Question & Answer
Lightweight Cryptography Workshop 2020
OCTOBER 19-21, 2020

Monday, October 19, 2020

Posted Questions

[11:58 AM]
Damian Vizar asked : A comment rather than a question: as has been already quite heatedly discussed during CAESAR, saying that an AEAD provides security *without* a nonce (assuming it is not internally randomized) is misleading: the user may falsely believe in getting full privacy even for plaintexts with low entropy.
4 upvotes | 0 answer | 0 reply

[11:26 AM]
gokull asked : I am very new to this . I have no basics in cryptography . will this workshop be useful for me ?
3 upvotes | 1 answer | 0 reply

Meltem Sonmez Turan answered -
Some of the talks are accessible for beginners. Hope you enjoy.

[11:41 AM]
prathusha k asked : can we get these ppt's?
3 upvotes | 1 answer | 0 reply

Kerry McKay answered -
Slides will be posted after the workshop. In about 2 weeks, the video will be available as well.

[11:40 AM]
Adam Brackmann asked : What was the selection criteria for the specific microarchitectures on which algorithms were benchmarked?
1 upvote | 0 answer | 0 reply

[12:59 PM]
Avik Chakraborti asked : @ Rishub. This is the answer to Rishub's question.

If you see Page 14 in the link
https://csrc.nist.gov/CSRC/media/Projects/lightweight-cryptography/documents/round-2/spec-doc-rnd2/saturnin-spec-round2.pdf, you can observe that,
Saturnin-BC generates Enc key stream and auth key stream parallely
1 upvote | 0 answer | 0 reply

[02:35 PM]
Dong Hoon Chang asked : In case of Romulus-H (hash function, not AEAD), do you also want to propose to use Skinny with reduced number of rounds? The reason why I ask this question because most of the existing analysis results on Skinny are presented when a key is secret.
1 upvote | 2 answers | 1 reply

Thomas Peyrin answered -
Indeed, but in the Hirose DBL construction, you have a huge constraint for the attacker: the same message/chaining variable will go to both TBC calls. So you don't have to control one, but two TBC calls at the same time.

Thomas Peyrin answered -
Also, more analysis on Romulus-H is incoming with this regards :-)

Dong Hoon Chang replied -
Ok I see. Thank for the clarification.

[11:28 AM]
Damian Vizar asked : A question for Kerry's talk:
The criteria for the selection of 3rd round candidates did not contain a mention of the lightweight use cases and constraints, which motivate the LWC project. I'd like to ask how will these be factored into the decision?
0 upvote | 1 answer | 0 reply

Meltem Sonmez Turan answered -
Thanks Damian, I will add this question to the open discussion session.

[11:36 AM]
Robert Moskowitz asked : As a protocol designer, I value looking beyond hashing to how the hash is efficient in use and protocol design. SHAKE provides a standardized variable hash output. KMAC provides a new approach over the venerable HMAC. It also seems to be superior to HKDF. Are these modes considered?
0 upvote | 0 answer | 0 reply

[11:42 AM]
Robert Moskowitz asked : Is comparing code size to SHA256 misleading? Is there any code size increase in the candidate to provide hashing over having to include SHA code?
0 upvote | 0 answer | 0 reply

[11:47 AM]
Robert Moskowitz asked : What if you do not need SHA at all?
0 upvote | 1 answer | 0 reply

Meltem Sonmez Turan answered -
AEAD benchmarking based on code size was independent of the hash selection,

[11:48 AM]
Robert Moskowitz asked : Yes! Combined compared to AES + SHA

0 upvote | 0 answer | 0 reply

[12:04 PM]
Damian Vizar asked : This time a question: Are you thinking of setting up a zoo-type website?
0 upvote | 3 answers | 1 reply

Meltem Sonmez Turan answered -
Is this question to the speaker, or to the NIST LWC team?

Damian Vizar replied -
Good point Meltem :) Originally to the speaker, but if NIST would reply affirmatively, we woud be interested to learn that as well!

Cagdas Calik answered -
We're going to publish the benchmarking framework code, the list of implementations benchmarked and the results on github.

Cagdas Calik answered -
The answer above was for the software benchmarking framework.

[12:15 PM]
Rishub Nagpal asked : Why is Saturnin considered a Stream cipher? AEAD uses a

nonce + counter, and a second pass over the ciphertext for tag generation.
0 upvote | 0 answer | 0 reply

[12:26 PM]
sebastien riou asked : I missed the intro and first talk. is there a way to see them in replay ?
0 upvote | 2 answers | 1 reply

Meltem Sonmez Turan answered -
Sebastien, in about 2 weeks, the videos of the talks will be available on the project webpage.

sebastien riou replied -
thanks, will the slides be available before that ?

Kerry McKay answered -
Yes. They will be available soon after the workshop.

[01:02 PM]
Avik Chakraborti asked : @ Rishub
Saturnin-BC generates Enc key stream and adds it to msg blocks in line 1
For all i = 0 to ℓ : $ci \leftarrow mi \oplus$ Saturnin1 (k, N||i + 1)

ci, N and ai s are processed using the auth key stream t generated by Saturnin-BC.Overall, it outputs Enc and Auth key streams to compute ciphertxt and tag
0 upvote | 0 answer | 0 reply

[01:03 PM]
Avik Chakraborti asked : @Rishub
It is a two pass scheme and the counter is independent of the message blocks.

0 upvote | 0 answer | 0 reply

[01:06 PM]
Rishub Nagpal asked : (Not sure how to make replies in a comment thread)

Thanks @Avik, I understand now.
0 upvote | 0 answer | 0 reply

[01:20 PM]
Damian Vizar asked : The remark about not-ideal permutation-based sponge is very interesting; do you see any practical implications/benefits of trading a cryptographic permutation for a TBC in a sponge mode?

0 upvote | 0 answer | 0 reply

[01:57 PM]
Dong Hoon Chang asked : Rectangle block cipher is based on 128-bit block. On the other hand, KNOT permutation is based on 256-bit block. According to your result, the differential probablity for 14-round is about 2^-60 for both cases. I wonder why the probability would be similar with the same number of rounds?
0 upvote | 4 answers | 3 replies

Wentao Zhang answered -
The block length of RECTANGLE is 64 bits. The probability of the best differential trail for 14-round RECTANGLE is 2^-61.

For the knot members with 256-bit state, the probability of the best differential trail for 14 rounds is 2^-71.

Dong Hoon Chang replied -
Intuitively, it seems that 256-bit provides more flexibility compared to 64-bit from the attacker's point of view. But your result seesm to be opposite. Could you explain the intuitive reason why 256-bit version (KNOT perm) is worse than 64-bit version in terms of differential probability?

Wentao Zhang answered -
In terms of differential probability, the 256-bit KNOT permutation is better than RECTANGLE. Both for 14 rounds, the probability of the best differential trail is 2^-71 for the 256-bit KNOT permutation, and 2^-61 for RECTANGLE. The lower this probability, the better :-)

Dong Hoon Chang replied -
can you explain why KNOT is better from the design (internal structure or round function) point of view?

Wentao Zhang answered -
KNOT and RECTANGLE are different symmetric-key primitives. The design of the KNOT permutations inherits the design of RECTANGLE, but with different state size. For a same number of rounds (more than 8 rounds), the 256-bit KNOT permutation is better than RECTANGLE w.r.t. differential attack. In my opinion, the state size is the main reason, the bigger the state size, the better the diffusion for long rounds. Hope it can be helpful to you, and thanks for your interest in KNOT.

Dong Hoon Chang replied -
Ok I see. Thank you for your opinion.

Wentao Zhang answered -
My pleasure.

[02:02 PM]
Dong Hoon Chang asked : Correcting: Rectangle blockcipher is based on 64-bit block.
0 upvote | 0 answer | 0 reply

[02:16 PM]
Gilles Van Assche asked : Comment: Keccak-f[200] is not "kind of" standardized, it is actually part of FIPS 202. :-)
0 upvote | 0 answer | 0 reply

[02:52 PM]
Florian Mendel asked : Thomas, is the threshold implementation of Romulus online and did you evaluate the implementation? How many shares are used to protect the key resp key schedule?
0 upvote | 3 answers | 2 replies

Thomas Peyrin answered -
the TI implementations of Romulus are built upon the original Skinny team TI implementations
https://eprint.iacr.org/2016/660.pdf

Thomas Peyrin answered -
so this is 3 shares

Florian Mendel replied -
Also for the key?

Thomas Peyrin answered -
yes !

Florian Mendel replied -
OK. Thanks.

Tuesday, October 20, 2020

Posted Questions

[11:09 AM]
Mustafa asked : One of the criticisms to memory encryption using standard AEAD is the memory overhead (about 32 bytes per chunk for the nonce and tag). For a chunk size of 64 bytes, that's 50% increase in memory. From your experience, do you think the cost is justifiable and practical?
0 upvote | 3 answers | 1 reply

sebastien riou answered -
chunk size of 64 bytes is a bit extreme but can make sense when performances have to be very high. we tend to favor larger chunk size to have reasonable overhead. That said, the cost is relatively low in any case because only a small part of the flash is used like that. The SOC application is typically using much more flash for non secure things (or things which are bulk encrypted).

sebastien riou answered -
anyway, the cost is practical, we + competitors are shipping products, it is not just a concept :-)

Mustafa replied -
I see thanks a lot, that answers my question

Kris Gaj answered -
Could you please define what time is being measured? What types of inputs and of what size are being processed?

[11:18 AM]
Kamyar Mohajerani asked : How much power, energy, or area efficiency gain do you think could be expected if a candidate supporting both AEAD and Hashing operations were to be used, instead of relying on SHA-3?
0 upvote | 1 answer | 1 reply

sebastien riou answered -
I do not have data to answer this section. Clearly on area there are huge gains to be made by replacing AES+SHA3 with LWC candidate which provide both

Kamyar Mohajerani replied -
Thank you for your answer. I realize I misunderstood one of the slides.

[11:23 AM]
Archanaa S Krishnan asked : @Sebastian Riou - Did Tiempo consider using software implementation of LWC candidates for existing SoCs with ext NVM?
0 upvote | 1 answer | 1 reply

sebastien riou answered -
we will consider this once one is NIST approved. before that we cannot sell our product based on LWC candidate

Archanaa S Krishnan replied -
Thank you.

[11:35 AM]
Gaëtan Cassiers asked : @Sebastien Riou: In presence of side-channel attack, do you need confidentiality for the plaintexts that are encrypted/decrypted while the SCA is performed ? Or is protection of confidentiality for plaintexts not manipulated during the SCA enough (assuming integrity is always guaranteed) ?

0 upvote | 1 answer | 1 reply

sebastien riou answered -
this is a bit of grey area. a sweet spot is reached when you have low general claims for the confidentiality of the plaintext. for example when you allow to leak hamming weight. for plaintext which happen to be really critical, we recommend to add another layer of encryption because merely transporting it over the internal bus is going to leak.

Gaëtan Cassiers replied -
Thank you.

[11:37 AM]
Meltem Sonmez Turan asked : @Kalikindan: Are the straight line programs for the second round candidates available online?
0 upvote | 0 answer | 0 reply

[11:39 AM]
Meltem Sonmez Turan asked : @Kalikindar Did you apply any heuristics to optimize gate counts?
0 upvote | 0 answer | 0 reply

[12:22 PM]
Mustafa asked : @Liliya, Thanks a lot for the talk, I think multi-key attacks and related-key attacks are interesting and should give more insight about the design. However, do you think getting 2^100 data complexity under one key is as hard/easy as to getting 2^100 data complexity using related keys?

0 upvote | 1 answer | 0 reply

Liliya Kraleva answered -
Thank you for the question. The related-key scenario is indeed more restrictive in practice, therefore it is harder to sample ciphertexts using related keys. However, a related-key attack can be applied, for exemple, in a scenario in which the Keys are not generated at random, or the used PRNG has some (unknown) weaknesses.

[01:03 PM]
Tetsu Iwata asked : Is it possible to see yesterday's chat/Q&A?
0 upvote | 1 answer | 1 reply

Kevin Hill answered -
Not on the blue Jeans platform now but Q & A questions will be avaliable after conference

Tetsu Iwata replied -
Thank you!

[01:22 PM]
Gilles Van Assche asked : @ Sebastian: Thanks for the presentation. Would it be possible to split time measurements into a fixed time (or #cycles) + a slope in time/byte (or cycles/byte)?
0 upvote | 0 answer | 0 reply

[01:32 PM]
Raghav Rohit asked : @Alexandre : Thanks for the talk. Any comments or suggestions on applying this technique to ciphers based on Generalized Feistel Network?
0 upvote | 1 answer | 1 reply

Alexandre Adomnicai answered -
We didn't have a look at such designs for the moment but maybe worth investigating it!

Raghav Rohit replied -
Thanks.

[01:35 PM]
sebastien riou asked : a remark: double computation can be use to protect against faults: https://patents.justia.com/patent/10341085
0 upvote | 1 answer | 0 reply

Alexandre Adomnicai answered -
Thanks for the comment. Yep indeed the block that is 'computed for nothing' can be probably used to build a countermeasure!

[01:50 PM]
Mustafa Khairallah asked : Thanks Patrick for the talk. It would be interesting to see the difference between SpoC64 and Ascon in your example of the CAESAR API when the use the same FIFO configuration, any idea about that case?
0 upvote | 0 answer | 0 reply

[02:20 PM]
Thomas Peyrin asked : question for the speaker and also for the NIST: the candidates are ranked according to throughput. Does that really make sense for lightweight ? Shouldn't we instead rank with other metrics such as throughput/area, etc. ?
0 upvote | 2 answers | 2 replies

sebastien riou answered -
Power x Area x Latency, no way to cheat that one

Thomas Peyrin replied -
fully agreed

Kris Gaj answered -
I do not. The users do not see Throughput/Area ot Power*Area*Time. They can only observe/experience values of individual metrics. These combined metrics make sense where there is a linear relationship between individual metrics. This kind of linear relationship does not exist for lightweight hardware architectures.

Thomas Peyrin replied -
most lightweight applications will only have a small throughput requirement ...

[02:27 PM]
Mustafa Khairallah asked : I don't think Throughput/Area or PxAxLatency necessarily mean falsely linearizing the throughput area relation, but rather it is an efficiency metric for a given implementation, wouldn't you agree? It does not provide false information about other implementations

0 upvote | 2 answers | 1 reply

Kris Gaj answered -
How has using more area helped to improve the throughput of Romulus? The users do

not see Throughput/Area ot Power*Area*Time. They can only observe/experience values of individual metrics.

Mustafa Khairallah replied -
Throughput/Area is not a metric of the area cost to increase the throughput but rather the cost you pay for a specific architecture, there is no evidence to suggest relation between throughput and area is linear.

You can have one Throughput/Area value for a given implementation and a completely different value for a different implementation of the same algorithm, as can be seen in your graphs. It is an implementation specific metric and not an algorithm specific metric. As you described increasing teh area DID increae the throughput of ROmulus up to a certain point, where the throughput afterwards plateaud, this exactly the point of combined implementation-specific metrics in my view :)

Kris Gaj answered -
If this is not an algorithm specific value, then why would you use it to compare algorithms? The meanigful comparisons are those in which selected metrics are limited by a certain threshold, and we try to optimize a single metric which has a physical meaning and can be observed by the users.

Wednesday, October 21, 2020

Posted Questions

[02:08 PM]
Adam Brackmann asked : Will ease of implementation be considered in the finalist selection?
5 upvotes | 2 answers | 0 reply

Cagdas Calik answered -
Easiness may be subjective, do you mean SW or HW, can you clarify?

Kerry McKay answered -
Ease of implementing unprotected, protected, or both?

[02:21 PM]
Tetsu Iwata asked : About additional features, some schemes claim 64-bit security, some claims 128-bit security, or more. Do you consider this in the selection?
5 upvotes | 0 answer | 0 reply

[01:24 PM]
Mustafa Khairallah asked : Thanks for the talk. You mention that WAGE has minimal interface 3200 GE on TSMC 65nm, can you please elaborate on what you mean by this statement?
3 upvotes | 1 answer | 1 reply

nusa zidaric answered -
some differences in api:
- we have 64-bit wide data I/Os

- envrionemt pads data to 64 bits
- WAGE HW is unaware of the length of AD, MSG
- we do not compare tag (left to environment)

Mustafa Khairallah replied -
I see.
Thanks a lot for your answer

[02:03 PM]
Luan asked : When you say that selection will consider additional features, do you mean things like SCA security, or using the primitives to instantiate (tweakable)block ciphers and XOFs for example?
3 upvotes | 0 answer | 0 reply

[02:18 PM]
Mustafa Khairallah asked : There are several designs that are secure within the NIST requirements in the single-key setting, but the data limits can be bypassed through targetting multi-users, either due to short tags, or weak keys, or other properties, where do you stand on that?

3 upvotes | 0 answer | 0 reply

[02:10 PM]
sebastien riou asked : a benchmarking platform using a software/hardware portioning would be nice. typically with core permutation in hardware and AEAD mode in software.
2 upvotes | 0 answer | 0 reply

[02:24 PM]
sebastien riou asked : Side channel and fault attacks will be taken into account for AEAD. Will it be the case also for hash ? (Do you consider HMAC based on LWC hash a target use case ?)
2 upvotes | 0 answer | 0 reply

[02:25 PM]
Rishub Nagpal asked : Is there a possibility of more than one finalist being standardized? For example, if one finalist excels at profile 2
2 upvotes | 0 answer | 0 reply

[02:26 PM]
Robert Moskowitz asked : Something like KMAC rather than HMAC. 1 hash function compared to 2 hash functions.
2 upvotes | 0 answer | 0 reply

[02:02 PM]
Adrian Neftali Sanchez asked : Which of the candidates do you consider the most promising?
1 upvote | 1 answer | 0 reply

Kerry McKay answered -
You will find out when we announce the finalists :)

[02:15 PM]
sebastien riou asked : this is mandatory in the real world. if you encrypt that's because you have something secret with some value. if it has value and it is not protected, people are going to get it...
1 upvote | 0 answer | 0 reply

[02:30 PM]
Siddaramappa asked : Thanks.
1 upvote | 0 answer | 0 reply

[11:25 AM]
Luan asked : Question for the ASCON presentation: You showed a single digit CPB for ascon, is that on an x86 processor?
0 upvote | 1 answer | 0 reply

Martin Schläffer answered -
Yes, this is on an AMD Ryzen 7 x86-64 processor. Results are taken from http://bench.cr.yp.to/results-aead.html

[11:36 AM]
Donghoon Chang asked : ISAP performs 1-round only (for performance gain) to process each bit of nonce to claim that ISAP is secure against DPA. But, we cannot expect any kind of security or randomness out of 1-round computation. So, your security claims seem to be heuristic. How about your opinion on this?
0 upvote | 2 answers | 3 replies

Robert Primas answered -
There are already quite a few publications that analyze isaps re-keying function (see e.g. the isap paper at the tosc special issue). Also, isap is not the only scheme that performs 1 permutation round. ketje is one such example that performs absorb and squeeze with higher rate (8 or 16-bit) and there are, to my knowlege, no attacks that can threaten ketje. isaps parameterization is even more conservative (since only 1-bit rate). we also specify variants of isap that perform more rounds in the rekeying function.

Donghoon Chang replied -
Can you specify where ISAP re-keying function is analyzed? I cannot find it from tosc 2020.

Donghoon Chang replied -
I am interested in what is the assumption on 1-round of permutation of ISAP to guarantee the security of ISAP re-keying against DPA.

Robert Primas answered -
section 4 of the tosc submission contains a summary of our analysis but also references other works. sections 4.3 and 4.6.1 of the design documents contain some statements

about the parameterization.

put simply, we require that an attacker can only reduce the entropy of a state per absorbed bit by a small amount so that recovery of the whole secret state would require a combination of information over several iterations of the permutation. we performed an analysis of differntial properties for 1-bit input differences.

Donghoon Chang replied -
ok I got your point. Thank you Robert.

[11:57 AM]
Rishub Nagpal asked : Would a ROM-based SBOX implementation mitigate this attack?
0 upvote | 0 answer | 0 reply

[12:00 PM]
Cyrus Minwalla asked : Have you explored any reinforcement learning attacks on Ascon?
0 upvote | 3 answers | 2 replies

Keyvan Ramezanpour answered -
Yes, we used a reinforcement learning attack for power analysis on Ascon. the results show that the efficiency is higher than classical DPA and CPA techniques

Cyrus Minwalla replied -
Very interesting. Could you point me to your paper/research?

Keyvan Ramezanpour answered -
Sure, the work is available on arXiv. Please find below the link to the paper:
https://arxiv.org/abs/2006.03995

Cyrus Minwalla replied -
Awesome, thank you.

Keyvan Ramezanpour answered -
you are welcome

[01:23 PM]
Donghoon Chang asked : @Ga¨etan: Can you share the full paper of your presentation?
0 upvote | 2 answers | 0 reply

Gaëtan Cassiers answered -
This paper contains most of the content of the talk: https://eprint.iacr.org/2020/211. Please see also the references section of the slides (they should be available soon on the workshop webpage).

Donghoon Chang answered -
Thank you Gaëtan.

[01:36 PM]
k asked : How can we get these presentations?

0 upvote | 1 answer | 0 reply

Sara Kerman answered -
They will be on the LWC 2020 webpage
https://csrc.nist.gov/events/2020/lightweight-cryptography-workshop-2020

[02:01 PM]
Mustafa Khairallah asked : Is mid-November a reasonable time for publishing the final benchmarking report, to meet your deadline?
0 upvote | 0 answer | 0 reply

[02:02 PM]
Damian Vizar asked : I would reiterate my question: how does the "ligtweight" use case requirements going to be applied to the selection criteria?
0 upvote | 1 answer | 1 reply

Kerry McKay answered -
It is coming up in the slides

Damian Vizar replied -
Thanks Kerry!

[02:02 PM]
Thomas Peyrin asked : will the NIST advise on the possible tweaks for the selected candidates (when the selection is announced) ?
0 upvote | 1 answer | 1 reply

Kerry McKay answered -
yes, we will provide guidance

Thomas Peyrin replied -
Thanks !