

The post-quantum Internet

Daniel J. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

Includes joint work with:

Tanja Lange

Technische Universiteit Eindhoven

Risk management

“Combining congruences” :
state-of-the-art pre-quantum
attack against original DH,
RSA, and some lattice systems.

Long history, including
many major improvements:
1975, CFRAC;
1977, linear sieve (LS);
1982, quadratic sieve (QS);
1990, number-field sieve (NFS);
1994, function-field sieve (FFS);
2006, medium-prime FFS/NFS;
2013, $x^q - x$ FFS.

post-quantum Internet

D. Bernstein

University of Illinois at Chicago &
Technische Universiteit Eindhoven

joint work with:

W. Ren

Technische Universiteit Eindhoven

1

Risk management

“Combining congruences” :
state-of-the-art pre-quantum
attack against original DH,
RSA, and some lattice systems.

Long history, including
many major improvements:

1975, CFRAC;

1977, linear sieve (LS);

1982, quadratic sieve (QS);

1990, number-field sieve (NFS);

1994, function-field sieve (FFS);

2006, medium-prime FFS/NFS;

2013, $x^q - x$ FFS.

2

Also made

>100 sc

Costs of

breaking

$\approx 2^{120}$, \approx

$\approx 2^{110}$, \approx

$\approx 2^{100}$, \approx

$\approx 2^{80}$, \approx

(FFS is

1

Risk management

“Combining congruences” :
state-of-the-art pre-quantum
attack against original DH,
RSA, and some lattice systems.

Long history, including
many major improvements:
1975, CFRAC;
1977, linear sieve (LS);
1982, quadratic sieve (QS);
1990, number-field sieve (NFS);
1994, function-field sieve (FFS);
2006, medium-prime FFS/NFS;
2013, $x^q - x$ FFS.

2

Also many smaller
>100 scientific pa
Costs of these algo
breaking RSA-102
 $\approx 2^{120}$, $\approx 2^{170}$, CF
 $\approx 2^{110}$, $\approx 2^{160}$, LS;
 $\approx 2^{100}$, $\approx 2^{150}$, QS
 $\approx 2^{80}$, $\approx 2^{112}$, NF
(FFS is not releva

Risk management

“Combining congruences” :
state-of-the-art pre-quantum
attack against original DH,
RSA, and some lattice systems.

Long history, including
many major improvements:
1975, CFRAC;
1977, linear sieve (LS);
1982, quadratic sieve (QS);
1990, number-field sieve (NFS);
1994, function-field sieve (FFS);
2006, medium-prime FFS/NFS;
2013, $x^q - x$ FFS.

ago &
hoven

hoven

Also many smaller improvements
>100 scientific papers.

Costs of these algorithms for
breaking RSA-1024, RSA-2048
 $\approx 2^{120}$, $\approx 2^{170}$, CFRAC;
 $\approx 2^{110}$, $\approx 2^{160}$, LS;
 $\approx 2^{100}$, $\approx 2^{150}$, QS;
 $\approx 2^{80}$, $\approx 2^{112}$, NFS.

(FFS is not relevant to RSA)

Risk management

“Combining congruences” :
state-of-the-art pre-quantum
attack against original DH,
RSA, and some lattice systems.

Long history, including
many major improvements:
1975, CFRAC;
1977, linear sieve (LS);
1982, quadratic sieve (QS);
1990, number-field sieve (NFS);
1994, function-field sieve (FFS);
2006, medium-prime FFS/NFS;
2013, $x^q - x$ FFS.

Also many smaller improvements:
>100 scientific papers.

Costs of these algorithms for
breaking RSA-1024, RSA-2048:

$\approx 2^{120}$, $\approx 2^{170}$, CFRAC;

$\approx 2^{110}$, $\approx 2^{160}$, LS;

$\approx 2^{100}$, $\approx 2^{150}$, QS;

$\approx 2^{80}$, $\approx 2^{112}$, NFS.

(FFS is not relevant to RSA.)

Risk management

“Combining congruences” :
state-of-the-art pre-quantum
attack against original DH,
RSA, and some lattice systems.

Long history, including
many major improvements:
1975, CFRAC;
1977, linear sieve (LS);
1982, quadratic sieve (QS);
1990, number-field sieve (NFS);
1994, function-field sieve (FFS);
2006, medium-prime FFS/NFS;
2013, $x^q - x$ FFS.

Also many smaller improvements:
>100 scientific papers.

Costs of these algorithms for
breaking RSA-1024, RSA-2048:

$\approx 2^{120}$, $\approx 2^{170}$, CFRAC;

$\approx 2^{110}$, $\approx 2^{160}$, LS;

$\approx 2^{100}$, $\approx 2^{150}$, QS;

$\approx 2^{80}$, $\approx 2^{112}$, NFS.

(FFS is not relevant to RSA.)

How much risk is there
of future breakthroughs?

Risk management

“Combining congruences” :
state-of-the-art pre-quantum
attack against original DH,
RSA, and some lattice systems.

Long history, including
many major improvements:
1975, CFRAC;
1977, linear sieve (LS);
1982, quadratic sieve (QS);
1990, number-field sieve (NFS);
1994, function-field sieve (FFS);
2006, medium-prime FFS/NFS;
2013, $x^q - x$ FFS.

Also many smaller improvements:
>100 scientific papers.

Costs of these algorithms for
breaking RSA-1024, RSA-2048:

$\approx 2^{120}$, $\approx 2^{170}$, CFRAC;

$\approx 2^{110}$, $\approx 2^{160}$, LS;

$\approx 2^{100}$, $\approx 2^{150}$, QS;

$\approx 2^{80}$, $\approx 2^{112}$, NFS.

(FFS is not relevant to RSA.)

How much risk is there
of future breakthroughs?

How much risk is there
of secret breakthroughs?

Management

"finding congruences":

state-of-the-art pre-quantum
attacks against original DH,
and some lattice systems.

History, including

major improvements:

CFRAC;

Linear sieve (LS);

Quadratic sieve (QS);

Number-field sieve (NFS);

Function-field sieve (FFS);

Medium-prime FFS/NFS;

l - x FFS.

2

Also many smaller improvements:
>100 scientific papers.

Costs of these algorithms for
breaking RSA-1024, RSA-2048:

$\approx 2^{120}$, $\approx 2^{170}$, CFRAC;

$\approx 2^{110}$, $\approx 2^{160}$, LS;

$\approx 2^{100}$, $\approx 2^{150}$, QS;

$\approx 2^{80}$, $\approx 2^{112}$, NFS.

(FFS is not relevant to RSA.)

How much risk is there
of future breakthroughs?

How much risk is there
of secret breakthroughs?

3

If we pursue
exploring
will we find
At least

2

quences":
 e-quantum
 ginal DH,
 ttice systems.
 ding
 vements:
 (LS);
 eve (QS);
 d sieve (NFS);
 d sieve (FFS);
 me FFS/NFS;

Also many smaller improvements:
 >100 scientific papers.

Costs of these algorithms for
 breaking RSA-1024, RSA-2048:

$\approx 2^{120}$, $\approx 2^{170}$, CFRAC;

$\approx 2^{110}$, $\approx 2^{160}$, LS;

$\approx 2^{100}$, $\approx 2^{150}$, QS;

$\approx 2^{80}$, $\approx 2^{112}$, NFS.

(FFS is not relevant to RSA.)

How much risk is there
 of future breakthroughs?

How much risk is there
 of secret breakthroughs?

3

If we put enough e
 exploring Attack M
 will we find the hig
 At least within ϵ ?

2

Also many smaller improvements:
>100 scientific papers.

Costs of these algorithms for
breaking RSA-1024, RSA-2048:

$\approx 2^{120}$, $\approx 2^{170}$, CFRAC;

$\approx 2^{110}$, $\approx 2^{160}$, LS;

$\approx 2^{100}$, $\approx 2^{150}$, QS;

$\approx 2^{80}$, $\approx 2^{112}$, NFS.

(FFS is not relevant to RSA.)

How much risk is there
of future breakthroughs?

How much risk is there
of secret breakthroughs?

3

If we put enough effort into
exploring Attack Mountain,
will we find the highest peak?
At least within ϵ ?

Also many smaller improvements:
 >100 scientific papers.

Costs of these algorithms for
 breaking RSA-1024, RSA-2048:

$\approx 2^{120}$, $\approx 2^{170}$, CFRAC;

$\approx 2^{110}$, $\approx 2^{160}$, LS;

$\approx 2^{100}$, $\approx 2^{150}$, QS;

$\approx 2^{80}$, $\approx 2^{112}$, NFS.

(FFS is not relevant to RSA.)

How much risk is there
 of future breakthroughs?

How much risk is there
 of secret breakthroughs?

If we put enough effort into
 exploring Attack Mountain,
 will we find the highest peak?
 At least within ϵ ?

Also many smaller improvements:
 >100 scientific papers.

Costs of these algorithms for
 breaking RSA-1024, RSA-2048:

$\approx 2^{120}$, $\approx 2^{170}$, CFRAC;

$\approx 2^{110}$, $\approx 2^{160}$, LS;

$\approx 2^{100}$, $\approx 2^{150}$, QS;

$\approx 2^{80}$, $\approx 2^{112}$, NFS.

(FFS is not relevant to RSA.)

How much risk is there
 of future breakthroughs?

How much risk is there
 of secret breakthroughs?

If we put enough effort into
 exploring Attack Mountain,
 will we find the highest peak?
 At least within ϵ ?

Combining-Congruences Mountain
 is a huge, foggy, high-dimensional
 mountain with many paths up.
 Scary: easy to imagine that
 we're not at the top yet.

Also many smaller improvements:
 >100 scientific papers.

Costs of these algorithms for
 breaking RSA-1024, RSA-2048:

$\approx 2^{120}$, $\approx 2^{170}$, CFRAC;

$\approx 2^{110}$, $\approx 2^{160}$, LS;

$\approx 2^{100}$, $\approx 2^{150}$, QS;

$\approx 2^{80}$, $\approx 2^{112}$, NFS.

(FFS is not relevant to RSA.)

How much risk is there
 of future breakthroughs?

How much risk is there
 of secret breakthroughs?

If we put enough effort into
 exploring Attack Mountain,
 will we find the highest peak?
 At least within ϵ ?

Combining-Congruences Mountain
 is a huge, foggy, high-dimensional
 mountain with many paths up.
 Scary: easy to imagine that
 we're not at the top yet.

18-year bet announced in 2014:
 Joux wins if RSA-2048 is broken
 first by pre-quantum algorithms;
 I win if RSA-2048 is broken
 first by quantum algorithms.

any smaller improvements:
scientific papers.

these algorithms for
RSA-1024, RSA-2048:

$\approx 2^{170}$, CFRAC;

$\approx 2^{160}$, LS;

$\approx 2^{150}$, QS;

$\approx 2^{112}$, NFS.

(not relevant to RSA.)

What risk is there
from breakthroughs?

What risk is there
from breakthroughs?

3

If we put enough effort into
exploring Attack Mountain,
will we find the highest peak?
At least within ϵ ?

Combining-Congruences Mountain
is a huge, foggy, high-dimensional
mountain with many paths up.
Scary: easy to imagine that
we're not at the top yet.

18-year bet announced in 2014:
Joux wins if RSA-2048 is broken
first by pre-quantum algorithms;
I win if RSA-2048 is broken
first by quantum algorithms.

4

Conservative
prefer more
less huge
more the

improvements:
pers.

gorithms for
4, RSA-2048:

RAC;

;
S.
(nt to RSA.)

there
oughs?

there
oughs?

3

If we put enough effort into
exploring Attack Mountain,
will we find the highest peak?
At least within ϵ ?

Combining-Congruences Mountain
is a huge, foggy, high-dimensional
mountain with many paths up.
Scary: easy to imagine that
we're not at the top yet.

18-year bet announced in 2014:
Joux wins if RSA-2048 is broken
first by pre-quantum algorithms;
I win if RSA-2048 is broken
first by quantum algorithms.

4

Conservative crypt
prefer mountains t
less huge, less fog
more thoroughly e

3

ments:

If we put enough effort into exploring Attack Mountain, will we find the highest peak? At least within ϵ ?

r

048:

Combining-Congruences Mountain is a huge, foggy, high-dimensional mountain with many paths up. Scary: easy to imagine that we're not at the top yet.

.)

18-year bet announced in 2014: Joux wins if RSA-2048 is broken first by pre-quantum algorithms; I win if RSA-2048 is broken first by quantum algorithms.

4

Conservative cryptographers prefer mountains that seem less huge, less foggy, more thoroughly explored.

If we put enough effort into exploring Attack Mountain, will we find the highest peak?
At least within ϵ ?

Combining-Congruences Mountain is a huge, foggy, high-dimensional mountain with many paths up.
Scary: easy to imagine that we're not at the top yet.

18-year bet announced in 2014:
Joux wins if RSA-2048 is broken first by pre-quantum algorithms;
I win if RSA-2048 is broken first by quantum algorithms.

Conservative cryptographers prefer mountains that seem less huge, less foggy, more thoroughly explored.

If we put enough effort into exploring Attack Mountain, will we find the highest peak?
At least within ϵ ?

Combining-Congruences Mountain is a huge, foggy, high-dimensional mountain with many paths up.
Scary: easy to imagine that we're not at the top yet.

18-year bet announced in 2014:
Joux wins if RSA-2048 is broken first by pre-quantum algorithms;
I win if RSA-2048 is broken first by quantum algorithms.

Conservative cryptographers prefer mountains that seem less huge, less foggy, more thoroughly explored.

1986 Miller "Use of elliptic curves in cryptography":
"It is extremely unlikely that an 'index calculus' attack [combining-congruences attack] on the elliptic curve method will ever be able to work."

If we put enough effort into exploring Attack Mountain, will we find the highest peak?

At least within ϵ ?

Combining-Congruences Mountain is a huge, foggy, high-dimensional mountain with many paths up.

Scary: easy to imagine that we're not at the top yet.

18-year bet announced in 2014:
Joux wins if RSA-2048 is broken first by pre-quantum algorithms;
I win if RSA-2048 is broken first by quantum algorithms.

Conservative cryptographers prefer mountains that seem less huge, less foggy, more thoroughly explored.

1986 Miller “Use of elliptic curves in cryptography”:
“It is extremely unlikely that an ‘index calculus’ attack [combining-congruences attack] on the elliptic curve method will ever be able to work.”

This is the core argument for ECC. Exceptions: rare curves with special structure—e.g., pairings.

Put enough effort into
 climbing Attack Mountain,
 you can find the highest peak?
 within ϵ ?

Combining-Congruences Mountain
 is, foggy, high-dimensional
 and has many paths up.
 It's easy to imagine that
 you can't get to the top yet.

A bet announced in 2014:
 If RSA-2048 is broken
 before quantum algorithms;
 If RSA-2048 is broken
 before quantum algorithms.

Conservative cryptographers
 prefer mountains that seem
 less huge, less foggy,
 more thoroughly explored.

1986 Miller “Use of
 elliptic curves in cryptography”:
 “It is extremely unlikely
 that an ‘index calculus’ attack
 [combining-congruences attack]
 on the elliptic curve method
 will ever be able to work.”

This is the core argument for
 ECC. Exceptions: rare curves with
 special structure—e.g., pairings.

2015 La
 bet your

The setting

It's 2050. Q

Evil Party A
 practically a
 vaccinations
 the past 70
 by law, but
 2020. Your
 doctor's pub
 public-key a

Organs are
 if they can
 presented w
 Statement.

(This is mea
 trust. Let's

effort into
Mountain,
highest peak?

ences Mountain
high-dimensional
ny paths up.
agine that
op yet.

nced in 2014:
2048 is broken
m algorithms;
is broken
algorithms.

Conservative cryptographers
prefer mountains that seem
less huge, less foggy,
more thoroughly explored.

1986 Miller “Use of
elliptic curves in cryptography”:
“It is extremely unlikely
that an ‘index calculus’ attack
[combining-congruences attack]
on the elliptic curve method
will ever be able to work.”

This is the core argument for
ECC. Exceptions: rare curves with
special structure—e.g., pairings.

2015 Lange: “Wo
bet your kidneys o

The setting

It's 2050. Quantum computers w

Evil Party A now runs the country
practically all 21st-century Internet
vaccinations are bad and jails any
the past 70 years. Doctor-patient
by law, but your health record fro
2020. Your health record is protec
doctor's public key, using our reco
public-key and authenticated sym

Organs are a scarce resource. Hos
if they can identify the donor (DM
presented with the donor's digital
Statement. They use our 2015 re

(This is meant to scare you, so th
trust. Let's make sure that this d

Conservative cryptographers prefer mountains that seem less huge, less foggy, more thoroughly explored.

1986 Miller “Use of elliptic curves in cryptography” :
 “It is extremely unlikely that an ‘index calculus’ attack [combining-congruences attack] on the elliptic curve method will ever be able to work.”

This is the core argument for ECC. Exceptions: rare curves with special structure—e.g., pairings.

2015 Lange: “Would you bet your kidneys on that?”

The setting

It's 2050. Quantum computers were built years ago.

Evil Party A now runs the country and has access to re practically all 21st-century Internet traffic. Evil Party A vaccinations are bad and jails anybody who was vaccina the past 70 years. Doctor-patient confidentiality is still by law, but your health record from birth has been onli 2020. Your health record is protected only by encryptio doctor's public key, using our recommendation from 20 public-key and authenticated symmetric encryption.

Organs are a scarce resource. Hospitals pay high prices if they can identify the donor (DNA tests are cheap) an presented with the donor's digitally signed Donor Volun Statement. They use our 2015 recommended signature

(This is meant to scare you, so that you recommend on trust. Let's make sure that this dystopia will not happe

Conservative cryptographers prefer mountains that seem less huge, less foggy, more thoroughly explored.

1986 Miller “Use of elliptic curves in cryptography”: “It is extremely unlikely that an ‘index calculus’ attack [combining-congruences attack] on the elliptic curve method will ever be able to work.”

This is the core argument for ECC. Exceptions: rare curves with special structure—e.g., pairings.

2015 Lange: “Would you bet your kidneys on that?”

The setting

It's 2050. Quantum computers were built years ago.

Evil Party A now runs the country and has access to records of practically all 21st-century Internet traffic. Evil Party A thinks vaccinations are bad and jails anybody who was vaccinated during the past 70 years. Doctor-patient confidentiality is still protected by law, but your health record from birth has been online since 2020. Your health record is protected only by encryption to your doctor's public key, using our recommendation from 2015 of public-key and authenticated symmetric encryption.

Organs are a scarce resource. Hospitals pay high prices for organs if they can identify the donor (DNA tests are cheap) and are presented with the donor's digitally signed Donor Volunteer Statement. They use our 2015 recommended signature system.

(This is meant to scare you, so that you recommend only what you trust. Let's make sure that this dystopia will not happen.)

ative cryptographers
mountains that seem
e, less foggy,
oroughly explored.

ller “Use of
curves in cryptography”:
tremely unlikely
‘index calculus’ attack
ing-congruences attack]
llyptic curve method
be able to work.”

he core argument for
ceptions: rare curves with
structure—e.g., pairings.

2015 Lange: “Would you
bet your kidneys on that?”

The setting

It's 2050. Quantum computers were built years ago.

Evil Party A now runs the country and has access to records of practically all 21st-century Internet traffic. Evil Party A thinks vaccinations are bad and jails anybody who was vaccinated during the past 70 years. Doctor-patient confidentiality is still protected by law, but your health record from birth has been online since 2020. Your health record is protected only by encryption to your doctor's public key, using our recommendation from 2015 of public-key and authenticated symmetric encryption.

Organs are a scarce resource. Hospitals pay high prices for organs if they can identify the donor (DNA tests are cheap) and are presented with the donor's digitally signed Donor Volunteer Statement. They use our 2015 recommended signature system.

(This is meant to scare you, so that you recommend only what you trust. Let's make sure that this dystopia will not happen.)

Risk of f
big unive
noticeab
terrifying

cryptographers
that seem
gy,
explored.

of
"cryptography":
unlikely
'attack
[quantum attacks]
method
to work."

argument for
rare curves with
-e.g., pairings.

2015 Lange: "Would you bet your kidneys on that?"

The setting

It's 2050. Quantum computers were built years ago.

Evil Party A now runs the country and has access to records of practically all 21st-century Internet traffic. Evil Party A thinks vaccinations are bad and jails anybody who was vaccinated during the past 70 years. Doctor-patient confidentiality is still protected by law, but your health record from birth has been online since 2020. Your health record is protected only by encryption to your doctor's public key, using our recommendation from 2015 of public-key and authenticated symmetric encryption.

Organs are a scarce resource. Hospitals pay high prices for organs if they can identify the donor (DNA tests are cheap) and are presented with the donor's digitally signed Donor Volunteer Statement. They use our 2015 recommended signature system.

(This is meant to scare you, so that you recommend only what you trust. Let's make sure that this dystopia will not happen.)

Risk of future attacks
big universal quantum
noticeable probability
terrifying impact.

2015 Lange: “Would you bet your kidneys on that?”

The setting

It's 2050. Quantum computers were built years ago.

Evil Party A now runs the country and has access to records of practically all 21st-century Internet traffic. Evil Party A thinks vaccinations are bad and jails anybody who was vaccinated during the past 70 years. Doctor-patient confidentiality is still protected by law, but your health record from birth has been online since 2020. Your health record is protected only by encryption to your doctor's public key, using our recommendation from 2015 of public-key and authenticated symmetric encryption.

Organs are a scarce resource. Hospitals pay high prices for organs if they can identify the donor (DNA tests are cheap) and are presented with the donor's digitally signed Donor Volunteer Statement. They use our 2015 recommended signature system.

(This is meant to scare you, so that you recommend only what you trust. Let's make sure that this dystopia will not happen.)

Risk of future attacker having
big universal quantum comp
noticeable probability;
terrifying impact.

2015 Lange: “Would you bet your kidneys on that?”

The setting

It's 2050. Quantum computers were built years ago.

Evil Party A now runs the country and has access to records of practically all 21st-century Internet traffic. Evil Party A thinks vaccinations are bad and jails anybody who was vaccinated during the past 70 years. Doctor-patient confidentiality is still protected by law, but your health record from birth has been online since 2020. Your health record is protected only by encryption to your doctor's public key, using our recommendation from 2015 of public-key and authenticated symmetric encryption.

Organs are a scarce resource. Hospitals pay high prices for organs if they can identify the donor (DNA tests are cheap) and are presented with the donor's digitally signed Donor Volunteer Statement. They use our 2015 recommended signature system.

(This is meant to scare you, so that you recommend only what you trust. Let's make sure that this dystopia will not happen.)

Risk of future attacker having big universal quantum computer: noticeable probability; terrifying impact.

2015 Lange: “Would you bet your kidneys on that?”

The setting

It's 2050. Quantum computers were built years ago.

Evil Party A now runs the country and has access to records of practically all 21st-century Internet traffic. Evil Party A thinks vaccinations are bad and jails anybody who was vaccinated during the past 70 years. Doctor-patient confidentiality is still protected by law, but your health record from birth has been online since 2020. Your health record is protected only by encryption to your doctor's public key, using our recommendation from 2015 of public-key and authenticated symmetric encryption.

Organs are a scarce resource. Hospitals pay high prices for organs if they can identify the donor (DNA tests are cheap) and are presented with the donor's digitally signed Donor Volunteer Statement. They use our 2015 recommended signature system.

(This is meant to scare you, so that you recommend only what you trust. Let's make sure that this dystopia will not happen.)

Risk of future attacker having big universal quantum computer: noticeable probability; terrifying impact.

Fortunately, we already know some confidence-inspiring post-quantum systems, including

- hash-based signatures;
- McEliece public-key encryption;
- AES-256 etc.

<https://pqcrypto.eu.org/docs/initial-recommendations.pdf>

Change: “Would you
put your kidneys on that?”

Quantum computers were built years ago.

Evil Party A now runs the country and has access to records of all 21st-century Internet traffic. Evil Party A thinks doctors are bad and jails anybody who was vaccinated during the last 10 years. Doctor-patient confidentiality is still protected. Your health record from birth has been online since 2000. Your health record is protected only by encryption to your public key, using our recommendation from 2015 of NIST SP 800-56A and authenticated symmetric encryption.

Organ donation is a scarce resource. Hospitals pay high prices for organs to identify the donor (DNA tests are cheap) and are often confused with the donor's digitally signed Donor Volunteer Certificate. They use our 2015 recommended signature system.

(We want to scare you, so that you recommend only what you think is best. We make sure that this dystopia will not happen.)

6

Risk of future attacker having a big universal quantum computer: noticeable probability; terrifying impact.

Fortunately, we already know some confidence-inspiring post-quantum systems, including

- hash-based signatures;
- McEliece public-key encryption;
- AES-256 etc.

<https://pqcrypto.eu.org/docs/initial-recommendations.pdf>

7

Applicat

Your com
new vers

Your com
signature
from the

Critical
Otherwis
insert m

e.g. Ope
signed u
ECC sig

ould you
on that?"

ere built years ago.

y and has access to records of
et traffic. Evil Party A thinks
body who was vaccinated during
confidentiality is still protected
m birth has been online since
cted only by encryption to your
ommendation from 2015 of
metric encryption.

spitals pay high prices for organs
(NA tests are cheap) and are
ly signed Donor Volunteer
commended signature system.

at you recommend only what you
(dystopia will not happen.)

6

Risk of future attacker having
big universal quantum computer:
noticeable probability;
terrifying impact.

Fortunately, we already know
some confidence-inspiring
post-quantum systems, including

- hash-based signatures;
- McEliece public-key encryption;
- AES-256 etc.

[https://pqcrypto.eu.org/docs/
initial-recommendations.pdf](https://pqcrypto.eu.org/docs/initial-recommendations.pdf)

7

Application: software

Your computer does
new version of its

Your computer checks
signature on the disk
from the OS manu

Critical use of crypt
Otherwise criminal
insert malware into

e.g. OpenBSD updat
signed using state-
ECC signature sys

Risk of future attacker having big universal quantum computer: noticeable probability; terrifying impact.

Fortunately, we already know some confidence-inspiring post-quantum systems, including

- hash-based signatures;
- McEliece public-key encryption;
- AES-256 etc.

<https://pqcrypto.eu.org/docs/initial-recommendations.pdf>

Application: software updates

Your computer downloads new version of its OS.

Your computer checks signature on the download from the OS manufacturer.

Critical use of crypto!

Otherwise criminals could insert malware into the OS.

e.g. OpenBSD updates are signed using state-of-the-art ECC signature system: Ed25519

Risk of future attacker having big universal quantum computer: noticeable probability; terrifying impact.

Fortunately, we already know some confidence-inspiring post-quantum systems, including

- hash-based signatures;
- McEliece public-key encryption;
- AES-256 etc.

<https://pqcrypto.eu.org/docs/initial-recommendations.pdf>

Application: software updates

Your computer downloads new version of its OS.

Your computer checks signature on the download from the OS manufacturer.

Critical use of crypto!
Otherwise criminals could insert malware into the OS.

e.g. OpenBSD updates are signed using state-of-the-art ECC signature system: Ed25519.

future attacker having
universal quantum computer:
non-negligible probability;
significant impact.

Concretely, we already know
confidence-inspiring
quantum systems, including
signature-based signatures;
post-quantum public-key encryption;
SHA-3 etc.

[https://pqcrypto.eu.org/docs/
quantum-recommendations.pdf](https://pqcrypto.eu.org/docs/quantum-recommendations.pdf)

7

Application: software updates

Your computer downloads
new version of its OS.

Your computer checks
signature on the download
from the OS manufacturer.

Critical use of crypto!

Otherwise criminals could
insert malware into the OS.

e.g. OpenBSD updates are
signed using state-of-the-art
ECC signature system: Ed25519.

8

Pre-quantum
needs to
post-quantum

7

hacker having
quantum computer:
security;

already know
inspiring
systems, including
signatures;
key encryption;

<https://www.openbsd.org/docs/updates.pdf>

Application: software updates

Your computer downloads
new version of its OS.

Your computer checks
signature on the download
from the OS manufacturer.

Critical use of crypto!

Otherwise criminals could
insert malware into the OS.

e.g. OpenBSD updates are
signed using state-of-the-art
ECC signature system: Ed25519.

8

Pre-quantum signature
needs to be replaced by
post-quantum signature

7

Application: software updates

Your computer downloads new version of its OS.

Your computer checks signature on the download from the OS manufacturer.

Critical use of crypto!

Otherwise criminals could insert malware into the OS.

e.g. OpenBSD updates are signed using state-of-the-art ECC signature system: Ed25519.

8

Pre-quantum signature systems need to be replaced with post-quantum signature systems.

Application: software updates

Your computer downloads new version of its OS.

Your computer checks signature on the download from the OS manufacturer.

Critical use of crypto!

Otherwise criminals could insert malware into the OS.

e.g. OpenBSD updates are signed using state-of-the-art ECC signature system: Ed25519.

Pre-quantum signature system P needs to be replaced with post-quantum signature system Q .

Application: software updates

Your computer downloads new version of its OS.

Your computer checks signature on the download from the OS manufacturer.

Critical use of crypto!

Otherwise criminals could insert malware into the OS.

e.g. OpenBSD updates are signed using state-of-the-art ECC signature system: Ed25519.

Pre-quantum signature system P needs to be replaced with post-quantum signature system Q .

Make auditors happier:
Replace P with $P + Q$.

$P + Q$ public key concatenates

P public key, Q public key.

$P + Q$ signature concatenates

P signature, Q signature.

Application: software updates

Your computer downloads new version of its OS.

Your computer checks signature on the download from the OS manufacturer.

Critical use of crypto!

Otherwise criminals could insert malware into the OS.

e.g. OpenBSD updates are signed using state-of-the-art ECC signature system: Ed25519.

Pre-quantum signature system P needs to be replaced with post-quantum signature system Q .

Make auditors happier:

Replace P with $P + Q$.

$P + Q$ public key concatenates

P public key, Q public key.

$P + Q$ signature concatenates

P signature, Q signature.

Want a tiny public key?

Replace public key with hash.

Include missing information

(\leq entire key) inside signature.

ion: software updates

computer downloads
sion of its OS.

computer checks
e on the download
e OS manufacturer.

use of crypto!
se criminals could
aware into the OS.

enBSD updates are
sing state-of-the-art
nature system: Ed25519.

8

Pre-quantum signature system P
needs to be replaced with
post-quantum signature system Q .

Make auditors happier:

Replace P with $P + Q$.

$P + Q$ public key concatenates

P public key, Q public key.

$P + Q$ signature concatenates

P signature, Q signature.

Want a tiny public key?

Replace public key with hash.

Include missing information

(\leq entire key) inside signature.

9

e.g. Ed25519

SPHINCOFF

≈ 50 million

≈ 1 million

Negligible

verify cost

8

Pre-quantum signature system P
needs to be replaced with
post-quantum signature system Q .

Make auditors happier:

Replace P with $P + Q$.

$P + Q$ public key concatenates

P public key, Q public key.

$P + Q$ signature concatenates

P signature, Q signature.

Want a tiny public key?

Replace public key with hash.

Include missing information

(\leq entire key) inside signature.

9

e.g. Ed25519+SPHINCS

SPHINCS-256 sign

≈ 50 million cycles

≈ 1 million cycles to

Negligible cost to

verify compared to

Pre-quantum signature system P
needs to be replaced with
post-quantum signature system Q .

Make auditors happier:

Replace P with $P + Q$.

$P + Q$ public key concatenates

P public key, Q public key.

$P + Q$ signature concatenates

P signature, Q signature.

Want a tiny public key?

Replace public key with hash.

Include missing information

(\leq entire key) inside signature.

e.g. Ed25519+SPHINCS-256

SPHINCS-256 signature is 4

≈ 50 million cycles to generate

≈ 1 million cycles to verify.

Negligible cost to sign, trans

verify compared to OS upda

Pre-quantum signature system P
needs to be replaced with
post-quantum signature system Q .

Make auditors happier:

Replace P with $P + Q$.

$P + Q$ public key concatenates

P public key, Q public key.

$P + Q$ signature concatenates

P signature, Q signature.

Want a tiny public key?

Replace public key with hash.

Include missing information

(\leq entire key) inside signature.

e.g. Ed25519+SPHINCS-256.

SPHINCS-256 signature is 41KB;

\approx 50 million cycles to generate;

\approx 1 million cycles to verify.

Negligible cost to sign, transmit,
verify compared to OS update.

Pre-quantum signature system P
needs to be replaced with
post-quantum signature system Q .

Make auditors happier:

Replace P with $P + Q$.

$P + Q$ public key concatenates

P public key, Q public key.

$P + Q$ signature concatenates

P signature, Q signature.

Want a tiny public key?

Replace public key with hash.

Include missing information

(\leq entire key) inside signature.

e.g. Ed25519+SPHINCS-256.

SPHINCS-256 signature is 41KB;

\approx 50 million cycles to generate;

\approx 1 million cycles to verify.

Negligible cost to sign, transmit,
verify compared to OS update.

+Ed25519: unnoticeable cost.

Some extra system complexity,

but the system includes

Ed25519 code anyway.

Pre-quantum signature system P
needs to be replaced with
post-quantum signature system Q .

Make auditors happier:

Replace P with $P + Q$.

$P + Q$ public key concatenates

P public key, Q public key.

$P + Q$ signature concatenates

P signature, Q signature.

Want a tiny public key?

Replace public key with hash.

Include missing information

(\leq entire key) inside signature.

e.g. Ed25519+SPHINCS-256.

SPHINCS-256 signature is 41KB;

\approx 50 million cycles to generate;

\approx 1 million cycles to verify.

Negligible cost to sign, transmit,
verify compared to OS update.

+Ed25519: unnoticeable cost.

Some extra system complexity,

but the system includes

Ed25519 code anyway.

Auditor sees very easily

that Ed25519+SPHINCS-256

security \geq Ed25519 security.

quantum signature system P
 to be replaced with
 quantum signature system Q .
 Auditors happier:
 P with $P + Q$.
 public key concatenates
 P public key, Q public key.
 signature concatenates
 P signature, Q signature.
 tiny public key?
 public key with hash.
 missing information
 (the key) inside signature.

e.g. Ed25519+SPHINCS-256.
 SPHINCS-256 signature is 41KB;
 ≈ 50 million cycles to generate;
 ≈ 1 million cycles to verify.
 Negligible cost to sign, transmit,
 verify compared to OS update.
 +Ed25519: unnoticeable cost.
 Some extra system complexity,
 but the system includes
 Ed25519 code anyway.
 Auditor sees very easily
 that Ed25519+SPHINCS-256
 security \geq Ed25519 security.

Does de
 mean th
 On the
 Pre-quant
 Hash-ba
 even mo
 than EC
 But und
 takes ex

signature system P
 ed with
 signature system Q .

opier:

+ Q .

concatenates
 public key.

concatenates
 signature.

c key?

y with hash.

formation

de signature.

e.g. Ed25519+SPHINCS-256.

SPHINCS-256 signature is 41KB;
 ≈ 50 million cycles to generate;
 ≈ 1 million cycles to verify.

Negligible cost to sign, transmit,
 verify compared to OS update.

+Ed25519: unnoticeable cost.

Some extra system complexity,
 but the system includes
 Ed25519 code anyway.

Auditor sees very easily
 that Ed25519+SPHINCS-256
 security \geq Ed25519 security.

Does deployment
 mean that we don't
 On the contrary!

Pre-quantum situa
 Hash-based signat
 even more confide
 than ECC signatur
 But understanding
 takes extra work f

em P

em Q .

tes

es

n.

ure.

e.g. Ed25519+SPHINCS-256.

SPHINCS-256 signature is 41KB;

≈ 50 million cycles to generate;

≈ 1 million cycles to verify.

Negligible cost to sign, transmit, verify compared to OS update.

+Ed25519: unnoticeable cost.

Some extra system complexity, but the system includes Ed25519 code anyway.

Auditor sees very easily

that Ed25519+SPHINCS-256

security \geq Ed25519 security.

Does deployment of $P + Q$ mean that we don't trust Q ?
On the contrary!

Pre-quantum situation:

Hash-based signatures are even more confidence-inspiring than ECC signatures.

But understanding this fact takes extra work for auditor.

e.g. Ed25519+SPHINCS-256.

SPHINCS-256 signature is 41KB;

\approx 50 million cycles to generate;

\approx 1 million cycles to verify.

Negligible cost to sign, transmit, verify compared to OS update.

+Ed25519: unnoticeable cost.

Some extra system complexity, but the system includes

Ed25519 code anyway.

Auditor sees very easily

that Ed25519+SPHINCS-256

security \geq Ed25519 security.

Does deployment of $P + Q$ mean that we don't trust Q ?

On the contrary!

Pre-quantum situation:

Hash-based signatures are even more confidence-inspiring than ECC signatures.

But understanding this fact takes extra work for auditor.

e.g. Ed25519+SPHINCS-256.

SPHINCS-256 signature is 41KB;

\approx 50 million cycles to generate;

\approx 1 million cycles to verify.

Negligible cost to sign, transmit, verify compared to OS update.

+Ed25519: unnoticeable cost.

Some extra system complexity, but the system includes

Ed25519 code anyway.

Auditor sees very easily

that Ed25519+SPHINCS-256

security \geq Ed25519 security.

Does deployment of $P + Q$ mean that we don't trust Q ?

On the contrary!

Pre-quantum situation:

Hash-based signatures are even more confidence-inspiring than ECC signatures.

But understanding this fact takes extra work for auditor.

Long-term situation:

Users see quantum computers easily breaking P . Simplify system by switching from $P + Q$ to Q .

Ed25519+SPHINCS-256.

Ed25519 signature is 41KB;

10 million cycles to generate;

10 million cycles to verify.

Relative cost to sign, transmit,
and verify compared to OS update.

Ed25519: unnoticeable cost.

Extra system complexity,
if system includes

signature code anyway.

Ed25519 sees very easily

Ed25519+SPHINCS-256

\geq Ed25519 security.

Does deployment of $P + Q$
mean that we don't trust Q ?
On the contrary!

Pre-quantum situation:

Hash-based signatures are
even more confidence-inspiring
than ECC signatures.

But understanding this fact
takes extra work for auditor.

Long-term situation:

Users see quantum computers
easily breaking P . Simplify system
by switching from $P + Q$ to Q .

IP: Inter

IP comm

limited-l

Each com

has a 4-

e.g. www

address

Your bro

addresse

gives pa

Hopefull

that pac

HINCS-256.

signature is 41KB;

to generate;

to verify.

sign, transmit,

OS update.

noticeable cost.

in complexity,

cludes

way.

easily

HINCS-256

.9 security.

Does deployment of $P + Q$
mean that we don't trust Q ?
On the contrary!

Pre-quantum situation:

Hash-based signatures are
even more confidence-inspiring
than ECC signatures.

But understanding this fact
takes extra work for auditor.

Long-term situation:

Users see quantum computers
easily breaking P . Simplify system
by switching from $P + Q$ to Q .

IP: Internet Protocol

IP communicates
limited-length bytes

Each computer on
has a 4-byte "IP a
e.g. `www.pqcrypt`
address `131.155.`

Your browser creat
addressed to `131.`
gives packet to the
Hopefully the Inte
that packet to `131`

6.

.1KB;

te;

smi,

te.

st.

ity,

6

Does deployment of $P + Q$
mean that we don't trust Q ?
On the contrary!

Pre-quantum situation:

Hash-based signatures are
even more confidence-inspiring
than ECC signatures.

But understanding this fact
takes extra work for auditor.

Long-term situation:

Users see quantum computers
easily breaking P . Simplify system
by switching from $P + Q$ to Q .

IP: Internet Protocol

IP communicates "packets":
limited-length byte strings.

Each computer on the Internet
has a 4-byte "IP address".
e.g. `www.pqcrypto.org` has
address `131.155.70.11`.

Your browser creates a packet
addressed to `131.155.70.11`
gives packet to the Internet.
Hopefully the Internet delivers
that packet to `131.155.70`

Does deployment of $P + Q$
mean that we don't trust Q ?
On the contrary!

Pre-quantum situation:
Hash-based signatures are
even more confidence-inspiring
than ECC signatures.

But understanding this fact
takes extra work for auditor.

Long-term situation:
Users see quantum computers
easily breaking P . Simplify system
by switching from $P + Q$ to Q .

IP: Internet Protocol

IP communicates “packets”:
limited-length byte strings.

Each computer on the Internet
has a 4-byte “IP address” .
e.g. `www.pqcrypto.org` has
address `131.155.70.11`.

Your browser creates a packet
addressed to `131.155.70.11`;
gives packet to the Internet.
Hopefully the Internet delivers
that packet to `131.155.70.11`.

deployment of $P + Q$
 that we don't trust Q ?
 contrary!

quantum situation:
 signed signatures are
 more confidence-inspiring
 than C signatures.

Understanding this fact
 extra work for auditor.

quantum situation:
 the quantum computers
 are breaking P . Simplify system
 by moving from $P + Q$ to Q .

IP: Internet Protocol

IP communicates "packets":
 limited-length byte strings.

Each computer on the Internet
 has a 4-byte "IP address".
 e.g. `www.pqcrypto.org` has
 address `131.155.70.11`.

Your browser creates a packet
 addressed to `131.155.70.11`;
 gives packet to the Internet.
 Hopefully the Internet delivers
 that packet to `131.155.70.11`.

DNS: Domain Name System

You actually
 connect

Browser
 by asking
 the pqc

Browser
 "Where

of $P + Q$

't trust Q ?

ation:

ures are

nice-inspiring

res.

g this fact

or auditor.

on:

n computers

Simplify system

$P + Q$ to Q .

IP: Internet Protocol

IP communicates “packets” :
limited-length byte strings.

Each computer on the Internet
has a 4-byte “IP address” .
e.g. `www.pqcrypto.org` has
address `131.155.70.11`.

Your browser creates a packet
addressed to `131.155.70.11`;
gives packet to the Internet.
Hopefully the Internet delivers
that packet to `131.155.70.11`.

DNS: Domain Name

You actually told y
connect to `www.p`

Browser learns “13
by asking a name
the `pqcrypto.org`

Browser \rightarrow `131.1`
“Where is `www.p`

IP: Internet Protocol

IP communicates “packets” :
limited-length byte strings.

Each computer on the Internet
has a 4-byte “IP address” .
e.g. `www.pqcrypto.org` has
address `131.155.70.11`.

Your browser creates a packet
addressed to `131.155.70.11`;
gives packet to the Internet.
Hopefully the Internet delivers
that packet to `131.155.70.11`.

DNS: Domain Name System

You actually told your browser
to connect to `www.pqcrypto.org`.

Browser learns “`131.155.70.11`”
by asking a name server,
the `pqcrypto.org` name server.

Browser → `131.155.71.14`
“Where is `www.pqcrypto.org`?”

IP: Internet Protocol

IP communicates “packets” :
limited-length byte strings.

Each computer on the Internet
has a 4-byte “IP address” .

e.g. `www.pqcrypto.org` has
address `131.155.70.11`.

Your browser creates a packet
addressed to `131.155.70.11`;
gives packet to the Internet.

Hopefully the Internet delivers
that packet to `131.155.70.11`.

DNS: Domain Name System

You actually told your browser to
connect to `www.pqcrypto.org`.

Browser learns “`131.155.70.11`”
by asking a name server,
the `pqcrypto.org` name server.

Browser → `131.155.71.143`:
“Where is `www.pqcrypto.org`?”

IP: Internet Protocol

IP communicates “packets” :
limited-length byte strings.

Each computer on the Internet
has a 4-byte “IP address” .

e.g. `www.pqcrypto.org` has
address `131.155.70.11`.

Your browser creates a packet
addressed to `131.155.70.11`;
gives packet to the Internet.

Hopefully the Internet delivers
that packet to `131.155.70.11`.

DNS: Domain Name System

You actually told your browser to
connect to `www.pqcrypto.org`.

Browser learns “`131.155.70.11`”
by asking a name server,
the `pqcrypto.org` name server.

Browser → `131.155.71.143`:
“Where is `www.pqcrypto.org`?”

IP packet from browser also
includes a return address:
the address of your computer.

`131.155.71.143` → browser:
“`131.155.70.11`”

Internet Protocol

communicates “packets”:
variable length byte strings.

Every computer on the Internet
has a 32-bit byte “IP address”.

www.pqcrypto.org has
IP address 131.155.70.11.

Browser creates a packet
addressed to 131.155.70.11;
sends packet to the Internet.
By the Internet delivers
packet to 131.155.70.11.

DNS: Domain Name System

You actually told your browser to
connect to `www.pqcrypto.org`.

Browser learns “131.155.70.11”
by asking a name server,
the `pqcrypto.org` name server.

Browser → 131.155.71.143:

“Where is `www.pqcrypto.org`?”

IP packet from browser also
includes a return address:
the address of your computer.

131.155.71.143 → browser:

“131.155.70.11”

Browser
address,
by asking

Browser
“Where

199.19.

“Ask th
name se

col

“packets” :
e strings.

the Internet
address” .

o.org has
70.11.

tes a packet
155.70.11;
e Internet.

ernet delivers
1.155.70.11.

DNS: Domain Name System

You actually told your browser to
connect to `www.pqcrypto.org`.

Browser learns “131.155.70.11”
by asking a name server,
the `pqcrypto.org` name server.

Browser → 131.155.71.143:

“Where is `www.pqcrypto.org`?”

IP packet from browser also
includes a return address:
the address of your computer.

131.155.71.143 → browser:

“131.155.70.11”

Browser learns the
address, “131.155
by asking the .org

Browser → 199.1
“Where is `www.p`

199.19.54.1 →

“Ask the `pqcryp`
name server, 13

DNS: Domain Name System

You actually told your browser to connect to `www.pqcrypto.org`.

Browser learns “131.155.70.11” by asking a name server, the `pqcrypto.org` name server.

Browser → 131.155.71.143:
“Where is `www.pqcrypto.org`?”

IP packet from browser also includes a return address: the address of your computer.

131.155.71.143 → browser:
“131.155.70.11”

Browser learns the name-server address, “131.155.71.143” by asking the `.org` name server.

Browser → 199.19.54.1:
“Where is `www.pqcrypto.org`?”

199.19.54.1 → browser:
“Ask the `pqcrypto.org` name server, 131.155.71.143”

DNS: Domain Name System

You actually told your browser to connect to `www.pqcrypto.org`.

Browser learns “131.155.70.11” by asking a name server, the `pqcrypto.org` name server.

Browser → 131.155.71.143:
“Where is `www.pqcrypto.org`?”

IP packet from browser also includes a return address: the address of your computer.

131.155.71.143 → browser:
“131.155.70.11”

Browser learns the name-server address, “131.155.71.143”, by asking the `.org` name server.

Browser → 199.19.54.1:
“Where is `www.pqcrypto.org`?”

199.19.54.1 → browser:
“Ask the `pqcrypto.org` name server, 131.155.71.143”

DNS: Domain Name System

You actually told your browser to connect to `www.pqcrypto.org`.

Browser learns “131.155.70.11” by asking a name server, the `pqcrypto.org` name server.

Browser → 131.155.71.143:
“Where is `www.pqcrypto.org`?”

IP packet from browser also includes a return address: the address of your computer.

131.155.71.143 → browser:
“131.155.70.11”

Browser learns the name-server address, “131.155.71.143”, by asking the `.org` name server.

Browser → 199.19.54.1:
“Where is `www.pqcrypto.org`?”

199.19.54.1 → browser:
“Ask the `pqcrypto.org` name server, 131.155.71.143”

Browser learns “199.19.54.1”, the `.org` server address, by asking the root name server.

DNS: Domain Name System

You actually told your browser to connect to `www.pqcrypto.org`.

Browser learns “131.155.70.11” by asking a name server, the `pqcrypto.org` name server.

Browser → 131.155.71.143:
“Where is `www.pqcrypto.org`?”

IP packet from browser also includes a return address: the address of your computer.

131.155.71.143 → browser:
“131.155.70.11”

Browser learns the name-server address, “131.155.71.143”, by asking the `.org` name server.

Browser → 199.19.54.1:
“Where is `www.pqcrypto.org`?”

199.19.54.1 → browser:
“Ask the `pqcrypto.org` name server, 131.155.71.143”

Browser learns “199.19.54.1”, the `.org` server address, by asking the root name server.

Browser learned root address by consulting the Bible.

Domain Name System

usually told your browser to
to `www.pqcrypto.org`.

learns “131.155.70.11”
g a name server,
`pqcrypto.org` name server.

→ 131.155.71.143:
“Where is `www.pqcrypto.org`?”

et from browser also
a return address:
ress of your computer.

5.71.143 → browser:
55.70.11”

TCP: Tr

Packets

(Actually

Oldest II

≥576. U

often 15

Browser learns the name-server
address, “131.155.71.143”,
by asking the `.org` name server.

Browser → 199.19.54.1:
“Where is `www.pqcrypto.org`?”

199.19.54.1 → browser:

“Ask the `pqcrypto.org`
name server, 131.155.71.143”

Browser learns “199.19.54.1”,
the `.org` server address,
by asking the root name server.

Browser learned root address
by consulting the Bible.

name System

your browser to
pqcrypto.org.

131.155.70.11"

server,

g name server.

155.71.143:

pqcrypto.org?"

rowser also

address:

r computer.

→ browser:

,

Browser learns the name-server
address, "131.155.71.143",
by asking the .org name server.

Browser → 199.19.54.1:

"Where is www.pqcrypto.org?"

199.19.54.1 → browser:

"Ask the pqcrypto.org
name server, 131.155.71.143"

Browser learns "199.19.54.1",
the .org server address,
by asking the root name server.

Browser learned root address
by consulting the Bible.

TCP: Transmission

Packets are limited

(Actually depends

Oldest IP standard

≥576. Usually 149

often 1500, somet

Browser learns the name-server address, "131.155.71.143", by asking the .org name server.

Browser → 199.19.54.1:

"Where is www.pqcrypto.org?"

199.19.54.1 → browser:

"Ask the pqcrypto.org name server, 131.155.71.143"

Browser learns "199.19.54.1", the .org server address, by asking the root name server.

Browser learned root address by consulting the Bible.

TCP: Transmission Control

Packets are limited to 1280

(Actually depends on network

Oldest IP standards required

≥576. Usually 1492 is safe,

often 1500, sometimes more

Browser learns the name-server address, “131.155.71.143”, by asking the .org name server.

Browser → 199.19.54.1:

“Where is www.pqcrypto.org?”

199.19.54.1 → browser:

“Ask the pqcrypto.org name server, 131.155.71.143”

Browser learns “199.19.54.1”, the .org server address, by asking the root name server.

Browser learned root address by consulting the Bible.

TCP: Transmission Control Protocol

Packets are limited to 1280 bytes.

(Actually depends on network. Oldest IP standards required ≥ 576 . Usually 1492 is safe, often 1500, sometimes more.)

Browser learns the name-server address, “131.155.71.143”, by asking the .org name server.

Browser → 199.19.54.1:

“Where is `www.pqcrypto.org`?”

199.19.54.1 → browser:

“Ask the `pqcrypto.org` name server, 131.155.71.143”

Browser learns “199.19.54.1”, the .org server address, by asking the root name server.

Browser learned root address by consulting the Bible.

TCP: Transmission Control Protocol

Packets are limited to 1280 bytes.

(Actually depends on network. Oldest IP standards required ≥ 576 . Usually 1492 is safe, often 1500, sometimes more.)

The page you’re downloading from `pqcrypto.org` doesn’t fit.

Browser learns the name-server address, “131.155.71.143”, by asking the .org name server.

Browser → 199.19.54.1:

“Where is `www.pqcrypto.org`?”

199.19.54.1 → browser:

“Ask the `pqcrypto.org` name server, 131.155.71.143”

Browser learns “199.19.54.1”, the .org server address, by asking the root name server.

Browser learned root address by consulting the Bible.

TCP: Transmission Control Protocol

Packets are limited to 1280 bytes.

(Actually depends on network. Oldest IP standards required ≥ 576 . Usually 1492 is safe, often 1500, sometimes more.)

The page you’re downloading from `pqcrypto.org` doesn’t fit.

Browser actually makes “TCP connection” to `pqcrypto.org`. Inside that connection: sends HTTP request, receives response.

learns the name-server
 “131.155.71.143”,
 g the .org name server.
 → 199.19.54.1:
 is `www.pqcrypto.org`?
 .54.1 → browser:
 e `pqcrypto.org`
 rver, 131.155.71.143”
 learns “199.19.54.1”,
 g server address,
 g the root name server.
 learned root address
 ulting the Bible.

TCP: Transmission Control Protocol

Packets are limited to 1280 bytes.

(Actually depends on network.

Oldest IP standards required

≥ 576 . Usually 1492 is safe,
 often 1500, sometimes more.)

The page you’re downloading
 from `pqcrypto.org` doesn’t fit.

Browser actually makes “TCP
 connection” to `pqcrypto.org`.

Inside that connection: sends
 HTTP request, receives response.

Browser
 “SYN 16
 Server –
 “ACK 16
 Browser
 “ACK 74
 Server n
 for this
 Browser
 counting
 Server s
 counting

e name-server
 5.71.143",
 g name server.
 9.54.1:
 pqcrypto.org?"
 browser:
 to.org
 1.155.71.143"
 99.19.54.1",
 address,
 name server.
 bot address
 Bible.

TCP: Transmission Control Protocol

Packets are limited to 1280 bytes.

(Actually depends on network.

Oldest IP standards required

≥ 576 . Usually 1492 is safe,

often 1500, sometimes more.)

The page you're downloading
 from pqcrypto.org doesn't fit.

Browser actually makes "TCP
 connection" to pqcrypto.org.

Inside that connection: sends
 HTTP request, receives response.

Browser \rightarrow server:

"SYN 168bb5d9"

Server \rightarrow browser:

"ACK 168bb5da, S"

Browser \rightarrow server:

"ACK 747bfa42"

Server now allocates
 for this TCP conn

Browser splits data
 counting bytes fro

Server splits data
 counting bytes fro

TCP: Transmission Control Protocol

Packets are limited to 1280 bytes.

(Actually depends on network.

Oldest IP standards required ≥ 576 . Usually 1492 is safe, often 1500, sometimes more.)

The page you're downloading from `pqcrypto.org` doesn't fit.

Browser actually makes "TCP connection" to `pqcrypto.org`.

Inside that connection: sends HTTP request, receives response.

Browser \rightarrow server:

"SYN 168bb5d9"

Server \rightarrow browser:

"ACK 168bb5da, SYN 747bf1"

Browser \rightarrow server:

"ACK 747bfa42"

Server now allocates buffers for this TCP connection.

Browser splits data into packets counting bytes from 168bb5

Server splits data into packets counting bytes from 747bfa

TCP: Transmission Control Protocol

Packets are limited to 1280 bytes.

(Actually depends on network.

Oldest IP standards required

≥ 576 . Usually 1492 is safe,
often 1500, sometimes more.)

The page you're downloading
from `pqcrypto.org` doesn't fit.

Browser actually makes "TCP
connection" to `pqcrypto.org`.

Inside that connection: sends
HTTP request, receives response.

Browser \rightarrow server:

"SYN 168bb5d9"

Server \rightarrow browser:

"ACK 168bb5da, SYN 747bfa41"

Browser \rightarrow server:

"ACK 747bfa42"

Server now allocates buffers
for this TCP connection.

Browser splits data into packets,
counting bytes from 168bb5da.

Server splits data into packets,
counting bytes from 747bfa42.

Transmission Control Protocol

are limited to 1280 bytes.

y depends on network.

P standards required

Usually 1492 is safe,

00, sometimes more.)

When you're downloading

crypto.org doesn't fit.

Browser actually makes "TCP

connection" to pqcrypto.org.

Browser sends request:

Browser sends request, receives response.

Browser → server:

"SYN 168bb5d9"

Server → browser:

"ACK 168bb5da, SYN 747bfa41"

Browser → server:

"ACK 747bfa42"

Server now allocates buffers
for this TCP connection.

Browser splits data into packets,
counting bytes from 168bb5da.

Server splits data into packets,
counting bytes from 747bfa42.

Main feature

"reliable"

Internet

or delivered

Doesn't

computer

inside each

Computer

if data is

Complicated

retransmission

avoiding

Transmission Control Protocol

limited to 1280 bytes.

on network.

bytes required

92 is safe,

times more.)

downloading

arg doesn't fit.

makes "TCP

crypto.org.

tion: sends

ceives response.

Browser → server:

"SYN 168bb5d9"

Server → browser:

"ACK 168bb5da, SYN 747bfa41"

Browser → server:

"ACK 747bfa42"

Server now allocates buffers
for this TCP connection.

Browser splits data into packets,
counting bytes from 168bb5da.

Server splits data into packets,
counting bytes from 747bfa42.

Main feature advertisement

"reliable data stream"

Internet sometimes

or delivers packets

Doesn't confuse TCP

computer checks to

inside each TCP p

Computer retransmits

if data is not acknowledged

Complicated rules

retransmission schedule

avoiding network congestion

Protocol

bytes.

rk.

|

e.)

g

t fit.

CP

rg.

ls

ponse.

Browser → server:

“SYN 168bb5d9”

Server → browser:

“ACK 168bb5da, SYN 747bfa41”

Browser → server:

“ACK 747bfa42”

Server now allocates buffers
for this TCP connection.

Browser splits data into packets,
counting bytes from 168bb5da.

Server splits data into packets,
counting bytes from 747bfa42.

Main feature advertised by T
“reliable data streams” .

Internet sometimes loses pack
or delivers packets out of ord
Doesn't confuse TCP connec
computer checks the counte
inside each TCP packet.

Computer retransmits data
if data is not acknowledged.
Complicated rules to decide
retransmission schedule,
avoiding network congestion

Browser → server:

“SYN 168bb5d9”

Server → browser:

“ACK 168bb5da, SYN 747bfa41”

Browser → server:

“ACK 747bfa42”

Server now allocates buffers
for this TCP connection.

Browser splits data into packets,
counting bytes from 168bb5da.

Server splits data into packets,
counting bytes from 747bfa42.

Main feature advertised by TCP:
“reliable data streams” .

Internet sometimes loses packets
or delivers packets out of order.

Doesn't confuse TCP connections:
computer checks the counter
inside each TCP packet.

Computer retransmits data
if data is not acknowledged.
Complicated rules to decide
retransmission schedule,
avoiding network congestion.

→ server:

8bb5d9”

→ browser:

8bb5da, SYN 747bfa41”

→ server:

7bfa42”

ow allocates buffers

TCP connection.

splits data into packets,

g bytes from 168bb5da.

splits data into packets,

g bytes from 747bfa42.

Main feature advertised by TCP:

“reliable data streams” .

Internet sometimes loses packets

or delivers packets out of order.

Doesn't confuse TCP connections:

computer checks the counter

inside each TCP packet.

Computer retransmits data

if data is not acknowledged.

Complicated rules to decide

retransmission schedule,

avoiding network congestion.

Stream-l

<http://>

uses HT

<https://>

uses HT

Your bro

- finds a

- makes

- inside

- builds

- by exc

- inside

- sends

Main feature advertised by TCP:
“reliable data streams” .

Internet sometimes loses packets
or delivers packets out of order.
Doesn't confuse TCP connections:
computer checks the counter
inside each TCP packet.

Computer retransmits data
if data is not acknowledged.
Complicated rules to decide
retransmission schedule,
avoiding network congestion.

Stream-level crypt

<http://www.pqc.com>

uses HTTP over T

<https://www.pqc.com>

uses HTTP over T

Your browser

- finds address 13
- makes TCP conn
- inside the TCP c
builds a TLS con
by exchanging c
- inside the TLS c
sends HTTP rec

Main feature advertised by TCP:
“reliable data streams” .

Internet sometimes loses packets
or delivers packets out of order.
Doesn't confuse TCP connections:
computer checks the counter
inside each TCP packet.

Computer retransmits data
if data is not acknowledged.
Complicated rules to decide
retransmission schedule,
avoiding network congestion.

Stream-level crypto

<http://www.pqcrypto.org>
uses HTTP over TCP.

<https://www.pqcrypto.org>
uses HTTP over TLS over T

Your browser

- finds address 131.155.70
- makes TCP connection;
- inside the TCP connection
builds a TLS connection
by exchanging crypto keys
- inside the TLS connection
sends HTTP request etc.

Main feature advertised by TCP:
“reliable data streams”.

Internet sometimes loses packets
or delivers packets out of order.

Doesn't confuse TCP connections:
computer checks the counter
inside each TCP packet.

Computer retransmits data
if data is not acknowledged.
Complicated rules to decide
retransmission schedule,
avoiding network congestion.

Stream-level crypto

<http://www.pqcrypto.org>

uses HTTP over TCP.

<https://www.pqcrypto.org>

uses HTTP over TLS over TCP.

Your browser

- finds address 131.155.70.11;
- makes TCP connection;
- inside the TCP connection,
builds a TLS connection
by exchanging crypto keys;
- inside the TLS connection,
sends HTTP request etc.

signature advertised by TCP:
"data streams".

sometimes loses packets
reorders packets out of order.

can confuse TCP connections:
server checks the counter
for each TCP packet.

server retransmits data
if not acknowledged.
uses various rules to decide
transmission schedule,
and network congestion.

Stream-level crypto

<http://www.pqcrypto.org>

uses HTTP over TCP.

<https://www.pqcrypto.org>

uses HTTP over TLS over TCP.

Your browser

- finds address 131.155.70.11;
- makes TCP connection;
- inside the TCP connection,
builds a TLS connection
by exchanging crypto keys;
- inside the TLS connection,
sends HTTP request etc.

What happens if
server forges a
signature pointing
to a different server.
Or a TCP connection
with bogus data.

DNS software
TCP software
TLS software
something else
but has no idea.

Browser
can't make a
connection
but this is
Huge data

rtised by TCP:

ams”.

s loses packets

s out of order.

TCP connections:

the counter

packet.

mits data

nowledged.

to decide

edule,

congestion.

Stream-level crypto

<http://www.pqcrypto.org>

uses HTTP over TCP.

<https://www.pqcrypto.org>

uses HTTP over TLS over TCP.

Your browser

- finds address 131.155.70.11;
- makes TCP connection;
- inside the TCP connection, builds a TLS connection by exchanging crypto keys;
- inside the TLS connection, sends HTTP request etc.

What happens if a

forges a DNS pack

pointing to fake se

Or a TCP packet

with bogus data?

DNS software is fo

TCP software is fo

TLS software sees

something has gon

but has no way to

Browser using TLS

make a whole new

but this is slow an

Huge damage from

TCP:

Stream-level crypto<http://www.pqcrypto.org>

uses HTTP over TCP.

<https://www.pqcrypto.org>

uses HTTP over TLS over TCP.

Your browser

- finds address 131.155.70.11;
- makes TCP connection;
- inside the TCP connection, builds a TLS connection by exchanging crypto keys;
- inside the TLS connection, sends HTTP request etc.

What happens if attacker forges a DNS packet pointing to fake server? Or a TCP packet with bogus data?

DNS software is fooled.
TCP software is fooled.
TLS software sees that something has gone wrong, but has no way to recover.

Browser using TLS can make a whole new connection but this is slow and fragile.
Huge damage from forged p

Stream-level crypto

<http://www.pqcrypto.org>

uses HTTP over TCP.

<https://www.pqcrypto.org>

uses HTTP over TLS over TCP.

Your browser

- finds address 131.155.70.11;
- makes TCP connection;
- inside the TCP connection, builds a TLS connection by exchanging crypto keys;
- inside the TLS connection, sends HTTP request etc.

What happens if attacker forges a DNS packet pointing to fake server? Or a TCP packet with bogus data?

DNS software is fooled.
TCP software is fooled.
TLS software sees that something has gone wrong, but has no way to recover.

Browser using TLS can make a whole new connection, but this is slow and fragile.
Huge damage from forged packet.

level crypto

[/www.pqcrypto.org](http://www.pqcrypto.org)

TP over TCP.

[/www.pqcrypto.org](http://www.pqcrypto.org)

TP over TLS over TCP.

rowser

address 131.155.70.11;

TCP connection;

the TCP connection,

a TLS connection

changing crypto keys;

the TLS connection,

HTTP request etc.

What happens if attacker
forges a DNS packet
pointing to fake server?
Or a TCP packet
with bogus data?

DNS software is fooled.
TCP software is fooled.
TLS software sees that
something has gone wrong,
but has no way to recover.

Browser using TLS can
make a whole new connection,
but this is slow and fragile.
Huge damage from forged packet.

Modern
CurveCF
Google's
encrypt

Discard
immedia
Retransm
authenti

What happens if attacker forges a DNS packet pointing to fake server? Or a TCP packet with bogus data?

DNS software is fooled.

TCP software is fooled.

TLS software sees that something has gone wrong, but has no way to recover.

Browser using TLS can make a whole new connection, but this is slow and fragile.

Huge damage from forged packet.

Modern trend (e.g. CurveCP; see also Google's QUIC): Authenticate and encrypt each packet.

Discard forged packets immediately: no damage.

Retransmit packet if not *authenticated* acknowledgment.

What happens if attacker forges a DNS packet pointing to fake server? Or a TCP packet with bogus data?

DNS software is fooled.
TCP software is fooled.
TLS software sees that something has gone wrong, but has no way to recover.

Browser using TLS can make a whole new connection, but this is slow and fragile.
Huge damage from forged packet.

Modern trend (e.g., DNSCurve, CurveCP; see also MinimalTLS, Google's QUIC): Authenticate and encrypt each packet separately.

Discard forged packet immediately: no damage.
Retransmit packet if no *authenticated* acknowledgment.

What happens if attacker forges a DNS packet pointing to fake server?
Or a TCP packet with bogus data?

DNS software is fooled.
TCP software is fooled.
TLS software sees that something has gone wrong, but has no way to recover.

Browser using TLS can make a whole new connection, but this is slow and fragile.
Huge damage from forged packet.

Modern trend (e.g., DNSCurve, CurveCP; see also MinimaLT, Google's QUIC): Authenticate and encrypt each packet separately.

Discard forged packet immediately: no damage.

Retransmit packet if no *authenticated* acknowledgment.

What happens if attacker forges a DNS packet pointing to fake server?
Or a TCP packet with bogus data?

DNS software is fooled.
TCP software is fooled.
TLS software sees that something has gone wrong, but has no way to recover.

Browser using TLS can make a whole new connection, but this is slow and fragile.
Huge damage from forged packet.

Modern trend (e.g., DNSCurve, CurveCP; see also MinimaLT, Google's QUIC): Authenticate and encrypt each packet separately.

Discard forged packet immediately: no damage.
Retransmit packet if no *authenticated* acknowledgment.

Engineering advantage:
Packet-level crypto works for more protocols than stream-level crypto.

What happens if attacker forges a DNS packet pointing to fake server?
Or a TCP packet with bogus data?

DNS software is fooled.
TCP software is fooled.
TLS software sees that something has gone wrong, but has no way to recover.

Browser using TLS can make a whole new connection, but this is slow and fragile.
Huge damage from forged packet.

Modern trend (e.g., DNSCurve, CurveCP; see also MinimaLT, Google's QUIC): Authenticate and encrypt each packet separately.

Discard forged packet immediately: no damage.
Retransmit packet if no *authenticated* acknowledgment.

Engineering advantage:
Packet-level crypto works for more protocols than stream-level crypto.

Disadvantage:
Crypto must fit into packet.

happens if attacker
 DNS packet
 to fake server?
 CP packet
 gus data?
 ftware is fooled.
 ftware is fooled.
 ftware sees that
 ng has gone wrong,
 no way to recover.
 using TLS can
 whole new connection,
 is slow and fragile.
 mage from forged packet.

Modern trend (e.g., DNSCurve,
 CurveCP; see also MinimaLT,
 Google's QUIC): Authenticate and
 encrypt each packet separately.

Discard forged packet
 immediately: no damage.
 Retransmit packet if no
authenticated acknowledgment.

Engineering advantage:
 Packet-level crypto
 works for more protocols
 than stream-level crypto.

Disadvantage:
 Crypto must fit into packet.

The KE
 Original
 Message
 as m^e m

attacker
 packet
 server?

poled.
 poled.
 that
 ne wrong,
 recover.
 S can
 y connection,
 d fragile.
 n forged packet.

Modern trend (e.g., DNSCurve, CurveCP; see also MinimaLT, Google's QUIC): Authenticate and encrypt each packet separately.

Discard forged packet immediately: no damage.
 Retransmit packet if no *authenticated* acknowledgment.

Engineering advantage:
 Packet-level crypto works for more protocols than stream-level crypto.

Disadvantage:
 Crypto must fit into packet.

The KEM+AE ph

Original view of R
 Message m is encr
 as $m^e \bmod pq$.

Modern trend (e.g., DNSCurve, CurveCP; see also MinimaLT, Google's QUIC): Authenticate and encrypt each packet separately.

Discard forged packet immediately: no damage.

Retransmit packet if no *authenticated* acknowledgment.

Engineering advantage:

Packet-level crypto works for more protocols than stream-level crypto.

Disadvantage:

Crypto must fit into packet.

The KEM+AE philosophy

Original view of RSA:
Message m is encrypted as $m^e \bmod pq$.

Modern trend (e.g., DNSCurve, CurveCP; see also MinimaLT, Google's QUIC): Authenticate and encrypt each packet separately.

Discard forged packet immediately: no damage.

Retransmit packet if no *authenticated* acknowledgment.

Engineering advantage:

Packet-level crypto

works for more protocols than stream-level crypto.

Disadvantage:

Crypto must fit into packet.

The KEM+AE philosophy

Original view of RSA:

Message m is encrypted as $m^e \bmod pq$.

Modern trend (e.g., DNSCurve, CurveCP; see also MinimaLT, Google's QUIC): Authenticate and encrypt each packet separately.

Discard forged packet immediately: no damage.

Retransmit packet if no *authenticated* acknowledgment.

Engineering advantage:

Packet-level crypto

works for more protocols than stream-level crypto.

Disadvantage:

Crypto must fit into packet.

The KEM+AE philosophy

Original view of RSA:

Message m is encrypted as $m^e \bmod pq$.

“Hybrid” view of RSA,

including random padding:

Choose random AES-GCM key k .

Randomly pad k as r .

Encrypt r as $r^e \bmod pq$.

Encrypt m under k .

Modern trend (e.g., DNSCurve, CurveCP; see also MinimaLT, Google's QUIC): Authenticate and encrypt each packet separately.

Discard forged packet immediately: no damage.

Retransmit packet if no *authenticated* acknowledgment.

Engineering advantage:

Packet-level crypto

works for more protocols than stream-level crypto.

Disadvantage:

Crypto must fit into packet.

The KEM+AE philosophy

Original view of RSA:

Message m is encrypted as $m^e \bmod pq$.

“Hybrid” view of RSA,

including random padding:

Choose random AES-GCM key k .

Randomly pad k as r .

Encrypt r as $r^e \bmod pq$.

Encrypt m under k .

Fragile, many problems:

e.g., Coppersmith attack,

Bleichenbacher attack,

bogus OAEP security proof.

trend (e.g., DNSCurve,
 P; see also MinimaLT,
 s QUIC): Authenticate and
 each packet separately.

forged packet

tely: no damage.

mit packet if no

icated acknowledgment.

ring advantage:

level crypto

r more protocols

eam-level crypto.

ntage:

must fit into packet.

The KEM+AE philosophy

Original view of RSA:

Message m is encrypted
 as $m^e \bmod pq$.

“Hybrid” view of RSA,

including random padding:

Choose random AES-GCM key k .

Randomly pad k as r .

Encrypt r as $r^e \bmod pq$.

Encrypt m under k .

Fragile, many problems:

e.g., Coppersmith attack,

Bleichenbacher attack,

bogus OAEP security proof.

Shoup’s

“Key en

Choose

Encrypt

Define k

“Data en

Encrypt

m under

Authent

any mod

Much ea

Also gen

$P + Q$:

..., DNSCurve,
 MinimaLT,
 Authenticate and
 et separately.

cket
 amage.
 if no
 nowledgment.

tage:

otocols
 crypto.

to packet.

The KEM+AE philosophy

Original view of RSA:

Message m is encrypted
 as $m^e \bmod pq$.

“Hybrid” view of RSA,
 including random padding:
 Choose random AES-GCM key k .

Randomly pad k as r .

Encrypt r as $r^e \bmod pq$.

Encrypt m under k .

Fragile, many problems:
 e.g., Coppersmith attack,
 Bleichenbacher attack,
 bogus OAEP security proof.

Shoup’s “KEM+D

“Key encapsulation

Choose random r

Encrypt r as $r^e \bmod pq$

Define $k = H(r, r^e \bmod pq)$

“Data encapsulation

Encrypt and authen

m under AES-GCM

Authenticator catch

any modification o

Much easier to ge

Also generalizes ni

$P + Q$: hash conc

The KEM+AE philosophy

Original view of RSA:

Message m is encrypted
as $m^e \bmod pq$.

“Hybrid” view of RSA,

including random padding:

Choose random AES-GCM key k .

Randomly pad k as r .

Encrypt r as $r^e \bmod pq$.

Encrypt m under k .

Fragile, many problems:

e.g., Coppersmith attack,

Bleichenbacher attack,

bogus OAEP security proof.

Shoup’s “KEM+DEM” view

“Key encapsulation mechanism”

Choose random $r \bmod pq$.

Encrypt r as $r^e \bmod pq$.

Define $k = H(r, r^e \bmod pq)$

“Data encapsulation mechanism”

Encrypt and authenticate

m under AES-GCM key k .

Authenticator catches
any modification of $r^e \bmod pq$

Much easier to get right.

Also generalizes nicely.

$P + Q$: hash concatenation.

The KEM+AE philosophy

Original view of RSA:

Message m is encrypted
as $m^e \bmod pq$.

“Hybrid” view of RSA,
including random padding:
Choose random AES-GCM key k .
Randomly pad k as r .
Encrypt r as $r^e \bmod pq$.
Encrypt m under k .

Fragile, many problems:
e.g., Coppersmith attack,
Bleichenbacher attack,
bogus OAEP security proof.

Shoup’s “KEM+DEM” view:

“Key encapsulation mechanism”:

Choose random $r \bmod pq$.

Encrypt r as $r^e \bmod pq$.

Define $k = H(r, r^e \bmod pq)$.

“Data encapsulation mechanism”:

Encrypt and authenticate
 m under AES-GCM key k .

Authenticator catches
any modification of $r^e \bmod pq$.

Much easier to get right.

Also generalizes nicely.

$P + Q$: hash concatenation.

M+AE philosophy

view of RSA:

m is encrypted
mod pq .

' view of RSA,

g random padding:

random AES-GCM key k .

ly pad k as r .

r as $r^e \bmod pq$.

m under k .

many problems:

oppersmith attack,

bacher attack,

AEP security proof.

Shoup's "KEM+DEM" view:

"Key encapsulation mechanism":

Choose random $r \bmod pq$.

Encrypt r as $r^e \bmod pq$.

Define $k = H(r, r^e \bmod pq)$.

"Data encapsulation mechanism":

Encrypt and authenticate

m under AES-GCM key k .

Authenticator catches

any modification of $r^e \bmod pq$.

Much easier to get right.

Also generalizes nicely.

$P + Q$: hash concatenation.

DEM se

weak sim

of securi

authenti

Chou: Is

for mult

Answer:

KEM+A

(But nee

AES-GC

aim for t

More co

Use KEM

n -time s

philosophy

SA:

rypted

RSA,

padding:

AES-GCM key k .

as r .

od pq .

k .

blems:

attack,

tack,

urity proof.

Shoup's "KEM+DEM" view:

"Key encapsulation mechanism":

Choose random $r \bmod pq$.

Encrypt r as $r^e \bmod pq$.

Define $k = H(r, r^e \bmod pq)$.

"Data encapsulation mechanism":

Encrypt and authenticate

m under AES-GCM key k .

Authenticator catches

any modification of $r^e \bmod pq$.

Much easier to get right.

Also generalizes nicely.

$P + Q$: hash concatenation.

DEM security hyp

weak single-messa

of security for secr

authenticated encr

Chou: Is it safe to

for multiple messa

Answer: KEM+AE

$KEM+AE \Rightarrow KEM$

(But need literatu

AES-GCM, Salsa2

aim for full AE sec

More complicated

Use KEM+DEM t

n -time secret key

Shoup's "KEM+DEM" view:

"Key encapsulation mechanism":

Choose random $r \bmod pq$.

Encrypt r as $r^e \bmod pq$.

Define $k = H(r, r^e \bmod pq)$.

"Data encapsulation mechanism":

Encrypt and authenticate
 m under AES-GCM key k .

Authenticator catches
any modification of $r^e \bmod pq$.

Much easier to get right.

Also generalizes nicely.

$P + Q$: hash concatenation.

key k .

DEM security hypothesis:

weak single-message version
of security for secret-key
authenticated encryption.

Chou: Is it safe to reuse k
for multiple messages?

Answer: KEM+AE is safe;
KEM+AE \Rightarrow KEM+" n DEM"
(But need literature on this!
AES-GCM, Salsa20-Poly1305
aim for full AE security goal)

More complicated alternative
Use KEM+DEM to encrypt
 n -time secret key m ; reuse r

Shoup's "KEM+DEM" view:

"Key encapsulation mechanism":

Choose random $r \bmod pq$.

Encrypt r as $r^e \bmod pq$.

Define $k = H(r, r^e \bmod pq)$.

"Data encapsulation mechanism":

Encrypt and authenticate
 m under AES-GCM key k .

Authenticator catches
any modification of $r^e \bmod pq$.

Much easier to get right.

Also generalizes nicely.

$P + Q$: hash concatenation.

DEM security hypothesis:
weak single-message version
of security for secret-key
authenticated encryption.

Chou: Is it safe to reuse k
for multiple messages?

Answer: KEM+AE is safe;
 $\text{KEM+AE} \Rightarrow \text{KEM+}n\text{DEM}$.
(But need literature on this!)
AES-GCM, Salsa20-Poly1305, etc.
aim for full AE security goal.

More complicated alternative:
Use KEM+DEM to encrypt an
 n -time secret key m ; reuse m .

“KEM+DEM” view:

“encapsulation mechanism”:

random $r \bmod pq$.

r as $r^e \bmod pq$.

$k = H(r, r^e \bmod pq)$.

“decapsulation mechanism”:

decrypt and authenticate

with AES-GCM key k .

MAC indicator catches

modification of $r^e \bmod pq$.

easier to get right.

generalizes nicely.

hash concatenation.

DEM security hypothesis:

weak single-message version

of security for secret-key

authenticated encryption.

Chou: Is it safe to reuse k

for multiple messages?

Answer: KEM+AE is safe;

$\text{KEM+AE} \Rightarrow \text{KEM+“}n\text{DEM”}$.

(But need literature on this!)

AES-GCM, Salsa20-Poly1305, etc.

aim for full AE security goal.

More complicated alternative:

Use KEM+DEM to encrypt an

n -time secret key m ; reuse m .

DNSCur

Server k

Client k

server's

Client —

packet c

where k

E is aut

q is DNS

Server —

packet c

where r

DEM" view:

on mechanism":

mod pq .

od pq .

$e \bmod pq$).

on mechanism":

entiate

M key k .

ches

of $r^e \bmod pq$.

t right.

cely.

atenation.

DEM security hypothesis:

weak single-message version
of security for secret-key
authenticated encryption.

Chou: Is it safe to reuse k
for multiple messages?

Answer: KEM+AE is safe;

KEM+AE \Rightarrow KEM+"nDEM".

(But need literature on this!)

AES-GCM, Salsa20-Poly1305, etc.

aim for full AE security goal.

More complicated alternative:

Use KEM+DEM to encrypt an
 n -time secret key m ; reuse m .

DNSCurve: ECDH

Server knows ECD

Client knows ECD
server's public key

Client \rightarrow server:

packet containing

where $k = H(cS)$;

E is authenticated

q is DNS query.

Server \rightarrow client:

packet containing

where r is DNS re

DEM security hypothesis:
 weak single-message version
 of security for secret-key
 authenticated encryption.

Chou: Is it safe to reuse k
 for multiple messages?

Answer: KEM+AE is safe;
 $\text{KEM+AE} \Rightarrow \text{KEM+“}n\text{DEM”}$.
 (But need literature on this!)
 AES-GCM, Salsa20-Poly1305, etc.
 aim for full AE security goal.

More complicated alternative:
 Use KEM+DEM to encrypt an
 n -time secret key m ; reuse m .

DNSSCurve: ECDH for DNS

Server knows ECDH secret k
 Client knows ECDH secret k
 server's public key $S = sG$.

Client \rightarrow server:
 packet containing $cG, E_k(0, q)$,
 where $k = H(cS)$;
 E is authenticated cipher;
 q is DNS query.

Server \rightarrow client:
 packet containing $E_k(1, r)$
 where r is DNS response.

DEM security hypothesis:

weak single-message version
of security for secret-key
authenticated encryption.

Chou: Is it safe to reuse k
for multiple messages?

Answer: KEM+AE is safe;
KEM+AE \Rightarrow KEM+“ n DEM”.
(But need literature on this!)
AES-GCM, Salsa20-Poly1305, etc.
aim for full AE security goal.

More complicated alternative:
Use KEM+DEM to encrypt an
 n -time secret key m ; reuse m .

DNSSCurve: ECDH for DNS

Server knows ECDH secret key s .

Client knows ECDH secret key c ,
server's public key $S = sG$.

Client \rightarrow server:

packet containing $cG, E_k(0, q)$

where $k = H(cS)$;

E is authenticated cipher;

q is DNS query.

Server \rightarrow client:

packet containing $E_k(1, r)$

where r is DNS response.

curity hypothesis:
 single-message version
 ty for secret-key
 cated encryption.

s it safe to reuse k
 iple messages?

KEM+AE is safe;
 AE \Rightarrow KEM+“nDEM”.

ed literature on this!)

M, Salsa20-Poly1305, etc.

full AE security goal.

mplicated alternative:

M+DEM to encrypt an

ecret key m ; reuse m .

DNSSCurve: ECDH for DNS

Server knows ECDH secret key s .

Client knows ECDH secret key c ,
 server's public key $S = sG$.

Client \rightarrow server:

packet containing $cG, E_k(0, q)$

where $k = H(cS)$;

E is authenticated cipher;

q is DNS query.

Server \rightarrow client:

packet containing $E_k(1, r)$

where r is DNS response.

Client ca
 across m
 but this
 Let's ass

othesis:
 ge version
 ret-key
 ryption.
 o reuse k
 ges?
 E is safe;
 $M + "nDEM"$.
 re on this!)
 0-Poly1305, etc.
 curity goal.
 alternative:
 o encrypt an
 m ; reuse m .

DNSECurve: ECDH for DNS

Server knows ECDH secret key s .

Client knows ECDH secret key c ,
 server's public key $S = sG$.

Client \rightarrow server:

packet containing $cG, E_k(0, q)$

where $k = H(cS)$;

E is authenticated cipher;

q is DNS query.

Server \rightarrow client:

packet containing $E_k(1, r)$

where r is DNS response.

Client can reuse c
 across multiple qu
 but this leaks met
 Let's assume one-t

DNSCurve: ECDH for DNS

Server knows ECDH secret key s .

Client knows ECDH secret key c ,
server's public key $S = sG$.

Client \rightarrow server:

packet containing $cG, E_k(0, q)$

where $k = H(cS)$;

E is authenticated cipher;

q is DNS query.

Server \rightarrow client:

packet containing $E_k(1, r)$

where r is DNS response.

Client can reuse c
across multiple queries,
but this leaks metadata.
Let's assume one-time c .

DNSCurve: ECDH for DNS

Server knows ECDH secret key s .

Client knows ECDH secret key c ,
server's public key $S = sG$.

Client \rightarrow server:

packet containing $cG, E_k(0, q)$

where $k = H(cS)$;

E is authenticated cipher;

q is DNS query.

Server \rightarrow client:

packet containing $E_k(1, r)$

where r is DNS response.

Client can reuse c
across multiple queries,
but this leaks metadata.
Let's assume one-time c .

DNSECurve: ECDH for DNS

Server knows ECDH secret key s .

Client knows ECDH secret key c ,
server's public key $S = sG$.

Client \rightarrow server:

packet containing $cG, E_k(0, q)$

where $k = H(cS)$;

E is authenticated cipher;

q is DNS query.

Server \rightarrow client:

packet containing $E_k(1, r)$

where r is DNS response.

Client can reuse c
across multiple queries,
but this leaks metadata.
Let's assume one-time c .

KEM+AE view:

Client is sending $k = H(cS)$
encapsulated as cG .

This is an "ECDH KEM".

DNSCurve: ECDH for DNS

Server knows ECDH secret key s .

Client knows ECDH secret key c ,
server's public key $S = sG$.

Client \rightarrow server:

packet containing $cG, E_k(0, q)$

where $k = H(cS)$;

E is authenticated cipher;

q is DNS query.

Server \rightarrow client:

packet containing $E_k(1, r)$

where r is DNS response.

Client can reuse c
across multiple queries,
but this leaks metadata.
Let's assume one-time c .

KEM+AE view:

Client is sending $k = H(cS)$
encapsulated as cG .

This is an "ECDH KEM".

Client then uses k
to authenticate+encrypt.

Server also uses k
to authenticate+encrypt.

View: ECDH for DNS

knows ECDH secret key s .

knows ECDH secret key c ,
public key $S = sG$.

→ server:

containing $cG, E_k(0, q)$

$= H(cS)$;

authenticated cipher;

S query.

→ client:

containing $E_k(1, r)$

is DNS response.

Post-qua

“McElie

Client se

encapsul

Random

random

public ke

Client can reuse c

across multiple queries,

but this leaks metadata.

Let's assume one-time c .

KEM+AE view:

Client is sending $k = H(cS)$

encapsulated as cG .

This is an “ECDH KEM”.

Client then uses k

to authenticate+encrypt.

Server also uses k

to authenticate+encrypt.

Diffie-Hellman for DNS

Diffie-Hellman secret key s .

Diffie-Hellman secret key c ,
 $S = sG$.

$cG, E_k(0, q)$

cipher;

$E_k(1, r)$

response.

Client can reuse c
 across multiple queries,
 but this leaks metadata.
 Let's assume one-time c .

KEM+AE view:

Client is sending $k = H(cS)$
 encapsulated as cG .

This is an "ECDH KEM".

Client then uses k
 to authenticate+encrypt.

Server also uses k
 to authenticate+encrypt.

Post-quantum enc

"McEliece KEM":
 Client sends $k = H$
 encapsulated as S

Random $c \in \mathbf{F}_2^{5413}$
 random small $e \in$
 public key $S \in \mathbf{F}_2^{69}$

Client can reuse c
 across multiple queries,
 but this leaks metadata.
 Let's assume one-time c .

KEM+AE view:

Client is sending $k = H(cS)$
 encapsulated as cG .

This is an "ECDH KEM".

Client then uses k
 to authenticate+encrypt.

Server also uses k
 to authenticate+encrypt.

Post-quantum encrypted DM

"McEliece KEM":

Client sends $k = H(c, e, Sc)$
 encapsulated as $Sc + e$.

Random $c \in \mathbf{F}_2^{5413}$;

random small $e \in \mathbf{F}_2^{6960}$;

public key $S \in \mathbf{F}_2^{6960 \times 5413}$.

Client can reuse c
 across multiple queries,
 but this leaks metadata.
 Let's assume one-time c .

KEM+AE view:

Client is sending $k = H(cS)$
 encapsulated as cG .

This is an "ECDH KEM".

Client then uses k
 to authenticate+encrypt.

Server also uses k
 to authenticate+encrypt.

Post-quantum encrypted DNS

"McEliece KEM":

Client sends $k = H(c, e, Sc + e)$
 encapsulated as $Sc + e$.

Random $c \in \mathbf{F}_2^{5413}$;

random small $e \in \mathbf{F}_2^{6960}$;

public key $S \in \mathbf{F}_2^{6960 \times 5413}$.

Client can reuse c
 across multiple queries,
 but this leaks metadata.
 Let's assume one-time c .

KEM+AE view:

Client is sending $k = H(cS)$
 encapsulated as cG .

This is an "ECDH KEM".

Client then uses k
 to authenticate+encrypt.

Server also uses k
 to authenticate+encrypt.

Post-quantum encrypted DNS

"McEliece KEM":

Client sends $k = H(c, e, Sc + e)$
 encapsulated as $Sc + e$.

Random $c \in \mathbf{F}_2^{5413}$;

random small $e \in \mathbf{F}_2^{6960}$;

public key $S \in \mathbf{F}_2^{6960 \times 5413}$.

S has secret Goppa structure
 allowing server to decrypt.

Client can reuse c
across multiple queries,
but this leaks metadata.
Let's assume one-time c .

KEM+AE view:

Client is sending $k = H(cS)$
encapsulated as cG .

This is an "ECDH KEM".

Client then uses k
to authenticate+encrypt.

Server also uses k
to authenticate+encrypt.

Post-quantum encrypted DNS

"McEliece KEM":

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random $c \in \mathbf{F}_2^{5413}$;

random small $e \in \mathbf{F}_2^{6960}$;

public key $S \in \mathbf{F}_2^{6960 \times 5413}$.

S has secret Goppa structure
allowing server to decrypt.

"Niederreiter KEM", smaller:

Client sends $k = H(e, S'e)$

encapsulated as $S'e \in \mathbf{F}_2^{1547}$.

can reuse c
 multiple queries,
 leaks metadata.
 assume one-time c .

AE view:

sending $k = H(cS)$
 encapsulated as cG .

an "ECDH KEM".

then uses k
 authenticate+encrypt.

also uses k
 authenticate+encrypt.

Post-quantum encrypted DNS

"McEliece KEM":

Client sends $k = H(c, e, Sc + e)$
 encapsulated as $Sc + e$.

Random $c \in \mathbf{F}_2^{5413}$;

random small $e \in \mathbf{F}_2^{6960}$;

public key $S \in \mathbf{F}_2^{6960 \times 5413}$.

S has secret Goppa structure
 allowing server to decrypt.

"Niederreiter KEM", smaller:

Client sends $k = H(e, S'e)$
 encapsulated as $S'e \in \mathbf{F}_2^{1547}$.

"NTRU
 obviously
 Client se
 encapsul

Post-quantum encrypted DNS

“McEliece KEM”:

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random $c \in \mathbf{F}_2^{5413}$;

random small $e \in \mathbf{F}_2^{6960}$;

public key $S \in \mathbf{F}_2^{6960 \times 5413}$.

S has secret Goppa structure
allowing server to decrypt.

“Niederreiter KEM”, smaller:

Client sends $k = H(e, S'e)$

encapsulated as $S'e \in \mathbf{F}_2^{1547}$.

“NTRU KEM”,
obviously totally u
Client sends $k = H$
encapsulated as S

Post-quantum encrypted DNS

“McEliece KEM” :

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random $c \in \mathbf{F}_2^{5413}$;

random small $e \in \mathbf{F}_2^{6960}$;

public key $S \in \mathbf{F}_2^{6960 \times 5413}$.

S has secret Goppa structure
allowing server to decrypt.

“Niederreiter KEM”, smaller:

Client sends $k = H(e, S'e)$
encapsulated as $S'e \in \mathbf{F}_2^{1547}$.

“NTRU KEM” ,

obviously totally unrelated:
Client sends $k = H(c, e, Sc)$
encapsulated as $Sc + e$.

Post-quantum encrypted DNS

“McEliece KEM”:

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random $c \in \mathbf{F}_2^{5413}$;

random small $e \in \mathbf{F}_2^{6960}$;

public key $S \in \mathbf{F}_2^{6960 \times 5413}$.

S has secret Goppa structure
allowing server to decrypt.

“Niederreiter KEM”, smaller:

Client sends $k = H(e, S'e)$
encapsulated as $S'e \in \mathbf{F}_2^{1547}$.

“NTRU KEM”,

obviously totally unrelated:

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Post-quantum encrypted DNS

“McEliece KEM”:

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random $c \in \mathbf{F}_2^{5413}$;

random small $e \in \mathbf{F}_2^{6960}$;

public key $S \in \mathbf{F}_2^{6960 \times 5413}$.

S has secret Goppa structure
allowing server to decrypt.

“Niederreiter KEM”, smaller:

Client sends $k = H(e, S'e)$
encapsulated as $S'e \in \mathbf{F}_2^{1547}$.

“NTRU KEM”,

obviously totally unrelated:

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random small

$c, e \in (\mathbf{Z}/q)[x]/(x^n - 1)$;

public key $S \in (\mathbf{Z}/q)[x]/(x^n - 1)$.

Secretly $S = 3s/t$; small s, t .

Server recovers $3sc + te$,

then $te \pmod 3$, then e , then c .

Post-quantum encrypted DNS

“McEliece KEM”:

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random $c \in \mathbf{F}_2^{5413}$;

random small $e \in \mathbf{F}_2^{6960}$;

public key $S \in \mathbf{F}_2^{6960 \times 5413}$.

S has secret Goppa structure
allowing server to decrypt.

“Niederreiter KEM”, smaller:

Client sends $k = H(e, S'e)$
encapsulated as $S'e \in \mathbf{F}_2^{1547}$.

“NTRU KEM”,

obviously totally unrelated:

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random small

$c, e \in (\mathbf{Z}/q)[x]/(x^n - 1)$;

public key $S \in (\mathbf{Z}/q)[x]/(x^n - 1)$.

Secretly $S = 3s/t$; small s, t .

Server recovers $3sc + te$,

then $te \bmod 3$, then e , then c .

Can imitate Niederreiter in the
NTRU context: e.g. “Ring-LWR”.

Quantum encrypted DNS

“NTRU KEM”:

Client sends $k = H(c, e, Sc + e)$

encapsulated as $Sc + e$.

Client chooses $c \in \mathbf{F}_2^{5413}$;

chooses small $e \in \mathbf{F}_2^{6960}$;

uses public key $S \in \mathbf{F}_2^{6960 \times 5413}$.

Server uses secret Goppa structure

to help server to decrypt.

“Niederreiter KEM”, smaller:

Client sends $k = H(e, S'e)$

encapsulated as $S'e \in \mathbf{F}_2^{1547}$.

“NTRU KEM”,

obviously totally unrelated:

Client sends $k = H(c, e, Sc + e)$

encapsulated as $Sc + e$.

Random small

$c, e \in (\mathbf{Z}/q)[x]/(x^n - 1)$;

public key $S \in (\mathbf{Z}/q)[x]/(x^n - 1)$.

Secretly $S = 3s/t$; small s, t .

Server recovers $3sc + te$,

then $te \pmod 3$, then e , then c .

Can imitate Niederreiter in the

NTRU context: e.g. “Ring-LWR”.

Client —

packet c

(Combin

Server —

packet c

Encrypted DNS

$$H(c, e, Sc + e)$$
$$c + e.$$

$$3;$$
$$\mathbf{F}_2^{6960};$$
$$960 \times 5413.$$

a structure
decrypt.

"", smaller:

$$H(e, S'e)$$
$$e \in \mathbf{F}_2^{1547}.$$

26

"NTRU KEM",

obviously totally unrelated:

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random small

$$c, e \in (\mathbf{Z}/q)[x]/(x^n - 1);$$
$$\text{public key } S \in (\mathbf{Z}/q)[x]/(x^n - 1).$$

Secretly $S = 3s/t$; small s, t .

Server recovers $3sc + te$,
then $te \bmod 3$, then e , then c .

Can imitate Niederreiter in the
NTRU context: e.g. "Ring-LWR".

27

Client \rightarrow server:
packet containing
(Combine with EC)

Server \rightarrow client:
packet containing

“NTRU KEM”,

obviously totally unrelated:

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random small

$c, e \in (\mathbf{Z}/q)[x]/(x^n - 1)$;

public key $S \in (\mathbf{Z}/q)[x]/(x^n - 1)$.

Secretly $S = 3s/t$; small s, t .

Server recovers $3sc + te$,

then $te \bmod 3$, then e , then c .

Can imitate Niederreiter in the

NTRU context: e.g. “Ring-LWR”.

Client \rightarrow server:

packet containing $Sc + e, E_k$
(Combine with ECDH KEM)

Server \rightarrow client:

packet containing $E_k(1, r)$.

“NTRU KEM”,

obviously totally unrelated:

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random small

$c, e \in (\mathbf{Z}/q)[x]/(x^n - 1)$;

public key $S \in (\mathbf{Z}/q)[x]/(x^n - 1)$.

Secretly $S = 3s/t$; small s, t .

Server recovers $3sc + te$,

then $te \bmod 3$, then e , then c .

Can imitate Niederreiter in the
NTRU context: e.g. “Ring-LWR”.

Client \rightarrow server:

packet containing $Sc + e, E_k(0, q)$.
(Combine with ECDH KEM.)

Server \rightarrow client:

packet containing $E_k(1, r)$.

“NTRU KEM”,

obviously totally unrelated:

Client sends $k = H(c, e, Sc + e)$
encapsulated as $Sc + e$.

Random small

$c, e \in (\mathbf{Z}/q)[x]/(x^n - 1)$;
public key $S \in (\mathbf{Z}/q)[x]/(x^n - 1)$.

Secretly $S = 3s/t$; small s, t .

Server recovers $3sc + te$,
then $te \bmod 3$, then e , then c .

Can imitate Niederreiter in the
NTRU context: e.g. “Ring-LWR”.

Client \rightarrow server:

packet containing $Sc + e, E_k(0, q)$.
(Combine with ECDH KEM.)

Server \rightarrow client:

packet containing $E_k(1, r)$.

r states a server address
and the server’s public key.

What if the key is too long
to fit into a single packet?

One simple answer:

Client separately requests
each block of public key.

Can do many requests in parallel.

"KEM",

are totally unrelated:

depends $k = H(c, e, Sc + e)$

related as $Sc + e$.

small

$(\mathbf{Z}/q)[x]/(x^n - 1)$;

every $S \in (\mathbf{Z}/q)[x]/(x^n - 1)$.

$S = 3s/t$; small s, t .

recovers $3sc + te$,

mod 3, then e , then c .

quote Niederreiter in the

context: e.g. "Ring-LWR".

Client \rightarrow server:

packet containing $Sc + e, E_k(0, q)$.

(Combine with ECDH KEM.)

Server \rightarrow client:

packet containing $E_k(1, r)$.

r states a server address

and the server's public key.

What if the key is too long

to fit into a single packet?

One simple answer:

Client separately requests

each block of public key.

Can do many requests in parallel.

Confidentiality

Attacker

can't de

Integrity

Server n

but E_k i

Attacker

but can't

Attacker

Availabili

Client di

continue

eventual

unrelated:

$$H(c, e, Sc + e)$$

$$c + e.$$

$$(x^n - 1);$$

$$/q)[x]/(x^n - 1).$$

; small s, t .

$$c + te,$$

then c .

reiter in the

g. "Ring-LWR".

Client \rightarrow server:

packet containing $Sc + e, E_k(0, q)$.

(Combine with ECDH KEM.)

Server \rightarrow client:

packet containing $E_k(1, r)$.

r states a server address

and the server's public key.

What if the key is too long

to fit into a single packet?

One simple answer:

Client separately requests

each block of public key.

Can do many requests in parallel.

Confidentiality:

Attacker can't guess

can't decrypt $E_k(0, q)$

Integrity:

Server never signs

but E_k includes a

Attacker can send

but can't forge q

Attacker *can* replace

Availability:

Client discards forged

continues waiting

eventually retransmits

Client \rightarrow server:

packet containing $Sc + e, E_k(0, q)$.

(Combine with ECDH KEM.)

Server \rightarrow client:

packet containing $E_k(1, r)$.

r states a server address
and the server's public key.

What if the key is too long
to fit into a single packet?

One simple answer:

Client separately requests
each block of public key.

Can do many requests in parallel.

Confidentiality:

Attacker can't guess k ,
can't decrypt $E_k(0, q), E_k(1, r)$.

Integrity:

Server never signs anything,

but E_k includes authentication.

Attacker can send new queries,

but can't forge q or r .

Attacker *can* replay request.

Availability:

Client discards forgery,

continues waiting for reply,

eventually retransmits request.

Client \rightarrow server:

packet containing $Sc + e, E_k(0, q)$.
(Combine with ECDH KEM.)

Server \rightarrow client:

packet containing $E_k(1, r)$.

r states a server address
and the server's public key.

What if the key is too long
to fit into a single packet?

One simple answer:

Client separately requests
each block of public key.

Can do many requests in parallel.

Confidentiality:

Attacker can't guess k ,
can't decrypt $E_k(0, q), E_k(1, r)$.

Integrity:

Server never signs anything,
but E_k includes authentication.

Attacker can send new queries
but can't forge q or r .

Attacker *can* replay request.

Availability:

Client discards forgery,
continues waiting for reply,
eventually retransmits request.

→ server:

containing $Sc + e, E_k(0, q)$.
(Use with ECDH KEM.)

→ client:

containing $E_k(1, r)$.

a server address

server's public key.

the key is too long
to a single packet?

simple answer:

separately requests

block of public key.

many requests in parallel.

Confidentiality:

Attacker can't guess k ,
can't decrypt $E_k(0, q), E_k(1, r)$.

Integrity:

Server never signs anything,
but E_k includes authentication.

Attacker can send new queries
but can't forge q or r .

Attacker *can* replay request.

Availability:

Client discards forgery,
continues waiting for reply,
eventually retransmits request.

Cookies

What if
into same

Client se
containing

Server se

cookie

k) encry

Server c

Client se

Server re

Server se

$Sc + e, E_k(0, q)$.
(CDH KEM.)

$E_k(1, r)$.

address

public key.

too long

packet?

r:

requests

public key.

requests in parallel.

Confidentiality:

Attacker can't guess k ,
can't decrypt $E_k(0, q), E_k(1, r)$.

Integrity:

Server never signs anything,
but E_k includes authentication.

Attacker can send new queries
but can't forge q or r .

Attacker *can* replay request.

Availability:

Client discards forgery,
continues waiting for reply,
eventually retransmits request.

Cookies

What if $E_k(0, q)$ and
 $E_k(1, r)$ go into same packet?

Client sends short
packet containing a **cookie**.

Server sends $E_k(1, r)$ and
cookie r' : server's
secret key.

$E_k(0, q)$ encrypted from
server's secret key.
Server can now forward
request.

Client sends packet
containing $E_k(0, q)$ and
server's secret key.

Server sends $E_k(0, q)$ and
server's secret key.

Confidentiality:

Attacker can't guess k ,
can't decrypt $E_k(0, q), E_k(1, r)$.

Integrity:

Server never signs anything,
but E_k includes authentication.

Attacker can send new queries
but can't forge q or r .

Attacker *can* replay request.

Availability:

Client discards forgery,
continues waiting for reply,
eventually retransmits request.

Cookies

What if $E_k(0, q)$ doesn't fit
into same packet as $Sc + e$?

Client sends short $E_k(0, q')$
containing a **cookie request**.

Server sends $E_k(1, r')$ containing
cookie r' : server state (including
 k) encrypted from server to client.
Server can now forget state.

Client sends packet $r', E_k(2, q)$.
Server recovers state, decrypts $E_k(2, q)$.

Server sends $E_k(3, r)$.

Confidentiality:

Attacker can't guess k ,
can't decrypt $E_k(0, q)$, $E_k(1, r)$.

Integrity:

Server never signs anything,
but E_k includes authentication.

Attacker can send new queries
but can't forge q or r .

Attacker *can* replay request.

Availability:

Client discards forgery,
continues waiting for reply,
eventually retransmits request.

Cookies

What if $E_k(0, q)$ doesn't fit
into same packet as $Sc + e$?

Client sends short $E_k(0, q')$
containing a **cookie request** q' .

Server sends $E_k(1, r')$ containing
cookie r' : server state (including
 k) encrypted from server to itself.
Server can now forget state.

Client sends packet $r', E_k(2, q)$.
Server recovers state, decrypts.

Server sends $E_k(3, r)$.

ntiality:

r can't guess k ,

crypt $E_k(0, q), E_k(1, r)$.

:

ever signs anything,

ncludes authentication.

r can send new queries

t forge q or r .

r *can* replay request.

lity:

discards forgery,

es waiting for reply,

ly retransmits request.

Cookies

What if $E_k(0, q)$ doesn't fit into same packet as $Sc + e$?

Client sends short $E_k(0, q')$ containing a **cookie request** q' .

Server sends $E_k(1, r')$ containing **cookie** r' : server state (including k) encrypted from server to itself. Server can now forget state.

Client sends packet $r', E_k(2, q)$. Server recovers state, decrypts.

Server sends $E_k(3, r)$.

Client au

Same st
for prote

$C \rightarrow S$,

isn't spe

many pa

Cookies

What if $E_k(0, q)$ doesn't fit into same packet as $Sc + e$?

Client sends short $E_k(0, q')$ containing a **cookie request** q' .

Server sends $E_k(1, r')$ containing **cookie** r' : server state (including k) encrypted from server to itself. Server can now forget state.

Client sends packet $r', E_k(2, q)$. Server recovers state, decrypts.

Server sends $E_k(3, r)$.

Client authentication

Same strategy works for protecting confidentiality. $C \rightarrow S, S \rightarrow C$ data isn't special; reuse many packets each

Cookies

What if $E_k(0, q)$ doesn't fit into same packet as $Sc + e$?

Client sends short $E_k(0, q')$ containing a **cookie request** q' .

Server sends $E_k(1, r')$ containing **cookie** r' : server state (including k) encrypted from server to itself. Server can now forget state.

Client sends packet $r', E_k(2, q)$. Server recovers state, decrypts.

Server sends $E_k(3, r)$.

Client authentication

Same strategy works for protecting connections. $C \rightarrow S, S \rightarrow C$ data flow isn't special; reuse k for many packets each direction

Cookies

What if $E_k(0, q)$ doesn't fit into same packet as $Sc + e$?

Client sends short $E_k(0, q')$ containing a **cookie request** q' .

Server sends $E_k(1, r')$ containing **cookie** r' : server state (including k) encrypted from server to itself.

Server can now forget state.

Client sends packet $r', E_k(2, q)$.

Server recovers state, decrypts.

Server sends $E_k(3, r)$.

Client authentication

Same strategy works for protecting connections.

$C \rightarrow S, S \rightarrow C$ data flow isn't special; reuse k for many packets each direction.

Cookies

What if $E_k(0, q)$ doesn't fit into same packet as $Sc + e$?

Client sends short $E_k(0, q')$ containing a **cookie request** q' .

Server sends $E_k(1, r')$ containing **cookie** r' : server state (including k) encrypted from server to itself. Server can now forget state.

Client sends packet $r', E_k(2, q)$. Server recovers state, decrypts.

Server sends $E_k(3, r)$.

Client authentication

Same strategy works for protecting connections.

$C \rightarrow S, S \rightarrow C$ data flow isn't special; reuse k for many packets each direction.

Another TCP availability problem: server allocates buffers for each connection; runs out of memory.

Cookies

What if $E_k(0, q)$ doesn't fit into same packet as $Sc + e$?

Client sends short $E_k(0, q')$ containing a **cookie request** q' .

Server sends $E_k(1, r')$ containing **cookie** r' : server state (including k) encrypted from server to itself. Server can now forget state.

Client sends packet $r', E_k(2, q)$. Server recovers state, decrypts.

Server sends $E_k(3, r)$.

Client authentication

Same strategy works for protecting connections.

$C \rightarrow S, S \rightarrow C$ data flow isn't special; reuse k for many packets each direction.

Another TCP availability problem: server allocates buffers for each connection; runs out of memory.

Semi-solution: Allocate buffers only after client sends r' .

Cookies

What if $E_k(0, q)$ doesn't fit into same packet as $Sc + e$?

Client sends short $E_k(0, q')$ containing a **cookie request** q' .

Server sends $E_k(1, r')$ containing **cookie** r' : server state (including k) encrypted from server to itself. Server can now forget state.

Client sends packet $r', E_k(2, q)$. Server recovers state, decrypts.

Server sends $E_k(3, r)$.

Client authentication

Same strategy works for protecting connections.

$C \rightarrow S, S \rightarrow C$ data flow isn't special; reuse k for many packets each direction.

Another TCP availability problem: server allocates buffers for each connection; runs out of memory.

Semi-solution: Allocate buffers only after client sends r' .

Solution 1: Hashcash from client.

$E_k(0, q)$ doesn't fit
the packet as $Sc + e$?

sends short $E_k(0, q')$
using a **cookie request** q' .

sends $E_k(1, r')$ containing
 r' : server state (including
encrypted from server to itself.
can now forget state.

sends packet $r', E_k(2, q)$.
recovers state, decrypts.

sends $E_k(3, r)$.

Client authentication

Same strategy works
for protecting connections.

$C \rightarrow S, S \rightarrow C$ data flow
isn't special; reuse k for
many packets each direction.

Another TCP availability problem:
server allocates buffers for each
connection; runs out of memory.

Semi-solution: Allocate buffers
only after client sends r' .

Solution 1: Hashcash from client.

Solution
to avoid
Imitate

doesn't fit
as $Sc + e$?

$E_k(0, q')$
ie request q' .

(r') containing
state (including
server to itself.
target state.

et $r', E_k(2, q)$.
ate, decrypts.

(r) .

Client authentication

Same strategy works
for protecting connections.

$C \rightarrow S, S \rightarrow C$ data flow
isn't special; reuse k for
many packets each direction.

Another TCP availability problem:
server allocates buffers for each
connection; runs out of memory.

Semi-solution: Allocate buffers
only after client sends r' .

Solution 1: Hashcash from client.

Solution 2: Redo
to avoid state on s
Imitate NFS, not

Client authentication

Same strategy works
for protecting connections.

$C \rightarrow S, S \rightarrow C$ data flow
isn't special; reuse k for
many packets each direction.

Another TCP availability problem:
server allocates buffers for each
connection; runs out of memory.

Semi-solution: Allocate buffers
only after client sends r' .

Solution 1: Hashcash from client.

Solution 2: Redo protocols
to avoid state on server.

Imitate NFS, not HTTP.

Client authentication

Same strategy works for protecting connections. $C \rightarrow S, S \rightarrow C$ data flow isn't special; reuse k for many packets each direction.

Another TCP availability problem: server allocates buffers for each connection; runs out of memory.

Semi-solution: Allocate buffers only after client sends r' .

Solution 1: Hashcash from client.

Solution 2: Redo protocols to avoid state on server.

Imitate NFS, not HTTP.

Client authentication

Same strategy works for protecting connections.

$C \rightarrow S, S \rightarrow C$ data flow isn't special; reuse k for many packets each direction.

Another TCP availability problem: server allocates buffers for each connection; runs out of memory.

Semi-solution: Allocate buffers only after client sends r' .

Solution 1: Hashcash from client.

Solution 2: Redo protocols to avoid state on server.

Imitate NFS, not HTTP.

Solution 3 for, e.g., SSH: Authenticate client.

Server can authenticate client without signatures, same way client authenticates server:

- Send to client's public key encapsulation of new key k' .
- Hash k' into shared secret.

Authentication

strategy works

rejecting connections.

$S \rightarrow C$ data flow

special; reuse k for

packets each direction.

TCP availability problem:

allocates buffers for each

connection; runs out of memory.

solution: Allocate buffers

after client sends r' .

1: Hashcash from client.

Solution 2: Redo protocols
to avoid state on server.

Imitate NFS, not HTTP.

Solution 3 for, e.g., SSH:

Authenticate client.

Server can authenticate client
without signatures, same way
client authenticates server:

- Send to client's public key
encapsulation of new key k' .
- Hash k' into shared secret.

Big keys

McEliece

for long-

Is this si

Do we n

lower-co

such as

Size of a

in Alexa

Web pag

public ke

but publ

can be r

Solution 2: Redo protocols to avoid state on server.

Imitate NFS, not HTTP.

Solution 3 for, e.g., SSH:
Authenticate client.

Server can authenticate client without signatures, same way client authenticates server:

- Send to client's public key encapsulation of new key k' .
- Hash k' into shared secret.

Big keys

McEliece public key for long-term conf

Is this size a probl

Do we need to sw

lower-confidence a

such as NTRU or

Size of average we

in Alexa Top 1000

Web page often ne

public keys for sev

but public key for

can be reused for

Solution 2: Redo protocols to avoid state on server.

Imitate NFS, not HTTP.

Solution 3 for, e.g., SSH:

Authenticate client.

Server can authenticate client without signatures, same way client authenticates server:

- Send to client's public key encapsulation of new key k' .
- Hash k' into shared secret.

Big keys

McEliece public key is 1MB for long-term confidence too

Is this size a problem?

Do we need to switch to lower-confidence approaches such as NTRU or QC-MDPC

Size of average web page in Alexa Top 1000000: 1.8M

Web page often needs public keys for several servers but public key for a server can be reused for many pages

Solution 2: Redo protocols to avoid state on server.

Imitate NFS, not HTTP.

Solution 3 for, e.g., SSH:
Authenticate client.

Server can authenticate client without signatures, same way client authenticates server:

- Send to client's public key encapsulation of new key k' .
- Hash k' into shared secret.

Big keys

McEliece public key is 1MB for long-term confidence today.

Is this size a problem?

Do we need to switch to lower-confidence approaches such as NTRU or QC-MDPC?

Size of average web page in Alexa Top 1000000: 1.8MB.

Web page often needs public keys for several servers, but public key for a server can be reused for many pages.

2: Redo protocols
state on server.

NFS, not HTTP.

3 for, e.g., SSH:
authenticate client.

can authenticate client
signatures, same way

authenticates server:

to client's public key

generation of new key k' .

k' into shared secret.

Big keys

McEliece public key is 1MB
for long-term confidence today.

Is this size a problem?

Do we need to switch to
lower-confidence approaches
such as NTRU or QC-MDPC?

Size of average web page
in Alexa Top 1000000: 1.8MB.

Web page often needs
public keys for several servers,
but public key for a server
can be reused for many pages.

Most im
on reuse
switchin
and **pro**

Rational
subsequ
doesn't

e.g. Mic
switches

Safer: n

Easier to
new key

Big keys

McEliece public key is 1MB
for long-term confidence today.

Is this size a problem?

Do we need to switch to
lower-confidence approaches
such as NTRU or QC-MDPC?

Size of average web page
in Alexa Top 1000000: 1.8MB.

Web page often needs
public keys for several servers,
but public key for a server
can be reused for many pages.

Most important line
on reuse of public
switching to new
and **promptly era**

Rationale: “forward
subsequent theft of
doesn't allow decr

e.g. Microsoft SCh
switches keys ever

Safer: new key eve

Easier to impleme
new key every con

Big keys

McEliece public key is 1MB
for long-term confidence today.

Is this size a problem?

Do we need to switch to
lower-confidence approaches
such as NTRU or QC-MDPC?

Size of average web page
in Alexa Top 1000000: 1.8MB.

Web page often needs
public keys for several servers,
but public key for a server
can be reused for many pages.

Most important limitation
on reuse of public keys:
switching to new keys
and **promptly erasing old k**

Rationale: “forward secrecy”
subsequent theft of computer
doesn’t allow decryption.

e.g. Microsoft SChannel
switches keys every two hours

Safer: new key every minute

Easier to implement:
new key every connection.

Big keys

McEliece public key is 1MB
for long-term confidence today.

Is this size a problem?

Do we need to switch to
lower-confidence approaches
such as NTRU or QC-MDPC?

Size of average web page
in Alexa Top 1000000: 1.8MB.

Web page often needs
public keys for several servers,
but public key for a server
can be reused for many pages.

Most important limitation
on reuse of public keys:
switching to new keys
and **promptly erasing old keys.**

Rationale: “forward secrecy” —
subsequent theft of computer
doesn't allow decryption.

e.g. Microsoft SChannel
switches keys every two hours.

Safer: new key every minute.

Easier to implement:
new key every connection.

the public key is 1MB
 -term confidence today.

ze a problem?

eed to switch to
 nfidence approaches
 NTRU or QC-MDPC?

average web page
 Top 1000000: 1.8MB.

ge often needs
 eys for several servers,
 ic key for a server
 eused for many pages.

Most important limitation
 on reuse of public keys:
 switching to new keys
 and **promptly erasing old keys.**

Rationale: “forward secrecy” —
 subsequent theft of computer
 doesn't allow decryption.

e.g. Microsoft SChannel
 switches keys every two hours.

Safer: new key every minute.

Easier to implement:
 new key every connection.

What is
 a new ke
 If server
 key gen,
 client en
 server de

Most important limitation
on reuse of public keys:
switching to new keys
and **promptly erasing old keys.**

Rationale: “forward secrecy” —
subsequent theft of computer
doesn't allow decryption.

e.g. Microsoft SChannel
switches keys every two hours.

Safer: new key every minute.

Easier to implement:
new key every connection.

What is the perform
a new key every m
If server makes ne
key gen, ≤ 1 per m
client encrypts to
server decrypts.

Most important limitation
on reuse of public keys:
switching to new keys
and **promptly erasing old keys.**

Rationale: “forward secrecy” —
subsequent theft of computer
doesn't allow decryption.

e.g. Microsoft SChannel
switches keys every two hours.

Safer: new key every minute.

Easier to implement:
new key every connection.

What is the performance of
a new key every minute?

If server makes new key:
key gen, ≤ 1 per minute;
client encrypts to new key;
server decrypts.

Most important limitation
on reuse of public keys:
switching to new keys
and **promptly erasing old keys.**

Rationale: “forward secrecy” —
subsequent theft of computer
doesn't allow decryption.

e.g. Microsoft SChannel
switches keys every two hours.

Safer: new key every minute.

Easier to implement:
new key every connection.

What is the performance of
a new key every minute?

If server makes new key:
key gen, ≤ 1 per minute;
client encrypts to new key;
server decrypts.

Most important limitation
on reuse of public keys:
switching to new keys
and **promptly erasing old keys.**

Rationale: “forward secrecy” —
subsequent theft of computer
doesn't allow decryption.

e.g. Microsoft SChannel
switches keys every two hours.

Safer: new key every minute.

Easier to implement:
new key every connection.

What is the performance of
a new key every minute?

If server makes new key:
key gen, ≤ 1 per minute;
client encrypts to new key;
server decrypts.

If client makes new key:
client has key-gen cost;
server has encryption cost;
client has decryption cost.

Either way:
one key transmission for each
active client-server pair.

important limitation
of public keys:
going to new keys
promptly erasing old keys.

Example: “forward secrecy” —
even if an attacker steals the
private key, they cannot
decrypt old messages.

Microsoft SChannel
generates new keys every two hours.

generates new key every minute.

How to implement:
generate new key every
connection.

What is the performance of
generating a new key every minute?

If server makes new key:
key gen, ≤ 1 per minute;
client encrypts to new key;
server decrypts.

If client makes new key:
client has key-gen cost;
server has encryption cost;
client has decryption cost.

Either way:
one key transmission for each
active client-server pair.

How does
forward secrecy
work without
perfect forward secrecy?

mitation

keys:

keys

using old keys.

rd secrecy" —

of computer

ryption.

channel

y two hours.

ery minute.

nt:

nection.

What is the performance of
a new key every minute?

If server makes new key:

key gen, ≤ 1 per minute;

client encrypts to new key;

server decrypts.

If client makes new key:

client has key-gen cost;

server has encryption cost;

client has decryption cost.

Either way:

one key transmission for each

active client-server pair.

How does a *stateless*

encrypt to a new c

without storing th

What is the performance of a new key every minute?

If server makes new key:
key gen, ≤ 1 per minute;
client encrypts to new key;
server decrypts.

If client makes new key:
client has key-gen cost;
server has encryption cost;
client has decryption cost.

Either way:
one key transmission for each active client-server pair.

How does a *stateless* server encrypt to a new client key without storing the key?

What is the performance of a new key every minute?

If server makes new key:

key gen, ≤ 1 per minute;

client encrypts to new key;

server decrypts.

If client makes new key:

client has key-gen cost;

server has encryption cost;

client has decryption cost.

Either way:

one key transmission for each active client-server pair.

How does a *stateless* server encrypt to a new client key without storing the key?

What is the performance of a new key every minute?

If server makes new key:
key gen, ≤ 1 per minute;
client encrypts to new key;
server decrypts.

If client makes new key:
client has key-gen cost;
server has encryption cost;
client has decryption cost.

Either way:
one key transmission for each active client-server pair.

How does a *stateless* server encrypt to a new client key without storing the key?

Slice McEliece public key so that each slice of encryption produces separate small output.

Client sends slices (in parallel), receives outputs as cookies, sends cookies (in parallel).

Server combines cookies.

Continue up through tree.

Server generates randomness as secret function of key hash.

Statelessly verifies key hash.