2013
Bernstein–Jeffery–Lange–Meurer.

Don't expect proofs or provability for the best quantum algorithms to attack post-quantum crypto.

How do we obtain confidence in analysis of these algorithms?
Quantum experiments are hard.

Where's my big computer?

Analogy: Public hasn't carried out a 2^{80} NFS RSA-1024 experiment.

But public has carried out 2^{50} , 2^{60} , 2^{70} NFS experiments.
Hopefully not too much extrapolation error for 2^{80} .

Vastly larger extrapolation for the quantum situation.

Imagine attacker performing 2^{80} operations on 2^{40} qubits; compare to today's challenges of 2^1 , 2^2 , 2^3 , 2^4 , 2^5 , 2^6 qubits.

Where's my quantum computer?

Quantum algorithm design
going beyond textbook stage
algorithms without proofs.

Example: subset-sum

Time ≈ 0.241 from 2013

by Brønner–Jeffery–Lange–Meurer.

Expect proofs or provability
for best quantum algorithms
to break post-quantum crypto.

How do we obtain confidence
in the analysis of these algorithms?
Quantum experiments are hard.

Where's my big computer?

Analogy: Public hasn't carried out
a 2^{80} NFS RSA-1024 experiment.

But public has carried out
 2^{50} , 2^{60} , 2^{70} NFS experiments.

Hopefully not too much
extrapolation error for 2^{80} .

Vastly larger extrapolation
for the quantum situation.

Imagine attacker performing
 2^{80} operations on 2^{40} qubits;
compare to today's challenges
of 2^1 , 2^2 , 2^3 , 2^4 , 2^5 , 2^6 qubits.

Simulation

An algorithm
is a combination
of the algorithm
for a particular

Quantum computer?

Quantum design
textbook stage
without proofs.

Quantum
from 2013
-Lange–Meurer.

Proofs or provability
Quantum algorithms
Quantum crypto.

Confidence
in algorithms?
Problems are hard.

Where's my big computer?

Analogy: Public hasn't carried out
a 2^{80} NFS RSA-1024 experiment.

But public has carried out
 2^{50} , 2^{60} , 2^{70} NFS experiments.

Hopefully not too much
extrapolation error for 2^{80} .

Vastly larger extrapolation
for the quantum situation.

Imagine attacker performing
 2^{80} operations on 2^{40} qubits;
compare to today's challenges
of 2^1 , 2^2 , 2^3 , 2^4 , 2^5 , 2^6 qubits.

Simulation

An algorithm simulation
is a computer-assisted
of the algorithm's
for a particular input

Where's my big computer?

Analogy: Public hasn't carried out a 2^{80} NFS RSA-1024 experiment.

But public has carried out 2^{50} , 2^{60} , 2^{70} NFS experiments.

Hopefully not too much extrapolation error for 2^{80} .

Vastly larger extrapolation for the quantum situation.

Imagine attacker performing 2^{80} operations on 2^{40} qubits; compare to today's challenges of 2^1 , 2^2 , 2^3 , 2^4 , 2^5 , 2^6 qubits.

Simulation

An algorithm simulation is a computer-assisted proof of the algorithm's performance *for a particular input*.

Where's my big computer?

Analogy: Public hasn't carried out a 2^{80} NFS RSA-1024 experiment.

But public has carried out 2^{50} , 2^{60} , 2^{70} NFS experiments.

Hopefully not too much extrapolation error for 2^{80} .

Vastly larger extrapolation for the quantum situation.

Imagine attacker performing 2^{80} operations on 2^{40} qubits; compare to today's challenges of 2^1 , 2^2 , 2^3 , 2^4 , 2^5 , 2^6 qubits.

Simulation

An algorithm simulation is a computer-assisted proof of the algorithm's performance *for a particular input*.

Where's my big computer?

Analogy: Public hasn't carried out a 2^{80} NFS RSA-1024 experiment.

But public has carried out 2^{50} , 2^{60} , 2^{70} NFS experiments.

Hopefully not too much extrapolation error for 2^{80} .

Vastly larger extrapolation for the quantum situation.

Imagine attacker performing 2^{80} operations on 2^{40} qubits; compare to today's challenges of 2^1 , 2^2 , 2^3 , 2^4 , 2^5 , 2^6 qubits.

Simulation

An algorithm simulation is a computer-assisted proof of the algorithm's performance *for a particular input*.

Compared to traditional proofs:

Theorem statement is easier.

Steps in proof are easier.

Don't need to generalize beyond a single input.

Provability is guaranteed.

Proof has computer assistance, so less chance of error.

my big computer?

: Public hasn't carried out
NFS RSA-1024 experiment.

Public has carried out
 2^{70} NFS experiments.

Why not too much
simulation error for 2^{80} .

Larger extrapolation
quantum situation.

Attacker performing
operations on 2^{40} qubits;

vs today's challenges
 $2^2, 2^3, 2^4, 2^5, 2^6$ qubits.

Simulation

An algorithm simulation
is a computer-assisted proof
of the algorithm's performance
for a particular input.

Compared to traditional proofs:

Theorem statement is easier.

Steps in proof are easier.

Don't need to generalize
beyond a single input.

Provability is guaranteed.

Proof has computer assistance,
so less chance of error.

The start
of an alg

Comput
and t_0, t
such tha
algorithm

Prove th
matches

Special c
The com

the origi
plus prin

Particula

computer?

hasn't carried out
2024 experiment.

carried out
experiments.

much
for 2^{80} .

simulation
situation.

performing
 2^{40} qubits;
challenges
 $2^5, 2^6$ qubits.

Simulation

An algorithm simulation
is a computer-assisted proof
of the algorithm's performance
for a particular input.

Compared to traditional proofs:

Theorem statement is easier.

Steps in proof are easier.

Don't need to generalize
beyond a single input.

Provability is guaranteed.

Proof has computer assistance,
so less chance of error.

The standard structure
of an algorithm simulation

Compute s_0, s_1, s_2, \dots
and t_0, t_1, t_2, \dots

such that s_i represents
algorithm state at time t_i

Prove that the computed state
matches the original state

Special case: experimental
The computation is compared to
the original algorithm's output

plus printouts of state at each step

Particularly easy to verify

Simulation

An algorithm simulation is a computer-assisted proof of the algorithm's performance *for a particular input*.

Compared to traditional proofs:

Theorem statement is easier.

Steps in proof are easier.

Don't need to generalize beyond a single input.

Provability is guaranteed.

Proof has computer assistance, so less chance of error.

The standard structure of an algorithm simulation:

Compute s_0, s_1, s_2, \dots

and t_0, t_1, t_2, \dots

such that s_i represents algorithm state at time t_i .

Prove that the computation matches the original algorithm.

Special case: experiment.

The computation *is* the original algorithm plus printouts of state.

Particularly easy proof.

Simulation

An algorithm simulation is a computer-assisted proof of the algorithm's performance *for a particular input*.

Compared to traditional proofs:

Theorem statement is easier.

Steps in proof are easier.

Don't need to generalize beyond a single input.

Provability is guaranteed.

Proof has computer assistance, so less chance of error.

The standard structure of an algorithm simulation:

Compute s_0, s_1, s_2, \dots

and t_0, t_1, t_2, \dots

such that s_i represents algorithm state at time t_i .

Prove that the computation matches the original algorithm.

Special case: experiment.

The computation *is* the original algorithm plus printouts of state.

Particularly easy proof.

on

Algorithm simulation

Computer-assisted proof

Algorithm's performance

Particular input.

Compared to traditional proofs:

Proving a statement is easier.

Writing a proof are easier.

Need to generalize

From a single input.

Correctness is guaranteed.

Requires computer assistance,

Small chance of error.

The standard structure
of an algorithm simulation:

Compute s_0, s_1, s_2, \dots

and t_0, t_1, t_2, \dots

such that s_i represents
algorithm state at time t_i .

Prove that the computation
matches the original algorithm.

Special case: experiment.

The computation *is*
the original algorithm
plus printouts of state.
Particularly easy proof.

Simulation

“If you can

simulate a quantum

algorithm, you can

simulate a pre-quantum

algorithm.

have an algorithm

Simulation
structured proof
performance
output.

Additional proofs:

is easier.

easier.

generalize

output.

guaranteed.

with assistance,

error.

The standard structure
of an algorithm simulation:

Compute s_0, s_1, s_2, \dots

and t_0, t_1, t_2, \dots

such that s_i represents
algorithm state at time t_i .

Prove that the computation
matches the original algorithm.

Special case: experiment.

The computation *is*
the original algorithm
plus printouts of state.
Particularly easy proof.

Simulation of quantum

“If you can efficiently
simulate a quantum algorithm
using only pre-quantum computation,
then you have an efficient proof
of the correctness of the algorithm for the simulation.”

The standard structure
of an algorithm simulation:

Compute s_0, s_1, s_2, \dots

and t_0, t_1, t_2, \dots

such that s_i represents
algorithm state at time t_i .

Prove that the computation
matches the original algorithm.

Special case: experiment.

The computation *is*
the original algorithm
plus printouts of state.

Particularly easy proof.

Simulation of quantum algo

“If you can efficiently simulate
a quantum algorithm using a
pre-quantum computer then
you have an efficient pre-quantum
algorithm for the same problem.”

The standard structure
of an algorithm simulation:

Compute s_0, s_1, s_2, \dots

and t_0, t_1, t_2, \dots

such that s_i represents
algorithm state at time t_i .

Prove that the computation
matches the original algorithm.

Special case: experiment.

The computation *is*
the original algorithm
plus printouts of state.

Particularly easy proof.

Simulation of quantum algorithms

“If you can efficiently simulate
a quantum algorithm using a
pre-quantum computer then you
have an efficient pre-quantum
algorithm for the same problem.”

The standard structure
of an algorithm simulation:

Compute s_0, s_1, s_2, \dots

and t_0, t_1, t_2, \dots

such that s_i represents
algorithm state at time t_i .

Prove that the computation
matches the original algorithm.

Special case: experiment.

The computation *is*
the original algorithm
plus printouts of state.

Particularly easy proof.

Simulation of quantum algorithms

“If you can efficiently simulate
a quantum algorithm using a
pre-quantum computer then you
have an efficient pre-quantum
algorithm for the same problem.”

No, not necessarily!

The standard structure
of an algorithm simulation:

Compute s_0, s_1, s_2, \dots

and t_0, t_1, t_2, \dots

such that s_i represents
algorithm state at time t_i .

Prove that the computation
matches the original algorithm.

Special case: experiment.

The computation *is*
the original algorithm
plus printouts of state.

Particularly easy proof.

Simulation of quantum algorithms

“If you can efficiently simulate
a quantum algorithm using a
pre-quantum computer then you
have an efficient pre-quantum
algorithm for the same problem.”

No, not necessarily!

“Yes, you do! Simply run the
simulation on the same input and
extract the original algorithm’s
output from the final state.”

The standard structure
of an algorithm simulation:

Compute s_0, s_1, s_2, \dots

and t_0, t_1, t_2, \dots

such that s_i represents
algorithm state at time t_i .

Prove that the computation
matches the original algorithm.

Special case: experiment.

The computation *is*
the original algorithm
plus printouts of state.

Particularly easy proof.

Simulation of quantum algorithms

“If you can efficiently simulate
a quantum algorithm using a
pre-quantum computer then you
have an efficient pre-quantum
algorithm for the same problem.”

No, not necessarily!

“Yes, you do! Simply run the
simulation on the same input and
extract the original algorithm’s
output from the final state.”

Ah, but did I say that the
simulation takes only this input?

Standard structure

Algorithm simulation:

States s_0, s_1, s_2, \dots

Times t_1, t_2, \dots

State s_i represents

System state at time t_i .

What the computation

Does the original algorithm.

Case: experiment.

Computation *is*

Original algorithm

Outputs of state.

Trivially easy proof.

Simulation of quantum algorithms

“If you can efficiently simulate a quantum algorithm using a pre-quantum computer then you have an efficient pre-quantum algorithm for the same problem.”

No, not necessarily!

“Yes, you do! Simply run the simulation on the same input and extract the original algorithm’s output from the final state.”

Ah, but did I say that the simulation takes only this input?

Trapdoor

Input to
to be inp

Simulation
that mak
faster th

Typical e

- Algorithm
- Algorithm
- Simula

This is s
can try
understa

Simulation of quantum algorithms

“If you can efficiently simulate a quantum algorithm using a pre-quantum computer then you have an efficient pre-quantum algorithm for the same problem.”

No, not necessarily!

“Yes, you do! Simply run the simulation on the same input and extract the original algorithm’s output from the final state.”

Ah, but did I say that the simulation takes only this input?

Trapdoor simulation

Input to simulation
to be input to original

Simulation can use
that makes simulation
faster than original

Typical example:

- Algorithm input
- Algorithm output
- Simulation input

This is still useful:
can try many choices
understand algorithm

Simulation of quantum algorithms

“If you can efficiently simulate a quantum algorithm using a pre-quantum computer then you have an efficient pre-quantum algorithm for the same problem.”

No, not necessarily!

“Yes, you do! Simply run the simulation on the same input and extract the original algorithm’s output from the final state.”

Ah, but did I say that the simulation takes only this input?

Trapdoor simulation

Input to simulation doesn’t to be input to original algorithm

Simulation can use extra input that makes simulation much faster than original algorithm

Typical example:

- Algorithm input: $f(x)$.
- Algorithm output: x .
- Simulation input: x .

This is still useful:

can try many choices of x , understand algorithm for $f(x)$

Simulation of quantum algorithms

“If you can efficiently simulate a quantum algorithm using a pre-quantum computer then you have an efficient pre-quantum algorithm for the same problem.”

No, not necessarily!

“Yes, you do! Simply run the simulation on the same input and extract the original algorithm’s output from the final state.”

Ah, but did I say that the simulation takes only this input?

Trapdoor simulation

Input to simulation doesn’t have to be input to original algorithm.

Simulation can use extra input that makes simulation much faster than original algorithm.

Typical example:

- Algorithm input: $f(x)$.
- Algorithm output: x .
- Simulation input: x .

This is still useful:

can try many choices of x , understand algorithm for $f(x)$.

Simulation of quantum algorithms

can efficiently simulate
quantum algorithm using a
quantum computer then you
efficient pre-quantum
simulation for the same problem.”

necessarily!

you do! Simply run the
simulation on the same input and
compare the original algorithm's
output from the final state.”

did I say that the
simulation takes only this input?

Trapdoor simulation

Input to simulation doesn't have
to be input to original algorithm.

Simulation can use extra input
that makes simulation much
faster than original algorithm.

Typical example:

- Algorithm input: $f(x)$.
- Algorithm output: x .
- Simulation input: x .

This is still useful:

can try many choices of x ,
understand algorithm for $f(x)$.

For com

Often se
in tradit

Typical p
 $(x, i) \mapsto$
Formula

Simulati

Given x ,
for each
simulation

Doesn't
that wor
Proof ca

Quantum algorithms

...ntly simulate
...hm using a
...puter then you
...re-quantum
...same problem.”

...y!

...ply run the
...same input and
...l algorithm's
...nal state.”

...hat the
...nly this input?

Trapdoor simulation

Input to simulation doesn't have
to be input to original algorithm.

Simulation can use extra input
that makes simulation much
faster than original algorithm.

Typical example:

- Algorithm input: $f(x)$.
- Algorithm output: x .
- Simulation input: x .

This is still useful:

...can try many choices of x ,
...understand algorithm for $f(x)$.

For comparison:

Often see x inside
in traditional algorithm.

Typical proof has
 $(x, i) \mapsto (s_i, t_i)$.
Formula is proven

Simulation is more

Given x ,
for each i ,
simulation computes

Doesn't need uniform
that works for all x .

Proof can work “locally”

gorithms

ate

a

you

m

lem.”

ie

it and

n's

put?

Trapdoor simulation

Input to simulation doesn't have to be input to original algorithm.

Simulation can use extra input that makes simulation much faster than original algorithm.

Typical example:

- Algorithm input: $f(x)$.
- Algorithm output: x .
- Simulation input: x .

This is still useful:

can try many choices of x ,
understand algorithm for $f(x)$.

For comparison:

Often see x inside proofs in traditional algorithm analysis.

Typical proof has formula $(x, i) \mapsto (s_i, t_i)$.

Formula is proven inductively.

Simulation is more flexible.

Given x ,

for each i ,

simulation computes (s_i, t_i) .

Doesn't need unified formula that works for all x, i .

Proof can work “locally”.

Trapdoor simulation

Input to simulation doesn't have to be input to original algorithm.

Simulation can use extra input that makes simulation much faster than original algorithm.

Typical example:

- Algorithm input: $f(x)$.
- Algorithm output: x .
- Simulation input: x .

This is still useful:

can try many choices of x ,
understand algorithm for $f(x)$.

For comparison:

Often see x inside proofs
in traditional algorithm analyses.

Typical proof has formula
 $(x, i) \mapsto (s_i, t_i)$.

Formula is proven inductively.

Simulation is more flexible.

Given x ,

for each i ,

simulation computes (s_i, t_i) .

Doesn't need unified formula
that works for all x, i .

Proof can work "locally".

or simulation

simulation doesn't have
output to original algorithm.

simulation can use extra input
makes simulation much
more than original algorithm.

example:

algorithm input: $f(x)$.

algorithm output: x .

simulation input: x .

still useful:

many choices of x ,
and algorithm for $f(x)$.

For comparison:

Often see x inside proofs
in traditional algorithm analyses.

Typical proof has formula
 $(x, i) \mapsto (s_i, t_i)$.

Formula is proven inductively.

Simulation is more flexible.

Given x ,

for each i ,

simulation computes (s_i, t_i) .

Doesn't need unified formula
that works for all x, i .

Proof can work "locally".

Proof of

2014.04

Simulation

proof of

distinctness

on

n doesn't have
original algorithm.

e extra input
tion much
l algorithm.

$f(x)$.

at: x .

c: x .

ces of x ,

hm for $f(x)$.

For comparison:

Often see x inside proofs
in traditional algorithm analyses.

Typical proof has formula

$(x, i) \mapsto (s_i, t_i)$.

Formula is proven inductively.

Simulation is more flexible.

Given x ,

for each i ,

simulation computes (s_i, t_i) .

Doesn't need unified formula

that works for all x, i .

Proof can work "locally".

Proof of concept

2014.04 Chou \rightarrow A

Simulation shows

proof of 2003 Am

distinctness algorit

For comparison:

Often see x inside proofs
in traditional algorithm analyses.

Typical proof has formula

$$(x, i) \mapsto (s_i, t_i).$$

Formula is proven inductively.

Simulation is more flexible.

Given x ,

for each i ,

simulation computes (s_i, t_i) .

Doesn't need unified formula
that works for all x, i .

Proof can work "locally".

Proof of concept

2014.04 Chou \rightarrow Ambainis:
Simulation shows error in
proof of 2003 Ambainis
distinctness algorithm.

For comparison:

Often see x inside proofs
in traditional algorithm analyses.

Typical proof has formula

$$(x, i) \mapsto (s_i, t_i).$$

Formula is proven inductively.

Simulation is more flexible.

Given x ,

for each i ,

simulation computes (s_i, t_i) .

Doesn't need unified formula
that works for all x, i .

Proof can work "locally".

Proof of concept

2014.04 Chou \rightarrow Ambainis:
Simulation shows error in
proof of 2003 Ambainis
distinctness algorithm.

For comparison:

Often see x inside proofs
in traditional algorithm analyses.

Typical proof has formula

$(x, i) \mapsto (s_i, t_i)$.

Formula is proven inductively.

Simulation is more flexible.

Given x ,

for each i ,

simulation computes (s_i, t_i) .

Doesn't need unified formula
that works for all x, i .

Proof can work "locally".

Proof of concept

2014.04 Chou \rightarrow Ambainis:

Simulation shows error in
proof of 2003 Ambainis
distinctness algorithm.

Ambainis: Yes, thanks, will fix.

For comparison:

Often see x inside proofs
in traditional algorithm analyses.

Typical proof has formula

$$(x, i) \mapsto (s_i, t_i).$$

Formula is proven inductively.

Simulation is more flexible.

Given x ,

for each i ,

simulation computes (s_i, t_i) .

Doesn't need unified formula
that works for all x, i .

Proof can work "locally".

Proof of concept

2014.04 Chou \rightarrow Ambainis:

Simulation shows error in
proof of 2003 Ambainis
distinctness algorithm.

Ambainis: Yes, thanks, will fix.

2014.04 Chou \rightarrow Childs:

Simulation shows that 2003
Childs–Eisenberg distinctness
algorithm is non-functional;
need to take half angle.

For comparison:

Often see x inside proofs
in traditional algorithm analyses.

Typical proof has formula

$$(x, i) \mapsto (s_i, t_i).$$

Formula is proven inductively.

Simulation is more flexible.

Given x ,

for each i ,

simulation computes (s_i, t_i) .

Doesn't need unified formula
that works for all x, i .

Proof can work "locally".

Proof of concept

2014.04 Chou \rightarrow Ambainis:

Simulation shows error in
proof of 2003 Ambainis
distinctness algorithm.

Ambainis: Yes, thanks, will fix.

2014.04 Chou \rightarrow Childs:

Simulation shows that 2003
Childs–Eisenberg distinctness
algorithm is non-functional;
need to take half angle.

Childs: Yes. Typo, already
fixed in 2005 journal version.