From: Blumenthal, Uri - 0553 - MITLL <uri@ll.mit.edu> via pqc-forum@list.nist.gov
To: Patrick Longa <plonga@microsoft.com>, pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] New results for SIKE
Date: Tuesday, March 22, 2022 02:56:35 PM ET
Attachments: smime.p7m

It is great to see speedups of SIDH and SIKE.

But, if I'm correct, it accelerated by less than 30%? Given the performance difference between SIKE and other finalists is several orders of magnitude, I'm not sure it would enable use cases that currently cannot use SIKE because of its slow performance.

Thanks

P.S. P434 assembly code has a problem on Mac. I submitted an issue with a fix on GitHub.

--

V/R,

Uri

There are two ways to design a system. One is to make it so simple there are obviously no deficiencies.

The other is to make it so complex there are no obvious deficiencies.

- C. A. R. Hoare

> **From:** 'Patrick Longa' via pqc-forum
> **Reply-To:** Patrick Longa
> **Date:** Tuesday, March 22, 2022 at 13:52
> **To:** "pqc-forum@list.nist.gov"
> **Subject:** [pqc-forum] New results for SIKE
>
> Dear NIST and pqc-forum,
>
> I'm happy to share new results that show a significant speed up in the computation of SIKE.

The paper https://eprint.iacr.org/2022/367 presents algorithms that generalize interleaved Montgomery multiplications to the computation of extension field multiplications. This, in a more general sense, has implications for any scheme whose underlying arithmetic runs over an extension field GF(p^k) of large prime characteristic (e.g., it directly impacts all the recently proposed supersingular isogeny-based protocols like B-SIDH and SQISign, and also bilinear pairings).

Specifically for SIKE, our software implementation is sped up by approximately 1.3x (p434, level 1), 1.2x (p503, level 2) and 1.1x (p610, level 3). The alternative parameter p377 proposed in https://eprint.iacr.org/2020/1457 is also sped up by approximately 1.3x.

The new algorithms are already integrated to the SIDH library: https://github.com/microsoft/PQCrypto-SIDH

All these results correspond to a software implementation on a standard x64 computer. However, the algorithms are generic and are expected to benefit implementations on other platforms, including hardware, constrained devices, vectorized implementations, etc. (see the discussion in Section 6 of https://eprint.iacr.org/2022/367).

Kind regards,

Patrick

--

--

To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/35C98171-D30A-41A8-BCC3-5AE7F094960E%40ll.mit.edu.

From: Patrick Longa <plonga@microsoft.com> via pqc-forum <pqc-forum@list.nist.gov>
To: Blumenthal, Uri - 0553 - MITLL <uri@ll.mit.edu>, pqc-forum@list.nist.gov
Subject: RE: [pqc-forum] New results for SIKE
Date: Tuesday, March 22, 2022 03:42:51 PM ET

The big picture is that SIDH and SIKE are also orders of magnitude faster than what they used to be, so they have experienced a more dramatic improvement over only a few years (and still continue to surprise us with new speedups). It's true that isogeny-based crypto is not faster (and might never be) than structured lattices, but it provides other benefits (maybe the most important point is that we really need schemes that are based on completely different hardness problems).

I'll let others assess if the improvements are game changing for real-world applications. The proposed algorithms are generic and will potentially affect results for other platforms. As is almost always the case, new ideas open the possibility for even further improvements.

> P.S. P434 assembly code has a problem on Mac. I submitted an issue with a fix on GitHub.

Yeah, that might be the case (I don't have a Mac at hand to test the code).

Btw I don't see the submitted issue online yet.

Best,

Patrick

From: pqc-forum@list.nist.gov <pqc-forum@list.nist.gov> **On Behalf Of** Blumenthal, Uri - 0553 - MITLL
Sent: Tuesday, March 22, 2022 11:56 AM
To: Patrick Longa <plonga@microsoft.com>; pqc-forum@list.nist.gov
Subject: [EXTERNAL] Re: [pqc-forum] New results for SIKE

It is great to see speedups of SIDH and SIKE.

But, if I'm correct, it accelerated by less than 30%? Given the performance difference between SIKE and other finalists is several orders of magnitude, I'm not sure it would enable use cases that currently cannot use SIKE because of its slow performance.

Thanks

P.S. P434 assembly code has a problem on Mac. I submitted an issue with a fix on GitHub.

--

V/R,

Uri

There are two ways to design a system. One is to make it so simple there are obviously no deficiencies.

The other is to make it so complex there are no obvious deficiencies.

- C. A. R. Hoare

From: 'Patrick Longa' via pqc-forum <pqc-forum@list.nist.gov>
Reply-To: Patrick Longa <plonga@microsoft.com>
Date: Tuesday, March 22, 2022 at 13:52
To: "pqc-forum@list.nist.gov" <pqc-forum@list.nist.gov>
Subject: [pqc-forum] New results for SIKE

Dear NIST and pqc-forum,

I'm happy to share new results that show a significant speed up in the computation of SIKE.

The paper https://eprint.iacr.org/2022/367 presents algorithms that generalize interleaved Montgomery multiplications to the computation of extension field multiplications. This, in a more general sense, has implications for any scheme whose underlying arithmetic runs over an extension field GF(p^k) of large prime characteristic (e.g., it directly impacts all the recently proposed supersingular isogeny-based protocols like B-SIDH and SQISign, and also bilinear pairings).

Specifically for SIKE, our software implementation is sped up by approximately 1.3x (p434, level 1), 1.2x (p503, level 2) and 1.1x (p610, level 3). The alternative parameter p377 proposed in https://eprint.iacr.org/2020/1457 is also sped up by approximately 1.3x.

The new algorithms are already integrated to the SIDH library: https://github.com/microsoft/PQCrypto-SIDH

All these results correspond to a software implementation on a standard x64 computer. However, the algorithms are generic and are expected to benefit implementations on other platforms, including hardware, constrained devices, vectorized implementations, etc. (see the discussion in Section 6 of https://eprint.iacr.org/2022/367).

Kind regards,

Patrick

--
You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/PH0PR21MB1862656D3ED138AD67E3D542DA179%40PH0PR21MB1862.namprd21.prod.outlook.com.

--
You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/35C98171-D30A-41A8-BCC3-5AE7F094960E%40ll.mit.edu.

From: Blumenthal, Uri - 0553 - MITLL <uri@ll.mit.edu> via pqc-forum@list.nist.gov
To: Patrick Longa <plonga@microsoft.com>, pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] New results for SIKE
Date: Tuesday, March 22, 2022 04:28:33 PM ET
Attachments: smime.p7m

> The big picture is that SIDH and SIKE are also orders of magnitude faster than what they used to be, so they have experienced a more dramatic improvement over only a few years (and still continue to surprise us with new speedups).

Hopefully, this trend will continue.

> It's true that isogeny-based crypto is not faster (and might never be) than structured lattices,
>
> but it provides other benefits (maybe the most important point is that we really need schemes that are based on completely different hardness problems).

Yes, absolutely. And it's great to see that this approach progresses and improves.

> I'll let others assess if the improvements are game changing for real-world applications.

For my practical applications, it doesn't need to be faster than lattices. Even twice as slow [as lattices] wouldn't be a deal-breaker. But being 500x slower…

> > P.S. P434 assembly code has a problem on Mac. I submitted an issue with a fix on GitHub.
>
> Yeah, that might be the case (I don't have a Mac at hand to test the code).
>
> Btw I don't see the submitted issue online yet.

Sorry, I got delayed actually submitting it. It's there now: https://github.com/microsoft/PQCrypto-SIDH/issues/51

Thanks

> **From:** pqc-forum@list.nist.gov **On Behalf Of** Blumenthal, Uri - 0553 - MITLL
> **Sent:** Tuesday, March 22, 2022 11:56 AM

To: Patrick Longa ; pqc-forum@list.nist.gov
Subject: [EXTERNAL] Re: [pqc-forum] New results for SIKE

It is great to see speedups of SIDH and SIKE.

But, if I'm correct, it accelerated by less than 30%? Given the performance difference between SIKE and other finalists is several orders of magnitude, I'm not sure it would enable use cases that currently cannot use SIKE because of its slow performance.

Thanks

P.S. P434 assembly code has a problem on Mac. I submitted an issue with a fix on GitHub.

--

V/R,

Uri

There are two ways to design a system. One is to make it so simple there are obviously no deficiencies.

The other is to make it so complex there are no obvious deficiencies.

- C. A. R. Hoare

From: 'Patrick Longa' via pqc-forum <pqc-forum@list.nist.gov>
Reply-To: Patrick Longa <plonga@microsoft.com>
Date: Tuesday, March 22, 2022 at 13:52
To: "pqc-forum@list.nist.gov" <pqc-forum@list.nist.gov>
Subject: [pqc-forum] New results for SIKE

Dear NIST and pqc-forum,

I'm happy to share new results that show a significant speed up in the computation of SIKE.

The paper https://eprint.iacr.org/2022/367 presents algorithms that generalize interleaved Montgomery multiplications to the computation of extension field multiplications. This, in a more general sense, has implications for any scheme whose underlying arithmetic runs over an extension field GF(p^k) of large prime

characteristic (e.g., it directly impacts all the recently proposed supersingular isogeny-based protocols like B-SIDH and SQISign, and also bilinear pairings).

Specifically for SIKE, our software implementation is sped up by approximately 1.3x (p434, level 1), 1.2x (p503, level 2) and 1.1x (p610, level 3). The alternative parameter p377 proposed in https://eprint.iacr.org/2020/1457 is also sped up by approximately 1.3x.

The new algorithms are already integrated to the SIDH library: https://github.com/microsoft/PQCrypto-SIDH

All these results correspond to a software implementation on a standard x64 computer. However, the algorithms are generic and are expected to benefit implementations on other platforms, including hardware, constrained devices, vectorized implementations, etc. (see the discussion in Section 6 of https://eprint.iacr.org/2022/367).

Kind regards,

Patrick

--
You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/PH0PR21MB1862AEFB56CB98E4FC4B608ADA179%40PH0PR21MB1862.namprd21.prod.outlook.com.

--
You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/9C19D847-2E8F-4451-B41E-F05D8D4B0B6A%40ll.mit.edu.

From: Scott Fluhrer (sfluhrer) <sfluhrer@cisco.com> via pqc-forum <pqc-forum@list.nist.gov>
To: Blumenthal, Uri - 0553 - MITLL <uri@ll.mit.edu>, Patrick Longa <plonga@microsoft.com>, pqc-forum@list.nist.gov
Subject: RE: [pqc-forum] New results for SIKE
Date: Tuesday, March 22, 2022 04:49:23 PM ET

From: pqc-forum@list.nist.gov <pqc-forum@list.nist.gov> **On Behalf Of** Blumenthal, Uri - 0553 - MITLL
Sent: Tuesday, March 22, 2022 4:28 PM
To: Patrick Longa <plonga@microsoft.com>; pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] New results for SIKE

> The big picture is that SIDH and SIKE are also orders of magnitude faster than what they used to be, so they have experienced a more dramatic improvement over only a few years (and still continue to surprise us with new speedups).

Hopefully, this trend will continue.

> It's true that isogeny-based crypto is not faster (and might never be) than structured lattices,
>
> but it provides other benefits (maybe the most important point is that we really need schemes that are based on completely different hardness problems).

Yes, absolutely. And it's great to see that this approach progresses and improves.

> I'll let others assess if the improvements are game changing for real-world applications.

For my practical applications, it doesn't need to be faster than lattices. Even twice as slow [as lattices] wouldn't be a deal-breaker. But being 500x slower…

Depends on the application; for some (e.g. Wireless), the extra bandwidth of lattices certainly are. For those, SIKE (even with its computational overhead) may be a viable option, and any performance improvement makes it more likely

From: Blumenthal, Uri - 0553 - MITLL <uri@ll.mit.edu> via pqc-forum@list.nist.gov
To: Scott Fluhrer (sfluhrer) <sfluhrer@cisco.com>, Patrick Longa <plonga@microsoft.com>, pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] New results for SIKE
Date: Tuesday, March 22, 2022 04:59:09 PM ET
Attachments: smime.p7m

> For my practical applications, it doesn't need to be faster than lattices. Even twice as slow [as lattices] wouldn't be a deal-breaker. But being 500x slower...

> Depends on the application; for some (e.g. Wireless), the extra bandwidth of lattices certainly are. For those, SIKE (even with its computational overhead) may be a viable option, and any performance improvement makes it more likely

An excellent point. Public key and ciphertext sizes would be about ¼ of their lattice-based counterparts – which is a definite advantage.

Still, for Wireless it matters how fast you can close the link.

From: Reza Azarderakhsh <azarderakhsh@gmail.com> via pqc-forum@list.nist.gov
To: Blumenthal, Uri - 0553 - MITLL <uri@ll.mit.edu>
CC: Scott Fluhrer (sfluhrer) <sfluhrer@cisco.com>, Patrick Longa <plonga@microsoft.com>, pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] New results for SIKE
Date: Tuesday, March 22, 2022 09:49:09 PM ET

On Tue, Mar 22, 2022 at 4:58 PM Blumenthal, Uri - 0553 - MITLL <uri@ll.mit.edu> wrote:

> For my practical applications, it doesn't need to be faster than lattices. Even twice as slow [as lattices] wouldn't be a deal-breaker. But being 500x slower...

This is inaccurate. SIKE takes 139 million cycles (E+D) on ARM Cortex M4 while the fastest structured lattice (Kyber) takes about 1.2 million cycles which is only about x100 slower. Accelerating SIKE in hardware will make it down to 2-4 million cycles or even lower if you pay more for silicon. You cannot pay to reduce the key size of Kyber. In a real world example, if both SIKE and Kyber get accelerated in hardware, SIKE will be as efficient as Kyber in terms of both speed and area for the whole protocol, e.g., Mbed TLS. This is the case for even classical cryptography such as ECC. Currently, all secure embedded IoTs are coming with accelerators for ECC. The comparison could get even more interesting if you plan to add side-channel countermeasures which will make SIKE even more efficient as it is more SCA resistant friendly adding less overhead than Kyber.

Reza Azarderakhsh

> Depends on the application; for some (e.g. Wireless), the extra bandwidth of lattices certainly are. For those, SIKE (even with its computational overhead) may be a viable option, and any performance improvement makes it more likely

An excellent point. Public key and ciphertext sizes would be about ¼ of their lattice-based counterparts – which is a definite advantage.

Still, for Wireless it matters how fast you can close the link.

> To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
> To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/7CF55252-7928-47AD-9EDC-B1CF01752B0F%40ll.mit.edu.

--

You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/CAGPziQ2MeLofRMsYjsX2Fr_EXR0-Qx1i_AkWF3U_cDHRkH2osg%40mail.gmail.com.

From: Blumenthal, Uri - 0553 - MITLL <uri@ll.mit.edu> via pqc-forum@list.nist.gov
To: Reza Azarderakhsh <azarderakhsh@gmail.com>
CC: pqc-forum@list.nist.gov
Subject: Re: [pqc-forum] New results for SIKE
Date: Wednesday, March 23, 2022 12:15:05 AM ET
Attachments: smime.p7m

> For my practical applications, it doesn't need to be faster than lattices. Even twice as slow [as lattices] wouldn't be a deal-breaker. But being 500x slower…

This is inaccurate. SIKE takes 139 million cycles (E+D) on ARM Cortex M4 while the fastest structured lattice (Kyber) takes about 1.2 million cycles which is only about x100 slower.

I compared the actual performance (timing, not the CPU cycles) on Intel CPUs. Your point is taken.

Accelerating SIKE in hardware will make it down to 2-4 million cycles or even lower if you pay more for silicon.

Hardware (2-4M cycles) – do you mean ASIC, or FPGA?

You cannot pay to reduce the key size of Kyber. In a real-world example, if both SIKE and Kyber get accelerated in hardware, SIKE will be as efficient as Kyber in terms of both speed and area for the whole protocol, e.g., Mbed TLS.

Good point, thanks.

The comparison could get even more interesting if you plan to add side-channel countermeasures which will make SIKE even more efficient as it is more SCA resistant friendly adding less overhead than Kyber.

Indeed, but I haven't looked that deep (yet), not to mention that adding SCA would be somebody else's job. 😉

Thank you!

> Depends on the application; for some (e.g. Wireless), the extra bandwidth of lattices certainly are. For those, SIKE (even with its computational overhead) may be a viable option, and any performance improvement makes it more likely
>
> An excellent point. Public key and ciphertext sizes would be about ¼ of their lattice-based counterparts – which is a definite advantage.
>
> Still, for Wireless it matters how fast you can close the link.
>
> --
> You received this message because you are subscribed to the Google Groups "pqc-forum" group.
> To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
> To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/7CF55252-7928-47AD-9EDC-B1CF01752B0F%40ll.mit.edu.

--
You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/CAGPziQ2MeLofRMsYjsX2Fr_EXR0-Qx1i_AkWF3U_cDHRkH2osg%40mail.gmail.com.

--
You received this message because you are subscribed to the Google Groups "pqc-forum" group.
To unsubscribe from this group and stop receiving emails from it, send an email to pqc-forum+unsubscribe@list.nist.gov.
To view this discussion on the web visit https://groups.google.com/a/list.nist.gov/d/msgid/pqc-forum/9752382D-BE17-4D2D-A8D4-9AEBC7AC9DF2%40ll.mit.edu.