Dear authors of the paper "Key Recovery Attack against 2.5-round π-Cipher"

Thank you very much for sending us your analysis of Pi-Cipher. We are happy that Pi-Cipher is the subject of that analysis and that it attracted your interest. We highly value the time and efforts that you spent to analyze it. We carefully examined your attack and as a result we have the following remarks that can be summarized in **three main parts**:

1. **You claim that the attack works both when an SMN is used and without a SMN.** In particular, on page 8 you say: "We assume that a private message number PMN is used, but the attack also works if there is no PMN by removing the corresponding steps." (btw, there is a mistake, in the abbreviation, it is SMN – secret message number). **However, your attack works only when |SMN|=0.** Here is an elaboration:
   You state that you use guess and determine attack under the known plaintext-ciphertext scenario. Your findings until the recovery of the common internal state CIS'' are valid and we agree with them. But how you proceed with the attack after that is not correct when a SMN is used, and does not follow the instructions given by the official CAESAR submission call. You state that you already know SMN and with the previously obtained CIS'', you can easily recover the CIS and the key. But, there is the flaw. SMN is **a secret number** that should be decrypted from the ciphertext C_0 during the decryption phase, and it is not known. In order to proceed with the key recovery, you need to guess the SMN which is 128 bits in the pi16-Cipher096 version. So, basically your attack ends up with a complexity of 2^{128}. Additionally, knowing the CIS'' state you can only break the confidentiality of the plaintext under the nonce misuse scenario where PMN and SMN are repeated. But in our official CAESAR documentation we have the following statement: "**An intermediate level of nonce-misuse resistance**. In this case, the robustness is manifested when legitimate key holder uses the same key K, the same associated data AD, the same public message number PMN but different secret message numbers SMN_1 and SMN_2."
   Therefore, you should explicitly state that your key recovery attack does not work at all even with one round when a SMN is used or that your attack is a combination of a plaintext-ciphertext and "a known secret information" attack.

2. **You claim that your attack breaks 2.5 rounds of the cipher. However, there is a fundamental flaw in the setting when you split the round in half.**
   In particular, on page 5 you say: "From the specification of Pi-Cipher, the capacity part of the state is IScapacity =$IS_2||IS_4$, and the rate part of the state is ISrate = $IS_1||IS_3$."
   This is true, but it is valid **only for a full round, and not for a half round!** The designer's choice to have $IS_1||IS_3$ as an output/ISrate (i.e. the part that can be seen), is exactly to avoid the possibility to use IS_1 (O_1 in your description) and the constant C_1 (S_1 in your description) to go up in the Pi-function, and recover the previous state IS_1 (i.e. to recover D_1 in your description). If a full round is used, this is impossible, without a guess. In the specification, we have not claimed that the output of a half round is ISrate = $IS_1||IS_3$, so you can not assume that this is the output. A logical output of a half round is a symmetric one to the full round, i.e $IS_2||IS_4$, in which case you can not mount the attack 2.5 rounds, but only 1.5.
   Therefore, you can not claim that your attack works for 2.5 rounds, but it is valid to say 2 rounds. We expect that you will change this in your paper so that the reader is not wrongfully led to believe something that is simply not part of the specification. Another alternative is to declare in your paper that you have introduced additional deliberate alternation in the design, by exposing O_1 and O_3 as rate parts of a half round, while the designers of Pi-Cipher would choose O_2 and O_4.

3. **You should be very clear in your writing, which of the variant of the cipher you were able to attack.** More precisely:
   - Your attack reduces a security margin of the current v2 CAESAR lightweight variant of Pi-Cipher, Pi16-Cipher096 that has 3 rounds. It does not break it, nor does it invalidate the claims of the Pi-Cipher designers team. More precisely, with your findings you have a key recovery attack on the lightweight variant of pi16-Cipher096 **with 2 rounds (not 2.5) when a SMN is not used** and still it doesn't break the whole cipher.
   - Pi-Cipher in v2 of the CAESAR competition has 3 more variants, and your attack applied to those variants (even with the deliberate weakening in 2.5 rounds with O_1 and O_3 as rate parts) is significantly more expensive than a brute force attack. This is apparent demonstration of the robustness of the overall design of Pi-Cipher.
   - Your attack reduces the security margin of two v1 variants of CAESAR Pi-Cipher after their alternation to 2.5 rounds with O_1 and O_3 as rate parts, **under the condition that a SMN is not used**. It does not break them, nor does it invalidate the claims of the Pi-Cipher designers team.
   - As it is a common academic practice in all cryptographic papers where the findings and attacks do not invalidate the claims of the designers of the cipher – we would like to ask you to mention in the abstract of your paper that your findings do not invalidate the security claims of the Pi-Cipher team submitted to the CAESAR competition.
   - In your attack you are analyzing the lightweight variant of Pi-Cipher with a provocative low number of just 2 rounds proposed at the NIST Lightweight Cryptography Workshop 2015. This variant is not a part of any official proposal within the CAESAR documentation of Pi-Cipher. Note that last year on Nov 20, 2014, Pi-Cipher Team sent a variant to the CAESAR mailing list **even with only 1 round**. The reasons were to attract cryptographers to analyze Pi-Cipher. Here is an excerpt from that email:

     > *On 20/11/14 20:30, Hristina Mihajloska wrote:*
     > *Dear all,*
     >
     > *We've written software implementations for our π-Cipher with reduced rounds.*
     > *In order to make it simpler for the people who are doing cryptanalysis on our cipher, we decided to write an implementation (reference and lightly optimized not using SSE) with one round and two rounds of the 64-bit version of the π64-Cipher256 with no changes in the algorithm.*
     > *We would like to ask all the competitors if they are OK with our decision to include these implementations in the SUPERCOP.*
     > *…*
     >
     > *Best Regards,*
     > *π-Cipher Team*

Unfortunately, then we did not manage to attract attention by other cryptographers for analysis of Pi-Cipher v1 with just 1 or 2 rounds. Now, we are really happy that you were attracted to analyze that **non-CAESAR** lightweight proposal with just 2 rounds.

- We find very offensive your writing style of giving numerous statements on page 3 where you are mixing the official CAESAR proposals with a non-CAESAR lightweight variant in order to influence the official CAESAR decision Pi-Cipher to be eliminated. First you say:
  i. The attack is faster than exhaustive search of the key for the following variants (reduced to 2.5 rounds):
     Pi16-Cipher096 with 16-bit words and 96-bit key.
     This variant was proposed with 4 rounds in version 1, 3 rounds in version 2, and 2 rounds in the lightweight version.
     Pi16-Cipher128 with 16-bit words and 128-bit key.
     This variant was proposed with 4 rounds in version 1, and 2 rounds in the lightweight version.
     Pi32-Cipher256 with 32-bit words and 256-bit key.
     This variant was proposed with 4 rounds in version 1.

  ii. Then, after mixing official CAESAR and non-CAESAR variants of Pi-Cipher, a little bit further on the same page you are addressing the CAESAR competition with the following statement: "This kind of analysis is very important for the progress of the CAESAR competition, as the final portfolio of the selected authenticated ciphers should offer a high level of security. Thus, evaluating the security of the remaining candidates, leads to a more clear overview of which candidates are robust and which should be eliminated."
- Despite the facts that you haven't broken nor weakened any CAESAR variant of Pi-Cipher (neither from v1 nor from v2); despite the fact that your attacks applied on three weakened designs with 2.5 rounds and with O_1 and O_3 as rate parts are significantly more expensive than the brute force attack; despite the fact that additionally to the known plaintext attack scenario you use a "known secret information attack"; despite the other features by which Pi-Cipher is much more robust than AES-GCM, you are hand waving and falsely guiding the reader that Pi-Cipher should be eliminated from the CAESAR competition. We suggest you to change this part of your paper and to not lead the readers to draw conclusions that the designers of Pi-Cipher took risky decisions by lowering the number of rounds from 4 to 3 in the CAESAR v2 of Pi-Cipher.

As a conclusion, we find your results very interesting and we encourage you to publish them with the remarks that we are giving you in this note.
Also we would like to use this opportunity and to point out that right after DIAC 2015, we received a draft paper from the following team of researchers: Joseph Alley and Josef Pieprzyk from the School of Electrical Engineering and Computer Science - Queensland University of Technology, Brisbane, Australia. They had very similar attack strategy as yours on 2 rounds of Pi-Cipher, although we find that yours "guess and determine" attack is covering more details then theirs. As far as we know they still haven't decided to publish their results.

Best regards,
Pi-Cipher team