

# A Concrete Treatment of Efficient Continuous Group Key Agreement via Multi-Recipient PKEs

Thomas Prest

PQShield

Séminaire de Cryptographie de l'Université de Rennes 1

This is joint work with:

- Keitaro Hashimoto (Tokyo Institute of Technology et AIST, Japon)
- Shuichi Katsumata (AIST, Japon)
- Eamonn W. Postlethwaite (Royal Holloway, UK)
- Bas Westerbaan (PQShield Ltd., UK)




Full paper: <https://eprint.iacr.org/2021/1407>

# Why Secure Messaging?



Secure messaging applications are widespread...

---

<b>WhatsApp</b> 	2.5 Billions
<b>Facebook Messenger</b> 	1.3 Billions
<b>Telegram</b> 	500 Millions
<b>Snapchat</b> 	280 Millions





---

...and represent attractive targets for attackers:

- "Al Jazeera journalists 'hacked via NSO Group spyware'", BBC, 2020  
<https://www.bbc.com/news/technology-55396843>
- "Grand jury subpoena for Signal user data, Central District of California", Signal , 2020  
<https://signal.org/bigbrother/central-california-grand-jury/>
- etc.

Secure messaging applications are widespread...

---

<b>WhatsApp</b> 	2.5 Billions
<b>Facebook Messenger</b> 	1.3 Billions
<b>Telegram</b> 	500 Millions
<b>Snapchat</b> 	280 Millions

---

...and represent attractive targets for attackers:

- "Al Jazeera journalists 'hacked via NSO Group spyware'", BBC, 2020  
<https://www.bbc.com/news/technology-55396843>
- "Grand jury subpoena for Signal user data, Central District of California", Signal , 2020  
<https://signal.org/bigbrother/central-california-grand-jury/>
- etc.

This talk: the security of secure *group* messaging *protocols*

Specific constraints:

→ Asynchrony



Specific constraints:

- Asynchrony
- Minimise trust in the server  
( $\implies$  end-to-end encryption)



Specific constraints:

- Asynchrony
- Minimise trust in the server  
( $\implies$  end-to-end encryption)
- Large number of users ( $N \gg 1$ )






Specific constraints:

- Asynchrony
- Minimise trust in the server  
( $\implies$  end-to-end encryption)
- Large number of users ( $N \gg 1$ )
- Very long sessions ( $t \gg 1$ )



Specific constraints:


- Asynchrony
- Minimise trust in the server  
( $\implies$  end-to-end encryption)
- Large number of users ( $N \gg 1$ )
- Very long sessions ( $t \gg 1$ )




If each user has a probability  $\epsilon$  of being compromised by an attacker during a unit of time, a group conversation with  $N$  members over  $t$  units of time will be compromised with probability  $1 - (1 - \epsilon)^{Nt}$ .

- This probability is significant as soon as  $Nt = \Omega(1/\epsilon)$ .
- **Solution:** periodically refresh encryption keys (next slides).


Forward secrecy (FS) [CCG16, CGCD<sup>+</sup>17, ACD19]:



Post-Compromise Security (PCS) [CCG16, CGCD<sup>+</sup>17, ACD19]:



Post-Compromise Forward Security (PCFS) [ACDT20, ACJM20, AJM20]:




CGKAs (*Continuous Group Key Agreement*) concentrate the cryptographic mechanisms of secure group messaging protocols:

- Add a user
- Remove a user
- **Remove one's encryption keypair (*Ratcheting/Commit Message*)**


Prominent CGKAs:

- *Pairwise Channels* (Signal)
- *Sender Keys* (WhatsApp)
- *TreeKEM* [BBR18, Wei19, BBN19, ACDT20, AJM20, ACJM20, ACC<sup>+</sup>21] (IETF MLS draft standard [OBR<sup>+</sup>21, BBM<sup>+</sup>20])
- *Chained mKEM* [BBN19]

## Chained mKEM




## TreeKEM




In Chained mKEM and TreeKEM:

- To each node (○, ⊙) is associated an encryption keypair
- A user knows the decryption key of a node if and only if this node is in their path (i.e. the node is an ancestor of the user's node)

## Chained mKEM



## TreeKEM

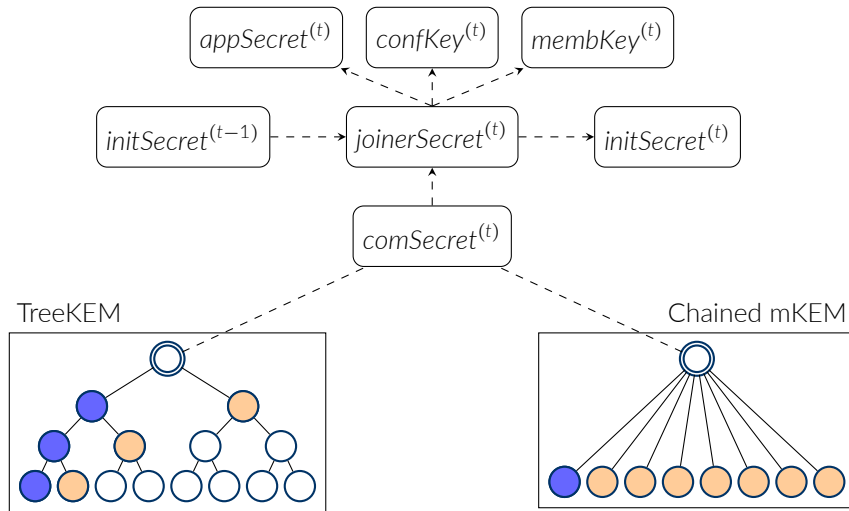


In Chained mKEM and TreeKEM:

- To each node ( $\bigcirc, \bigcirc$ ) is associated an encryption keypair
- A user knows the decryption key of a node if and only if this node is in their path (i.e. the node is an ancestor of the user's node)
- A commit message (here, sent by the leftmost user) contains:
  - an encryption key for each  $\bigcirc$
  - an asymmetric ciphertext for each  $\bigcirc$
  - two signatures (one that authenticates encryption keys, one that authenticates ciphertext)

Therefore it has a size  $O(N)$  for Chained mKEM, and  $\Omega(\log N)$  for TreeKEM.

# A change of epoch $(t - 1) \rightarrow t$



The larger a group size  $N$  is, the more *commit messages*:

- are necessary
- are costly

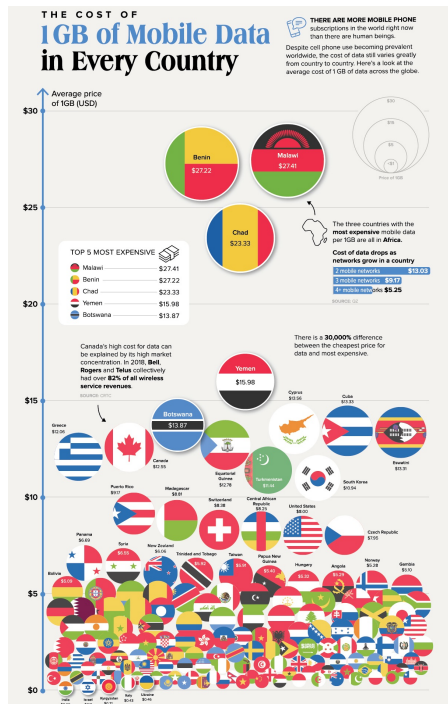
This tension is amplified by:

- the cost and impact for end users
- post-quantum cryptography ( $\times 10$  or more compared to classical cryptography).

Example with:

- TreeKEM
- Classic McEliece [ABC+20]
- 256 users

If each user sends one *commit message*, the bandwidth cost is 512 MiB per user.



<sup>1</sup><https://www.visualcapitalist.com/cost-of-mobile-data-worldwide/>



*Chained CmPKE*, un CGKA avec un coût asymétrique en envoi et réception (de *commit message*):

	Upload	Download	Total
Chained CmPKE	$O(N)$	$O(1)$	$O(N)$
TreeKEM	$\Omega(\log N)$	$\Omega(\log N)$	$\Omega(N \log N)$

Conceptually, the main change of Chained CmPKE is to make the server more active (but without entrusting it more).



Technically, Chained CmPKE is based on Chained mKEM, with two new ideas:

- 1 Use very efficient mPKEs (*multi-recipient PKE*)  $\Rightarrow$  reduce upload costs
- 2 The use of CmPKE (*committing mPKE*)  $\Rightarrow$  a cost  $O(1)$

# Ingredient N°1: *Committing* mPKEs



## Syntax of a mPKE (*multi-recipient PKE*):

$$\rightarrow \text{mEnc}(\mathbf{M}, (\text{ek}_i)_{i \in [N]}) \rightarrow (\text{ct}_0, (\widehat{\text{ct}}_i)_{i \in [N]})$$

$$\rightarrow \text{mDec}(\text{dk}_i, (\text{ct}_0, \widehat{\text{ct}}_i)) \rightarrow \{\mathbf{M} \text{ or } \perp\}$$

Recently revisited in [KKPP20], which inspired this work.

## The syntax of a CmPKE (*committing mPKE*) is identical:

$$\rightarrow \text{CmEnc}(\mathbf{M}, (\text{ek}_i)_{i \in [N]}) \rightarrow (\mathbf{T}, (\text{ct}_i)_{i \in [N]})$$

$$\rightarrow \text{CmDec}(\text{dk}_i, (\mathbf{T}, \text{ct}_i)) \rightarrow \{\mathbf{M} \text{ or } \perp\}$$

In addition, we require that  $\mathbf{T}$  is *committing*, i.e.  $\mathbf{T}$  is bound to a unique message  $\mathbf{M}$ . A related notion: *committing AEADs* [GLR17]

## We provide an (mPKE IND-CPA $\Rightarrow$ CmPKE IND-CCA) transform, with:

$$\rightarrow \text{ct}_i = \widehat{\text{ct}}_i.$$

$$\rightarrow \mathbf{T} = (\text{ct}_0, c) \text{ and } |c| = 32 \text{ bytes.}$$

Our transform uses *key-committing AEADs* [FOR17, GLR17, ADG<sup>+</sup>20].

In Chained mKEM, a commit message contains:

- 1 A new encryption key  $ek$
  - 2 An mPKE ciphertext:  $(ct_0, (\widehat{ct}_i)_{i \in [N]})$
  - 3 A signature  $sig_1 \leftarrow \text{Sign}(sk, ek)$
  - 4 A signature  $sig_2 \leftarrow \text{Sign}(sk, (ct_0, (\widehat{ct}_i)_{i \in [N]}))$
- 

In Chained CmPKE, a commit message contains:

## Upload:

- 1  $ek$
- 2 A CmPKE ciphertext:  $(T, (ct_i)_{i \in [N]})$
- 3  $sig_1 \leftarrow \text{Sign}(sk, ek)$
- 4  $sig_2 \leftarrow \text{Sign}(sk, T)$

## Download:

- 1  $ek$
- 2  $(T, ct_i)$
- 3  $sig_1$
- 4  $sig_2$

Intuitively, any attempt from the server to tamper with  $T$  or  $ct_i$  is detected upon signature verification or during decryption.

Ingredient Nº2:  
more efficient  
mPKEs

[KKPP20] highlighted the existence of very efficient mPKEs based on LWE, LWR and SIDH. Example with LPR-style schemes [LPR10, LP11]:

**Enc( $\mathbf{ek} = \mathbf{B}, \mathbf{M}$ )**

- 1 Sample short matrices  $\mathbf{R}, \mathbf{E}', \mathbf{E}''$
- 2  $\mathbf{U} \leftarrow \mathbf{R}\mathbf{A} + \mathbf{E}'$
- 3  $\mathbf{V} \leftarrow \mathbf{R}\mathbf{B} + \mathbf{E}'' + \text{Encode}(\mathbf{M})$
- 4  $\text{ct} := (\mathbf{U}, \mathbf{V})$

[KKPP20] highlighted the existence of very efficient mPKEs based on LWE, LWR and SIDH. Exemple with LPR-style schemes [LPR10, LP11]:

## Enc( $\mathbf{ek} = \mathbf{B}, \mathbf{M}$ )

- ① Sample short matrices  $\mathbf{R}, \mathbf{E}', \mathbf{E}''$
- ②  $\mathbf{U} \leftarrow \mathbf{RA} + \mathbf{E}'$
- ③  $\mathbf{V} \leftarrow \mathbf{RB} + \mathbf{E}'' + \text{Encode}(\mathbf{M})$
- ④  $\text{ct} := (\mathbf{U}, \mathbf{V})$

$\Rightarrow$

## mEnc( $\{\mathbf{ek}_1, \dots, \mathbf{ek}_N\}, \mathbf{M}$ )

- ① Sample short matrices  $\mathbf{R}, \mathbf{E}'$
- ②  $\mathbf{U} \leftarrow \mathbf{RA} + \mathbf{E}'$
- ③ For  $i = 1, \dots, N$ :
  - ① Sample a short matrix  $\mathbf{E}_i''$
  - ②  $\mathbf{V}_i \leftarrow \mathbf{RB}_i + \mathbf{E}_i'' + \text{Encode}(\mathbf{M})$
- ④  $(\text{ct}_0, (\widehat{\text{ct}}_i)_{i \in [N]}) := (\mathbf{U}, (\mathbf{V}_i)_{i \in [N]})$

Limitations of [KKPP20]:


- ➔ Naive parametrisations
- ➔ No concrete security analysis

We propose three re-parametrisations of lattice-based (m)PKEs:

**BilboKEM640** (inspired  
from FrodoKEM640)



**Illum512** (inspired from  
Kyber512)



**LPR757** (inspired from NTRU  
LPR653)



These mPKEs are tailored for the Chained CmPKE protocol:

- The  $(\widehat{ct}_i)_{i \in [N]}$  are extremely small
- This entails a small increase of the sizes of  $ek$  and  $ct_0$




Recall that  $\mathbf{M} = \text{Decode}(\mathbf{V}_i - \text{dk}_i \cdot \mathbf{U})$ . Our toolkit:

- 🔧 **Bit dropping:** cut the least significant bits of  $\mathbf{V}_i$ 
  - 👍 Reduces the size of  $\mathbf{V}_i$ , increases the LWEv error rate
  - 👎 Increases the *decryption failure rate*
- 🔧 **Coefficient dropping:** cut superfluous coefficients of  $\mathbf{V}$ 
  - 👍 Reduces the size of  $\mathbf{V}_i$
  - 👎 None!
- 🔧 **Increase the modulus  $q$** 
  - 👍 Allows to pack more bits of key per coefficient of  $\mathbf{V}_i$
  - 👎 Increases the size of  $\mathbf{U}$ , reduces the LWE error rate
- 🔧 **Error correcting codes** (discarded option)
  - 👍 Reduces the decryption failure rate
  - 👎 Timing attacks, delicate security analysis [DVV19, GJY19, DTVV19]

---

The main attacks to (re-)consider:



Communication costs of mPKEs based on existing ( gray background ) and new (fond blanc) parametrisations. Sécurité: NIST I ( $\geq$  AES-128).


Schéma	$ ek $	$ ct_0 $	$ \widehat{ct}_i $
Kyber512 [SAB <sup>+</sup> 20]	768 (+32)	640	<b>128</b>
Illum512	768	704	<b>48</b>
LPRime653 [BBC <sup>+</sup> 20]	865 (+32)	865 (+32)	<b>128</b>
LPRime757	1076	1076	<b>32</b>
Frodo640 [NAB <sup>+</sup> 20]	9600 (+16)	9600	<b>120</b>
Bilbo640	10240	10240	<b>24</b>
SIKEp434 [JAC <sup>+</sup> 20]	330	330	<b>16</b>

# Comparisons & Conclusion




# Chained CmPKE vs TreeKEM (upload and download)

Size of a commit message in KiB as a function of the group size.



(a) Bilbo640 vs. Frodo640




(b) SIKEp434 vs. SIKEp434


- Chained CmPKE (**upload**) with {Bilbo640, SIKEp434}
- Chained CmPKE (**download**) with {Bilbo640, SIKEp434}
- ◆ TreeKEM (**upload and download**) with {Frodo640, SIKEp434}

# Chained CmPKE vs TreeKEM (total normalised by $N$ )

Normalised cost of a commit message in KiB as a function of the group size.




(a) Bilbo640 vs. Frodo640




(b) SIKEp434 vs. SIKEp434

- Chained CmPKE with {Bilbo640, SIKEp434}
- ♦ TreeKEM with {Frodo640, SIKEp434}

# Computational cost of a commit message



(a) Commit (sender side)



(b) Process (receiver side)


Running time of some procedures as a function of the group size, for Illum512 (—■—), LPRime757 (—▲—), Bilbo640 (—●—), SIKEp434 (—◆—). Logarithmic scales. Timings obtained on an Apple M1 @3.2 GHz.

We proposed *Chained CmPKE*, a CGKA that is:

- Very fast
- More compact than TreeKEM:  $O(N)$  instead of  $\Omega(N \log N)$
- Simpler than TreeKEM
- Satisfying the same security notions (see paper)

As well as techniques that might be of independent interest:

- 1 *Committing mPKEs*
- 2 More efficient lattice-based mPKEs



Questions?





Martin R. Albrecht, Daniel J. Bernstein, Tung Chou, Carlos Cid, Jan Gilcher, Tanja Lange, Varun Maram, Ingo von Maurich, Rafael Misoczki, Ruben Niederhagen, Kenneth G. Paterson, Edoardo Persichetti, Christiane Peters, Peter Schwabe, Nicolas Sendrier, Jakub Szefer, Cen Jung Tjhai, Martin Tomlinson, and Wen Wang.

Classic McEliece.

Technical report, National Institute of Standards and Technology, 2020.

available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions>.



Joël Alwen, Margarita Capretto, Miguel Cueto, Chethan Kamath, Karen Klein, Ilia Markov, Guillermo Pascual-Perez, Krzysztof Pietrzak, Michael Walter, and Michelle Yeo.

Keep the dirt: Tainted treekem, adaptively and actively secure continuous group key agreement.  
In *2021 IEEE Symposium on Security and Privacy (S&P)*, pages 596–612, Los Alamitos, CA, USA, may 2021. IEEE Computer Society.



Joël Alwen, Sandro Coretti, and Yevgeniy Dodis.

The double ratchet: Security notions, proofs, and modularization for the Signal protocol.

In Yuval Ishai and Vincent Rijmen, editors, *EUROCRYPT 2019, Part I*, volume 11476 of *LNCS*, pages 129–158. Springer, Heidelberg, May 2019.



Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis.

Security analysis and improvements for the IETF MLS standard for group messaging.

In Daniele Micciancio and Thomas Ristenpart, editors, *CRYPTO 2020, Part I*, volume 12170 of *LNCS*, pages 248–277. Springer, Heidelberg, August 2020.



Joël Alwen, Sandro Coretti, Daniel Jost, and Marta Mularczyk.

Continuous group key agreement with active security.

In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020, Part II*, volume 12551 of *LNCS*, pages 261–290. Springer, Heidelberg, November 2020.



Ange Albertini, Thai Duong, Shay Gueron, Stefan Kölbl, Atul Luykx, and Sophie Schmieg.

How to abuse and fix authenticated encryption without key commitment.

Cryptology ePrint Archive, Report 2020/1456, 2020.

<https://eprint.iacr.org/2020/1456>.



Joël Alwen, Daniel Jost, and Marta Mularczyk.

On the insider security of MLS.

Cryptology ePrint Archive, Report 2020/1327, 2020.

<https://eprint.iacr.org/2020/1327>.



Daniel J. Bernstein, Billy Bob Brumley, Ming-Shing Chen, Chitchanok Chuengsatiansup, Tanja Lange, Adrian Marotzke, Bo-Yuan Peng, Nicola Tuveri, Christine van Vredendaal, and Bo-Yin Yang.

NTRU Prime.

Technical report, National Institute of Standards and Technology, 2020.

available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions>.



Richard Barnes, Benjamin Beurdouche, Jon Millican, Emad Omara, Katriel Cohn-Gordon, and Raphael Robert.

The Messaging Layer Security (MLS) Protocol.

Internet-Draft draft-ietf-mls-protocol-11, Internet Engineering Task Force, December 2020.  
Work in Progress.



Karthikeyan Bhargavan, Benjamin Beurdouche, and Prasad Naldurg.

Formal Models and Verified Protocols for Group Messaging: Attacks and Proofs for IETF MLS.  
Research report, Inria Paris, December 2019.



Karthikeyan Bhargavan, Richard Barnes, and Eric Rescorla.

TreeKEM: Asynchronous Decentralized Key Management for Large Dynamic Groups A protocol proposal for Messaging Layer Security (MLS).  
Research report, Inria Paris, May 2018.



Katriel Cohn-Gordon, Cas J. F. Cremers, and Luke Garratt.

On post-compromise security.

In Michael Hicks and Boris Köpf, editors, *CSF 2016 Computer Security Foundations Symposium*, pages 164–178. IEEE Computer Society Press, 2016.



Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila.

A formal security analysis of the signal messaging protocol.

In *2017 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 451–466, 2017.



Jan-Pieter D’Anvers, Marcel Tiepelt, Frederik Vercauteren, and Ingrid Verbauwhede.

Timing attacks on error correcting codes in post-quantum schemes.

In Begül Bilgin, Svetla Petkova-Nikova, and Vincent Rijmen, editors, *TIS@CCS*, pages 2–9. ACM, 2019.



Jan-Pieter D'Anvers, Frederik Vercauteren, and Ingrid Verbauwhede.

The impact of error dependencies on ring/mod-LWE/LWR based schemes.

In Jintai Ding and Rainer Steinwandt, editors, *Post-Quantum Cryptography - 10th International Conference, PQCrypto 2019*, pages 103–115. Springer, Heidelberg, 2019.



Pooya Farshim, Claudio Orlandi, and Razvan Rosie.

Security of symmetric primitives under incorrect usage of keys.

*IACR Transactions on Symmetric Cryptology*, pages 449–473, 2017.



Qian Guo, Thomas Johansson, and Jing Yang.

A novel CCA attack using decryption errors against LAC.

In Steven D. Galbraith and Shiho Moriai, editors, *ASIACRYPT 2019, Part I*, volume 11921 of *LNCS*, pages 82–111. Springer, Heidelberg, December 2019.



Paul Grubbs, Jiahui Lu, and Thomas Ristenpart.

Message franking via committing authenticated encryption.

In Jonathan Katz and Hovav Shacham, editors, *CRYPTO 2017, Part III*, volume 10403 of *LNCS*, pages 66–97. Springer, Heidelberg, August 2017.



David Jao, Reza Azarderakhsh, Matthew Campagna, Craig Costello, Luca De Feo, Basil Hess, Amir Jalali, Brian Koziel, Brian LaMacchia, Patrick Longa, Michael Naehrig, Joost Renes, Vladimir Soukharev, David Urbanik, Geovandro Pereira, Koray Karabina, and Aaron Hutchinson.

SIKE.

Technical report, National Institute of Standards and Technology, 2020.

available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions>.



Shuichi Katsumata, Kris Kwiatkowski, Federico Pintore, and Thomas Prest.

Scalable ciphertext compression techniques for post-quantum KEMs and their applications.

In Shiho Moriai and Huaxiong Wang, editors, *ASIACRYPT 2020, Part I*, volume 12491 of *LNCS*, pages 289–320. Springer, Heidelberg, December 2020.



Richard Lindner and Chris Peikert.

Better key sizes (and attacks) for LWE-based encryption.

In Aggelos Kiayias, editor, *CT-RSA 2011*, volume 6558 of *LNCS*, pages 319–339. Springer, Heidelberg, February 2011.



Vadim Lyubashevsky, Chris Peikert, and Oded Regev.

On ideal lattices and learning with errors over rings.

In Henri Gilbert, editor, *EUROCRYPT 2010*, volume 6110 of *LNCS*, pages 1–23. Springer, Heidelberg, May / June 2010.



Michael Naehrig, Erdem Alkim, Joppe Bos, Léo Ducas, Karen Easterbrook, Brian LaMacchia, Patrick Longa, Ilya Mironov, Valeria Nikolaenko, Christopher Peikert, Ananth Raghunathan, and Douglas Stebila.

FrodoKEM.

Technical report, National Institute of Standards and Technology, 2020.

available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions>.



Emad Omara, Benjamin Beurdouche, Eric Rescorla, Srinivas Inguva, Albert Kwon, and Alan Duric.  
The Messaging Layer Security (MLS) Architecture.  
Internet-Draft draft-ietf-mls-architecture-06, Internet Engineering Task Force, March 2021.  
Work in Progress.



Peter Schwabe, Roberto Avanzi, Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M. Schanck, Gregor Seiler, and Damien Stehlé.  
CRYSTALS-KYBER.

Technical report, National Institute of Standards and Technology, 2020.  
available at <https://csrc.nist.gov/projects/post-quantum-cryptography/round-3-submissions>.



Matthew Weidner.  
Group messaging for secure asynchronous collaboration.  
Mphil dissertation, University of Cambridge, Cambridge, UK, 2019.