

Structural Weaknesses in the Open Smart Grid Protocol

Klaus Kursawe Christiane Peters*

European Network for Cyber Security
The Hague, The Netherlands,
{klaus.kursawe,christiane.peters}@encs.eu

Abstract

The Open Smart Grid Protocol (OSGP) is currently deployed in various countries in large-scale Smart Metering projects. The protocol was developed by the OSGP Alliance and published as a standard by the European Telecommunications Standards Institute (ETSI). We identify several security issues in the OSG Protocol, primarily the use of a weak digest function and the way the protocol utilizes the RC4 algorithm for encryption. A straight-forward oracle attack triggers the leakage of key material of the digest function. We outline how an attacker can make use of the simple protocol structure to send maliciously altered messages with valid authentication tags to the meters.

Keywords: protocol analysis, authentication protocol, OSGP, ISO/IEC 14908, Advanced Smart Metering, RC4, oracle attack, bit-flipping attack

1 Introduction

In the past years, an increasing amount of critical infrastructure is getting digitalized, adding communication capabilities as well as local computation to numerous devices in energy and water distribution networks, transport systems, and manufacturing. This is partially with the goal of efficiency increase, but often also a necessary requirement to handle the changing environment – a shrinking workforce will result in a lack of certified engineers with the ability to manually control processes, increasingly complex systems require increasing remote maintenance, and changes in consumer behavior – such as local power generation or a switch to electric vehicles in the case of electricity distribution – require intelligent management.

One very visible digitalization project is the switch from analog to digital (smart) meters, which is currently performed in many countries worldwide. In addition to more precise billing, a smart meter can also give input to grid control algorithms, be used for electricity markets, communicate with the smart home (for example, to tune down air conditioning and heating systems at times when energy demand is high), or to remotely disconnect a household. In this, a previously disconnected and uncritical device turns into a connected device that can generate process-critical data and perform active switching.

Robustness of data and switching commands is vital – if a large amount of households are switched off simultaneously, the excess energy can harm the overall grid. Similarly, if grid

*This work was supported by the European Commission FP7 Programme under Contract No. 607109 (SEGRID). Date: 2015.06.15

maintenance algorithms rely on smart meter measurement data, wrong input can have effects far beyond just billing fraud. This poses a new challenge for meter manufacturers: design devices that are cheap, widely distributed, work over very low bandwidth channels such as low frequency powerline communication (PLC).

In this paper, we investigate the example of the Open Smart Grid Protocol (OSGP) [8]. The OSGP was one of the first powerline communication protocols for smart meters on the market, and is widely used for communication between Smart meters and the data concentrator which is a device that collects data from potentially several hundred meters in a PLC segment and forwards this data to the head end.

The OSGP resides in the application layer of the protocol stack defined by the ISO/IEC 14908 standard [13], also called ‘EN 14908’. While the OSGP is primarily used for smart-metering applications, it was designed for a wider usage in smart grid devices. The protocol stack is very lightweight. This comes at the price of avoiding NIST recommended cryptographic primitives (e.g., the Advanced Encryption Standard (AES) in an authenticated mode) in favor of less computationally intense ones: the RC4 stream cipher for encryption and a non-standard digest function for message authentication.

We found that the biggest issue is the combination of a stream cipher with a linear digest function, opening the possibility for attacks on cryptographic keys and messages. One can for example exploit this by using the receiver of a message as an ‘oracle’. Given a properly encrypted and authenticated message, an attacker guesses a combination of a message bit and a key bit, sends a specially crafted message to the receiver, and uses the device response — either a rejection or an acceptance of the message — to confirm or reject the guess. As this can be done for individual bits, an attacker can reconstruct the main device key (OMA key) with less than 300 message responses.

Overall, we identified four potential weaknesses in the OSG Protocol:

- **Use of RC4.** Over the last years a number of weaknesses have been identified in the RC4 stream cipher. One of the issues is a correlation between the key and the RC4 keystream that can be used to reconstruct the key. This weakness was successfully exploited using known initialization vectors (IVs) to break the WEP (Wired Equivalent Protection) standard which makes heavy use of RC4. The attacks [16, 28] break WEP within seconds and attack tools such as Aircrack-ng [1] are freely available and a standard tool for penetration testers [15].

The use of RC4 in the OSGP is similar to the way RC4 was used in the WEP standard. While a new key is generated per OSGP message, only the first 8 bytes of this key are in fact different from each other, while the remaining 8 bytes are constant. Furthermore, the initial 8 bytes are a constant xored with a value that is publicly known. While the authors are not aware of a concrete attack on this, it does imply a strong correlation between the used keys, and supplies an attacker with rich data to analyze. Furthermore, the setting is rather close to the one used in the WEP standard, which means there is a large body of experience in the field on statistical key recovery.

Another issue with RC4 in the OSGP is that it is used with weak authentication. Given the nature of a stream cipher, an attacker who gets hold of a ciphertext can alter bits in chosen positions. Using the knowledge of the message space of the EN 14908 standard an attacker can send altered messages that are valid with high probability. Moreover, she can authenticate these messages exploiting the weaknesses of the digest function that are described in the following.

- **Weak digest function.** The digest function used in the OSG Protocol for message authentication is basically linear and is implemented in a way that not only disables authentication but can also decrypt an intercepted message in time linear in the message length. The structural weaknesses in the digest function allow for a ‘meet-in-the-middle’ attack (effectively halving the key size for a brute-force attack on the key), modification of messages in a controlled way, recomputing the correct digest without need for encryption- or authentication keys, and using a device’s NACK messages¹ on a wrong digest to reconstruct both authentication key and message.
- **Undefined broadcast security.** While the broadcast function is called ‘Secure Broadcast’, there is very little security specified for it. This is particularly worrying as this mechanism is used to send firmware updates, which are distributed using the secure broadcast mechanism without any specified measures to provide source authentication.
- **Key usage.** While the protocol does use session keys for encryption, it uses one master key for authentication. However, this authentication key is used to derive the encryption session keys. Thus if the authentication key is compromised, all session keys are known to the attacker. As the authentication key is used with a weak authentication algorithm, it is fairly exposed to a number of attacks, and a compromise of this key is possible.

1.1 Exploiting the structural weaknesses

These structural weaknesses can be exploited in various ways. The attacks we described in this paper are not overly sophisticated, and should rather serve as an illustration of how weak cryptographic building blocks can enable an attacker to potentially hijack a whole PLC segment. A relatively simple attack breaks the OSGP digest function in a way that an attacker can not only forge messages, but—given access to one sufficiently long authentic message—can reconstruct cryptographic keys with low effort.

We describe how to recover the key which then can be used to authenticate modified messages. Due to the usage of the RC4 stream cipher and the simple message structure it is possible to send false commands that are properly authenticated using the OSGP digest.

1.2 Recent work and our contribution

The results of our analysis have been made available to the manufacturer at the end of 2013, and mitigation strategies were discussed with the manufacturer and utilities in an attempt to have the vulnerabilities fixed before publication. Beginning of 2015, we were contacted by two research teams who had independently analyzed the OSGP and identified the same vulnerabilities.

Jovanovic and Neves cryptanalyzed the OSGP digest function in [14] and Feiten and Sauer cryptanalyzed the RC4 usage in OSGP in [9].

Jovanovic and Neves ([14]) analyze the OSGP digest under a variety of assumptions. Their focus lies on applying different cryptanalytic techniques and their analysis goes deeper than ours. However, they do not discuss how to realize their assumptions in a real protocol run as done in this paper. Feiten and Sauer ([9]) provide a detailed statistical analysis on the use of RC4 similarly to the WEP attacks; their analysis complements our analysis.

¹A message containing a **N**egative **A**CKnowledgement character to indicate an error.

We focus on the attacks that use existing and practically occurring message structures and take protocol limitations into account. We also specify how to obtain the necessary input for both attacks within normal protocol usage.

We are currently investigating how to implement our attacks. The obstacles are of engineering nature, in particular, in how to establish communication between devices in a PLC network. We note that the goal of our analysis is not to provide a readily usable attack tool, but to demonstrate the structural weaknesses of the used building blocks. These should be seen like cracks in a dam — a last warning sign that something needs to be fixed before the real damage has been done.

1.3 Organization of the paper

In Section 2 we describe the general setup of the OSG Protocol and identify first issues. In Section 3 we discuss the usage of the RC4 stream cipher in the OSGP and potential attacks. Section 4 introduces the OSGP digest function and notation. Section 5 outlines an attack on the digest function that uses the meter as an oracle to fully recover the authentication key. Section 6 discusses how to make meet-in-the-middle attacks effective using the structure of the digest function. Section 7 provides lessons learned and advice how to replace the weak primitives by recognized cryptographic primitives in the standard while respecting the constraints of the PLC environment.

2 Overview of the OSGP

In this section we describe the protocol according to the specifications in the “Group specification GS OSG 001: Open Smart Grid Protocol” by the ETSI Open Smart Grid Industry Specification Group [8]. We outline the problems of using RC4 and the weak digest function for encryption and authentication in the OSGP.

The EN 14908 standard [13] defines communication for layers 2 to 7 in control networks. Applications lie amongst others in the field of smart metering, specifically in the communication between a data concentrator and smart meters in one PLC segment. Such a segment contains meters in a whole neighborhood, possibly several hundred devices. As shown in Figure 1 meters can act as repeaters to forward information to the targeted devices. The data concentrator acts as master with connected proxies and meters as slaves.

2.1 Application layer security

The OSGP was specifically designed to secure the application layer in smart meter communications. EN 14908 was designed to secure the transport and session layers. Applications use the EN 14908 communication protocol at layers 2 to 6 to facilitate the lower-layer communication between connected nodes and use the OSGP at layer 7 to provide message authentication and encryption. The OSGP digest function is supposed to provide message authentication. For encryption the RC4 stream cipher is used. Encryption keys are generated using the EN 14908 ‘encryption function’ which we will discuss in the next paragraph.

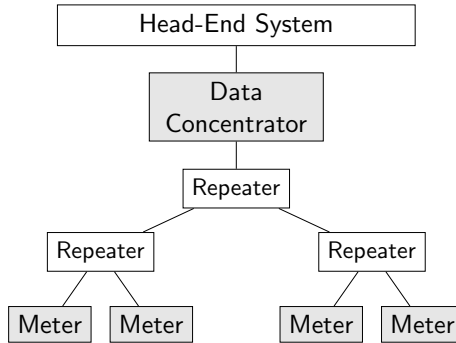


Figure 1: Head-end connected to PLC segment with data concentrator (DC), meters acting as repeaters, and meters as end nodes.

2.2 EN 14908 legacy

The EN 14908 defines security only at the transport layer. It specifies an ‘encryption function’ that is almost identical to the OSGP digest function. Section 11.12 in [13] admits that “The encryption algorithm [...] facilitates one way encoding rather than real encryption.” The randomness requirements in EN 14908 do not adhere to the state of the art in cryptographic randomness. Moreover, “Any 48-bit number is a valid encryption key.” In the sample implementation random numbers are generated using the `rand()` function. Needless to say, this function is not considered secure for cryptographic use, and should be replaced.

The main difference is that the OSGP digest function is initialized with an initial state of 0, while the EN 14908 encryption function receives a random value as a parameter to use for initial setup. Another difference is that the EN 14908 encryption function is used as a challenge-response protocol. Here the challenger sends a message and a random value, and receives back an encrypted and authenticated response message. This makes an attack rather easy, as the challenger has full freedom to choose which messages the server encrypts, and can then investigate the outcome. The challenger would then reuse the same randomness with an altered message as will be described in Section 5, and thus recover the keys from there.

2.3 Message structure

The OSGP inherits the message structure from EN 14908. The first one or two bytes of the APDU (Application Protocol Data Unit) define the destination and type of the message. The remaining bytes form the (possibly encrypted) message bytes and a digest.

Commands are organized in tables that are stored on the devices. A data concentrator initiates the interaction with connected meters by sending the addresses of the commands. The meter looks up the address in the corresponding table, executes the command associated with the table entry, and writes the requested results in a transaction-response table which is then sent back to the data concentrator.

Section 9.2 of the OSGP standard limits the buffer to a total of 114 bytes.

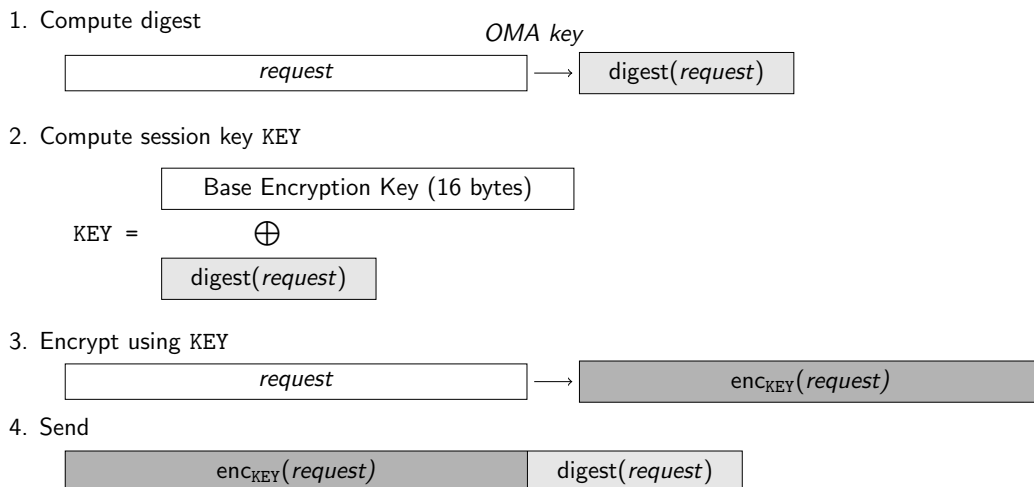


Figure 2: **Encryption and authentication process.** First the digest is computed on the request. Then encryption key is derived from the digest by xoring the digest to the first 8 bytes of the Base Encryption Key. Then the request is encrypted under this key. Ciphertext and digest are concatenated to form the message.

2.4 Device keys

Each OSGP device is manufactured with a unique symmetric key, called the Open Media Access Key (OMA key). This key can and is typically updated by the data concentrator when the device is placed in the field. This 96-bit OMA key is used for authentication of OSGP messages. It is also used for authentication of certain EN 14908 messages. Moreover, the OSGP derives encryption keys from the OMA key.

For encryption a 128-bit ‘Base Encryption Key’ (BEK) is set up as follows: the device encrypts the two fixed patterns 81 3F 52 9A 7B E3 89 BA and 72 B0 91 8D 44 05 AA 57, using the EN 14908 encryption function using the current OMA key. The resulting two digests of each 8 bytes are concatenated to form the BEK. If a malicious server manages to construct a request of the form of those two patterns it can challenge the sender with 0 and thus get half of the bytes of the base key as a response. Also note that the entropy of this 128-bit key are only 96 bits derived from the OMA key.

2.5 Encryption and decryption of OSGP messages

The general approach used in the OSGP for message authentication of encrypted messages is “encrypt-and-authenticate”, a method which has been shown to be generically insecure [17]. However, encryption and authentication are not completely independent operations. As shown in Figure 2, first a digest is computed on the plaintext message; then the plaintext is encrypted using the RC4 stream cipher with a session key derived from the digest. In particular, the sender sets up a session key by xoring the 8-byte digest of the request message to the first 8 bytes of the BEK of the receiving device. The receiver takes the digest of the request message and computes the session key in the same way. We will discuss in the following how this construction shows substantial structural similarities to the way RC4 keys were set up in the WEP standard, and how this might allow for attacks similar to the ones WEP is already

experiencing.

2.6 Updating keys

The EN 14908 uses an *increment key* function that seems rather weak, adding (without carry) a 6-byte key to the current key; a better idea would be to take the old and the new key into a hash function to derive the actual new key. In either case, once a key is compromised, or partial information about the key leaks, updating the key does not help at all.

2.7 Secure broadcast and firmware updates

The specification is somewhat unclear on the protection properties of the secure broadcast protocol; no authentication or encryption is defined in the specifications, and the basic protocols are not directly applicable, as they (correctly) require every message to add the receiver identity into the authenticated part.

In either case, there appears to be some risk here. As every OSGP enabled meter may act as a repeater, it is relatively easy for an attacker to get legitimate access to the broadcast medium, and it appears reasonably easy to obtain the valid broadcast sequence number. As the *secure broadcast* is also used for firmware updates, this can lead to a somewhat dangerous situation given that the standard does not define source authentication by means of digital signatures.

3 RC4 in the OSGP

The RC4 stream cipher was developed by Rivest at the end of the 1980s. The first biases in the keystream were identified in the 1990s [24] and since then cryptanalysts have published a stream of papers identifying flaws in the RC4 design ([11, 10, 19, 18, 25, 26, 29]). In the last years more and more real-world attacks on applications using RC4 for encryption were carried out successfully, showing that RC4 is in desperate need of replacement ([28, 2, 23]). An overview of the RC4 weaknesses can be found in [12]. Large companies such as Microsoft and Cisco strongly advise against the usage of RC4 and are currently phasing out code containing RC4 ([21, 4]).

The OSGP uses RC4 to encrypt certain messages. Given that not all messages are encrypted, one can assume that the encrypted messages contain critical commands and sensitive data.

3.1 Bit-flipping attack on OSGP messages

Bit-flipping on OSGP-encrypted messages is probably the most dangerous attack in the absence of strong authentication. Due to the simple message structure, the attacker can easily guess which part of the ciphertext corresponds to table entries which will be used to execute commands on the receiving device. Even if the attacker does not know which command will be executed instead, it is possible to disturb the communication and possibly execute commands that disrupt communication at all. To compensate the error caused by modifying one message byte, an attacker needs to change only two other message bytes and will succeed with reasonably high probability in generating a valid message. In principle, this could maliciously change a read-out message into a remote disconnect.

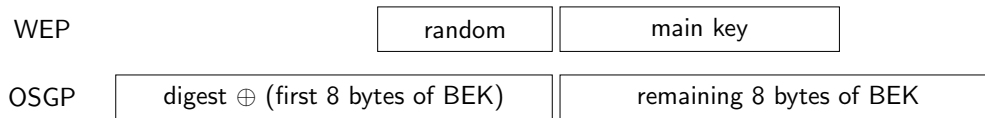


Figure 3: **Initializing vectors in WEP and OSGP.**

3.2 Related-key attacks

The RC4 stream cipher does not specify an initialization vector (IV). The common workaround is to use parts of the key as IV, e.g., to set a temporary key as `<init-vector|main key>`. This however is bad practice. An attacker who observes a certain number of IVs can reconstruct keys using correlated biases in the keystream. This was first observed by Roos [24] and formalized as attack by Fluhrer, Mantin, and Shamir [10], the famous ‘FMS attack’. The common workaround to thwart the FMS attack is to drop the first n bytes of the keystream where n ranges between 256 and 3,072. Klein [16] however showed that dropping 256 bytes is not enough and that a much smaller number of observed IVs is sufficient to recover the key. Klein’s attack was put into practice and practically broke the WEP standard. More subsequent work followed and most recently Ohigashi et al. [22] showed that “RC4-drop(n)” is insecure for even larger values of n .

3.3 Attacking RC4 in the OSGP

The usage of the RC4 in the OSGP is very similar to the setup of RC4 in WEP: the temporary key in the OSGP is a random-looking 8-byte value (the digest of the request) xored to the first 8 bytes of the main key followed by the remaining 8 bytes of that key; see Figure 3. While this construction is slightly longer (in WEP the IV consists of 3 random bytes and a key of 5 bytes), the OSGP model is weaker given that the request digest is known. Furthermore, the OSGP allows a receiver to manipulate the sequence number, and thus have a control on how the initialization vector changes. Klein’s related-key attack on WEP reconstructs the key after observing only 13,000 sessions in WEP, i.e., different keystreams with different initialization vectors.

Breaking the RC4 encryption in WEP can be carried out within seconds, as it is easy to observe a large number of sessions by means of packet spoofing. This is harder in the OSGP case, as there is no straightforward mechanism for an attacker to force the targeted system to generate encrypted messages.

However, in the OSG Protocol a receiver is supposed to reject a packet based on a wrong sequence number. The receiver not only rejects but also appends the correct 32-bit sequence number to the NACK message. Thus, an attacker who can fake those NACKs is able to arbitrarily generate messages with chosen sequence numbers, which then can be used to derive the key. In particular, the attacker can try out related keys differing in a few well-chosen bits.

Note that a more detailed statistical analysis of RC4 in the OSGP will be released soon in [9].

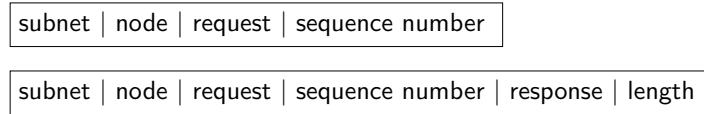


Figure 4: **OSGP digest input.** The digest is computed over the destination address (subnet, node), the request, the 32-bit sequence number, and in case of the response additionally over the response and its length.

3.4 More RC4 victims

Among the victims of recent attacks exploiting biases in the RC4 keystream were TLS cipher suites using RC4 for encryption and the HIVE hidden volume encryption system ([2, 23]). It should be noted that there is no easy patch against the attack of AlFardan et al. [2] such as for the more implementation dependent attacks such as BEAST [5] and Lucky 13 [3], which are attacks against padding errors in CBC mode. It is now generally recommended to avoid usage and phase out RC4 where it is still used.

4 The OSGP digest function

This section introduces notation and a first brief analysis of the OSGP digest function and serves as basis for the attacks outlined in the following section.

4.1 Usage of the digest function

This function computes an authentication tag called *digest* on the (unencrypted) message. More specifically, as depicted in Figure 4 the digests are computed on the destination address (subnet address and node), the message bytes (payload), an appended sequence number, and in case it is response message also on the response and its the response length.

4.2 Linearity of the digest function

The OSGP digest function outputs 8 bytes of digest that are generated by processing the message byte by byte while going linearly through all key bits. In fact, the digest function turns the 96-bit OMA key into a 144-bit key: after processing all bits of the OMA key the function goes back to the first key bit and processes the first 48 bits once more. While this does not add entropy—or additional effort, as any attacker that has the last 96 bits immediately has the remaining 48—it adds 48 additional rounds to the digest function.

4.3 Notation

For the ease of presentation we introduce the following notation. Where possible we try to stay close to the algorithm code given in "Annex E" of [8] computing the digest output in place.

i, j : Parameter i indexes *bytes*, parameter j indexes *bits*.

$k_i[j]$: This value denotes *bit* j of key byte i .

ℓ : We denote the number of *rounds* by ℓ . In total there are 144 rounds, exactly one per key bit.

$m[\ell]$: The value $m[\ell]$ holds the message *byte* processed in round ℓ , also called the *active message byte*. Given that there are more rounds than message bytes $m[\ell]$ is set to 0 once the algorithm has processed all message bytes.

$m_{i,j}$: When discussing the implementation it is useful to see the correspondence between the active key bit and the active message byte. Key bit $k_i[j]$ corresponds to the message byte $m[8 \cdot i + j]$ which we also denote by $m_{i,j}$.

$S_{i,j}$: Internal state byte corresponding to key bit $k_i[j]$. The eight state bytes corresponding to the bits of the last key byte form the final digest value. For $j = 0$, $S_{i,j-1}$ is actually $S_{i-1,7}$. In round 0, $S_{i-1,j}$ equals 0 for all j .

$C_b(X)$: The function updates the active state byte. There are two possible outcomes depending on input bit b . Depending on key bit b , the result is plus or minus $\neg X$ shifted cyclically one bit to the left or to the right.

4.4 Updating the state

The implementation uses an array of 8 bytes for generating the digest. The algorithm cycles through those bytes and updates them one by one according to the currently active message and key bytes. For our analysis we use a linear description of the state to indicate how state byte $S_{i,j}$ is modified with input from key bit $k_i[j]$ and message byte $m_{i,j}$.

The code of the reference implementation in [8] in condensed write-up using j rather than $7 - j$, $k[i][j]$ for key bit $k_i[j]$, $m[\cdot]$ for the active message byte, and $digest[j]$ for the variable holding state byte $S_{i,j}$, is given as

```
n = ~(digest[j]+j);
if (k[i][j])
    digest[j]= ((n<<1)+(n>>7));
else
    digest[j]= -((n>>1)+(n<<7));
digest[j]+= digest[(j-1)%8] + m[8*i+j];
```

In other words, the digest function adds the index j to the state $S_{i-1,j}$, negates the outcome, and stores the result in the temporary variable n . If key bit $k_i[j]$ is 1, shift the bits of n cyclically to the left and write the result to digest byte $S_{i,j}$. If key bit $k_i[j]$ is 0, shift the bits of n cyclically to the right and write the negated result to $S_{i,j}$. This is the only part depending on (a single bit of) the key. Finally, add state $S_{i,j-1}$ and the active message byte to $S_{i,j}$.

In our notation the OSGP digest function looks as follows:

```
for i = 0 to 17
  for j = 0 to 7
     $S_{i,j} = C_{k_i[j]}(\neg(S_{i-1,j} + j)) + S_{i,j-1} + m_{i,j}$ 
```

4.5 Observations on the update function $C_b(X)$

The function is linear in every way. In byte arithmetic, bit-wise negation of a byte X , i.e., computing $\sim X$, is the same as computing the value $-X-1$.

Due to the linearity of the $C_b(X)$ function, adding a '1' to the input X of $C_{k_i[j]}$ results in adding 2 or subtracting 128 to the output of the shift function. This is easily seen since when rotating value '1', the outcome equals

$$\begin{aligned}\text{"rotate 1 left"} &= (00000010) = 2 \text{ or} \\ \text{"rotate 1 right"} &= (10000000) = 128.\end{aligned}$$

5 The Meter as Oracle: reconstruction of Key and Message

This section outlines how we can decrypt an encrypted message and reconstruct the OMA key used for authentication. We first present the general idea and then sketch how to recover the authentication key by showing explicitly how the last key byte can be derived from the last 9 rounds of the digest computation.

There are two ways how we can probe the system to reveal information about an encrypted message and the authentication key: using the receiving device as an oracle and using chosen messages.

5.1 Receiver as an Oracle

Upon receiving an encrypted and authenticated message the receiver will try to verify the digest. If the digest does not match an error message is returned. The device will return a response message including a negative-acknowledge character (NACK), that is a fixed error code of length 1, and a digest computed on the response message (see Figure 5). In the OSG Protocol, this NACK message is explicitly never encrypted, turning the receiver into an oracle. Exploiting this, we make a controlled modification of the message and the digest, and use the receiver to test if the modification we made to the digest matches the modification we made to the message.

For encrypted messages there are four message/key-bit combinations given that we have to guess the plaintext bit (two options starting from the ciphertext) and two options with respect to the modification of the digest. Each oracle query will give us one bit of information about the key or the message. We can reconstruct the entire OMA key with at most $96 \cdot 3$ queries given that we need at most three guesses to check for all four message/key-bit combinations. As the message encryption key is derived from the digest of the first message in a conversation, this message is not suitable for this attack; any modification of the digest would change the key, and thus got the message to decrypt wrongly (leading to a surely wrong digest). The response to the request however is independent of its own message digest, and thus can be used directly.

5.2 Chosen message attacks

If we can manipulate the sender to send a specific message, we can send two similar messages and compare the digests. This also gives us one bit of information per message pair, although a more elaborate attack is feasible where more bits are recovered at once. To some extent,

type	m	digest(m)
type	m NACK 1	digest(m NACK 1)

Figure 5: **Oracle attack.** On sending a request message with digest, the response of the meter contains the request, the one-byte error code, its length, and a digest computed on these partially known values.

message manipulation is easy. If the receiver sends the error message ‘wrong sequence number’, the sender will send the new message with the new sequence number, giving an attacker the chance to have the same message resend arbitrarily often, with the bytes representing the sequence number being of the attackers choice. Other attacks may be possible on an application level, if the protocol is used as a carrier for a more user defined protocol. Chosen-message attacks were also applied in [14]; however, it is not entirely clear how to realize their strong assumptions on the message structure in a real protocol execution. A potential way to realize the attacks might be given by the way the receiver handles sequence numbers as described in the previous section. This would allow to fix 32 bits of the message.

5.3 Reconstruction the key from the back

To outline how to decrypt a message and to reconstruct the current OMA key bit by bit we analyze the final 9 rounds of the digest function.

We first start analyzing a plaintext message with a digest and then extend the attack to an encrypted message.

Using the notation introduced in Section 4, the last 9 rounds of the protocol look as follows:

$$\begin{aligned}
S_{16,7} &= C_{k_{16}[7]}(S_{15,7} + 7) + S_{16,6} + m_{16,7} \\
S_{17,0} &= C_{k_{17}[0]}(S_{16,0} + 0) + S_{16,7} + m_{17,0} \\
S_{17,1} &= C_{k_{17}[1]}(S_{16,1} + 1) + S_{17,0} + m_{17,1} \\
S_{17,2} &= C_{k_{17}[2]}(S_{16,2} + 2) + S_{17,1} + m_{17,2} \\
S_{17,3} &= C_{k_{17}[3]}(S_{16,3} + 3) + S_{17,2} + m_{17,3} \\
S_{17,4} &= C_{k_{17}[4]}(S_{16,4} + 4) + S_{17,3} + m_{17,4} \\
S_{17,5} &= C_{k_{17}[5]}(S_{16,5} + 5) + S_{17,4} + m_{17,5} \\
S_{17,6} &= C_{k_{17}[6]}(S_{16,6} + 6) + S_{17,5} + m_{17,6} \\
S_{17,7} &= C_{k_{17}[7]}(S_{16,7} + 7) + S_{17,6} + m_{17,7}
\end{aligned}$$

To determine the last key bit $k_{17}[7]$, we modify a bit in message $m_{16,7}$. Suppose for now we know how to add the value ‘1’ to this message. This means that $S_{16,7}$ can be increased by 1, as then is $S_{17,6}$. This yields two possibilities for $S_{17,7}$: in the shift function $C_{k_{17}[7]}$ the modified bit is either shifted to the right and added (add 128), or shifted to the left and subtracted (subtract 2).

Thus, the correct value for the new $S_{17,7}$ is either the original value plus 128, or the original value minus 2. Using either a chosen message attack (let the sender authenticate two messages that differ in $m_{17,7}$), or using the meter as an oracle (adding 1 to $m_{17,7}$ and testing

which of the two possible digests the receiver accepts), we can verify which choice was taken in $C_{k_{17}[7]}$.

This way, we have now determined one key bit. With this key bit fixed, we can repeat the same method modifying $m_{16,6}$ and determining the choice of $CS_{17,6}$, thus successively determining all key bits of the authentication key.

When dealing with an encrypted message, we cannot be sure whether flipping the last bit adds or subtracts 1. This, however, only doubles the number of possibilities — the options are now +128, -128, +2 or -2. By testing which of the four options hold, we not only get the key bit, but implicitly also decrypt one message bit of $m_{11,7}$. The others can be decrypted the same way by adding/subtracting 2, 4, 8,

5.4 OSGP aspects in key recovery

There is a minor complication in that the message sent to the requester is not the same message that the digest is computed on. Rather, it is the ‘response’ part of the message, with other fields such as subnet, node, request, and sequence number. For simplicity, we assume we have a backward compatible device for which we do not need to include the response length. This way, the response sent back to the requester consists of exactly the last bytes of the message used for the digest.

For longer messages the first and last parts are digested with the same key bytes. In this case, we can directly reconstruct the key using above method. For shorter messages, the outlined method only allows us to recover some of the key bits. However, we do still profit from the fact that some key bits are used twice, as this means we do not need to go back to the beginning of the message to get the OMA key.

In summary, we should expect to be able to get most key bits with this method, and then get the remaining ones by different means (e.g., through a brute force attack). As the attack on the last 8 bytes, this attack only works on the response message, and cannot work directly on the request.

To fully recover the key with the outlined method we would need messages with 144 bytes. Otherwise, the receiver of a message will fill in the trailing zeros themselves, and thus not allow the attacker to manipulate them. If such a message does not exist due to the length restrictions on request and response tables, this can be achieved by using a receiver with no message length byte, which will accept a message of every length².

We note that the idea of oracle attacks is certainly not new. Interestingly, many mistakes made in the OSGP were already identified for the WEP protocol: Tews and Beck [27] were able to decrypt messages in WEP by using an oracle attack which they named “Chopchop”. Similar to the attack here, the Chopchop attack does not exploit RC4 but uses the NACKs returned by an access point to decrypt a message. The Chopchop attacker modifies the message starting from the last byte and adapting the CRC32 checksum.

5.5 Reconstruction the key from the front

As noted in the discussion, the above oracle attack reconstructs the key from back to front and therefore needs to start from messages that have as many bytes as there are key bits,

²It is not clear in the specification how exactly the message length byte is used, so it might be possible to do this for every receiver.

starting with the last key bit corresponding to the final message byte and working its way to the front.

It is however possible to reconstruct the key from the first key bits that depend on the first message bytes. Rather than making a modification on the message, then on the digest and then to check the result with the meter as oracle, one can use a ciphertext-only attack: The attacker takes an intercepted message and introduces an error in message byte $m[0]$, say setting $m[0]$ to $m[0] + 1$, then corrects this by setting $m[1]$ to $m[1] - 1$, so the state byte $S_{0,1}$ is unaffected. To correct the error in state byte $S_{1,7}$ the attacker has to modify $m[8]$. There are two options to test: either $m[8]$ increases by 2 if $k_1[7]$ is 1 or it is decreased by 128 if $k_1[7]$ is 0. This turns into an oracle as follows:

```
m[i]    = (m[i]+1) %256
m[i+1]  = (m[i+1]-1) %256
m[i+8]  = (m[i+8]+2) %256
if oracle(m,digest):
    return 1
else return 0
```

This has been also shown with differential cryptanalysis in [14]. Our notation is simpler as it treats message bytes as integers. This can occasionally go wrong when dealing with carries. However, our experiments have shown that this simple attack almost always works. One could also simply use another attack message.

Another advantage of the forward decryption is that it is a ciphertext-only attack where the digest stays the same under message modifications. In particular, the message will decrypt properly as the encryption key is derived from the digest.

In summary, the forward reconstruction recovers the full 96-bit key with only 96+8 payload bytes which is well under the buffer restriction of 113 bytes and is thus a practical attack breaking the authentication and in particular retrieving the main device key of the OSGP meter. If for some reason it is not feasible to use messages of 104 payload bytes, the attacker can also use shorter messages to recover the first $96 - x$ key bits and use brute force to guess the remaining x key bits.

6 Meet-in-the-middle attack

If an attacker gets hold of one plaintext message and one corresponding digest, she can test all possible authentication keys to find the correct one (though it may be that several keys fit for a specific message, so in reality an attacker would probably need two message-digest pairs). In theory, given a 96-bit key, the effort for such a calculation should be 2^{96} computations.

However, the digest function is linear in the sense that every execution of the internal loop depends on exactly one key bit and the internal state. Thus, the active round is independent of all key bits that have not been processed up to that round. The attacker can easily compute in both directions. To this end, the attacker guesses the last bits of the key, and runs the algorithm backwards for those bits. She then creates a table with all the internal states for all combinations of trailing key bits. Then, she runs the algorithm forward with the leading key bits, and verifies if the resulting internal state is a value in the state table.

In theory, this way one can compute from both sides using 2^{48} bit operations, giving an overall effort of 2^{49} bit operations. As we need to keep a table of the computed values in at least one direction, that would require a memory of $2^{48} \cdot 8$ bytes plus overhead, i.e., 2048 Terabyte of memory³. A slightly asymmetric search —for example, a table of 2^{36} digest entries will require only 512 gigabyte of fast memory, requiring 2^{59} bit operations to process, which is well in range for a moderately advanced attacker.

We note that many more optimizations can be applied here. For the purpose of the preparation it is sufficient to notice that the algorithm easily computes in both directions.

6.1 OSGP aspects in meet-in-the-middle attacks

This attack does need a reasonably advanced attacker, mostly because the outlined attacks require a comparatively high amount of resources. In terms of getting access to a message and the digest, it seems relatively straightforward —the digest is communicated in the clear, and the communication is predictable enough that an attacker should be able to predict some of the plaintext messages (there is no randomness in any of the messages, so simple knowledge of the protocols will be sufficient here).

We note that many more optimizations and combinations with other attack methodologies can be applied here making attacks much more efficient. We refer to the independently developed techniques in [14].

6.2 Meet-in-the-middle attacks for the EN14908 encryption key

In the original EN 14908 protocol, the encryption function uses only a 48-bit key. As the "meet-in-the-middle" attack still works, this reduces the effective key length to 24 bit, which can be brute forced within seconds and requires only 4 megabytes of memory for the tables. It is also easy for an attacker to obtain a plaintext/ciphertext pair, as this is exactly what the challenge-response pair is. In OSGP, a 96-bit version of EN 14908 is suggested, which is still in reach of an advanced attacker, but does need sizable resources to break.

6.3 Exploiting invertibility of the digest for simple brute-force attacks

Because of the linear character of the digest function, it can be calculated backwards and forwards. In addition to the meet-in-the-middle attack, this also allows a substantial speedup in brute-forcing keys. The attacker does not need to compute the entire digest with a new key each time (which takes 144 iterations). Instead, she can take an existing digest, backtrack the calculation of the last key bit, and then compute the digest of the next key with one bit difference with only 2 computation steps. Overall, the computing required to brute force n bits of the key is equivalent to the traversal of a binary tree with depth n . The efficiency of this is $2 \cdot 2^{n+1}$ steps of computation, as opposed to $2^n \cdot n$ for normal brute force search.

7 Conclusion

We identified a number of vulnerabilities in the OSG Protocol, some of which can be used to completely recover the authentication- and encryption keys. The outlined attacks are just

³The calculations are based on the interpretation of 'kilo' as 1,024. Thus 2^{10} bytes are one kilobyte, 2^{20} bytes are one megabyte, 2^{30} bytes are one gigabyte, and 2^{40} bytes are one terabyte.

examples how to exploit the inherent weaknesses of the digest function and the RC4 stream cipher.

The encryption and message-authentication mechanisms were chosen to be lightweight, but not to be secure. The OSGP digest function does not adhere to any crypto standard, deviating from any best practice in the design of a secure message authentication code (MAC). The digest function is weak in the sense that each key bit affects only one message byte directly; this property allows both forward and backward computation. In addition to replacing the digest function, adding some randomness to the message text may make it harder for an attacker to guess a plaintext message. Also RC4 should not be used due to the known issues. While RC4 is probably friendlier to non-volatile memory it also needs a lengthy initialization to reduce the risk of biases in the key stream.

We believe that the only clean way to resolve the security issues in the OSGP is to completely replace the cryptographic algorithms involved, especially the digest function.

Moreover, the standard does not seem to provide entity authentication that is assurance on the source of messages. Broadcasts are ‘authenticated’ using the weak digest function which provides (at best) message authentication but not source authentication as would be done by appending a digital signature. Thus, the absence of public-key cryptography to secure broadcasted firmware poses a very high risk.

7.1 Lessons to take away

A standard such as the OSGP should use state-of-the-art cryptographic primitives. In particular, algorithm agility is crucial as even those cryptographic primitives might expire during the lifetime of those devices.

The state-of-the-art solution would be to use an authenticated cipher such as AES-CCM [6] or AES-GCM [20, 7]. This would also simplify the problem in OSGP with setting up encryption and authentication keys.

Acknowledgment

The authors would like to thank the anonymous reviewers for their helpful remarks and suggestions.

The authors would like to thank Philipp Jovanovic and Samuel Neves for fruitful and enjoyable discussions.

The authors would like to thank Echelon/NES for many insightful discussions on the security of OSGP smart meters.

References

- [1] Aircrack-ng. <http://www.aircrack-ng.org/>. (Cited on page 2).
- [2] Nadhem J. AlFardan, Daniel J. Bernstein, Kenneth G. Paterson, Bertram Poettering, and Jacob C. N. Schuldt. On the security of RC4 in TLS. In Samuel T. King, editor, *Proceedings of the 22th USENIX Security Symposium, Washington, DC, USA, August 14-16, 2013*, pages 305–320. USENIX Association, 2013. (Cited on pages 7 and 9).

- [3] Nadhem J. AlFardan and Kenneth G. Paterson. Lucky thirteen: Breaking the TLS and DTLS record protocols. In *2013 IEEE Symposium on Security and Privacy, SP 2013, Berkeley, CA, USA, May 19-22, 2013*, pages 526–540. IEEE Computer Society, 2013. (Cited on page 9).
- [4] Cisco. Next Generation Encryption, 2014. Published April 2012, last updated April 2014. http://www.cisco.com/web/about/security/intelligence/nextgen_crypto.html. (Cited on page 7).
- [5] Thai Duong and Juliano Rizzo. Here come the XOR Ninjas, 2011. Unpublished manuscript. (Cited on page 9).
- [6] Morris J. Dworkin. SP 800-38C. Recommendation for Block Cipher Modes of Operation: The CCM Mode for Authentication and Confidentiality. Technical report, Gaithersburg, MD, United States, 2004. (Cited on page 16).
- [7] Morris J. Dworkin. SP 800-38D. Recommendation for Block Cipher Modes of Operation: Galois/Counter Mode (GCM) and GMAC. Technical report, Gaithersburg, MD, United States, 2007. (Cited on page 16).
- [8] European Telecommunications Standards Institute (ETSI). Group specification GS OSG 001: Open Smart Grid Protocol, January 2012. (Cited on pages 2, 4, 9, and 10).
- [9] Linus Feiten and Matthias Sauer. Extracting the RC4 secret key of the Open Smart Grid Protocol with a known-cipher-stream attack. 2015. Unpublished. Received in a private communication. (Cited on pages 3 and 8).
- [10] Scott R. Fluhrer, Itsik Mantin, and Adi Shamir. Weaknesses in the key scheduling algorithm of RC4. In Serge Vaudenay and Amr M. Youssef, editors, *Selected Areas in Cryptography, 8th Annual International Workshop, SAC 2001 Toronto, Ontario, Canada, August 16-17, 2001, Revised Papers*, volume 2259 of *Lecture Notes in Computer Science*, pages 1–24. Springer, 2001. (Cited on pages 7 and 8).
- [11] Scott R. Fluhrer and David A. McGrew. Statistical analysis of the alleged RC4 keystream generator. In Bruce Schneier, editor, *Fast Software Encryption, 7th International Workshop, FSE 2000, New York, NY, USA, April 10-12, 2000, Proceedings*, volume 1978 of *Lecture Notes in Computer Science*, pages 19–30. Springer, 2000. (Cited on page 7).
- [12] Sourav Sen Gupta, Subhamoy Maitra, Goutam Paul, and Santanu Sarkar. (Non-) Random Sequences from (Non-) Random Permutations – Analysis of RC4 Stream Cipher. *J. Cryptology*, 27(1):67–108, 2014. (Cited on page 7).
- [13] International Organization for Standardization. ISO/IEC 14908-1:2012:Information technology – Control network protocol – Part 1: Protocol stack, 2012. (Cited on pages 2, 4, and 5).
- [14] Philipp Jovanovic and Samuel Neves. Dumb Crypto in Smart Grids: Practical Cryptanalysis of the Open Smart Grid Protocol. Conference Proceedings of Fast Software Encryption, FSE 2015, 2015. Cryptology ePrint Archive, Report 2015/428. (Cited on pages 3, 12, 14, and 15).

- [15] Kali Linux. <https://www.kali.org/>. (Cited on page 2).
- [16] Andreas Klein. Attacks on the RC4 stream cipher. *Des. Codes Cryptography*, 48(3):269–286, 2008. (Cited on pages 2 and 8).
- [17] Hugo Krawczyk. The Order of Encryption and Authentication for Protecting Communications (or: How Secure Is SSL?). In Joe Kilian, editor, *Advances in Cryptology - CRYPTO 2001, 21st Annual International Cryptology Conference, Santa Barbara, California, USA, August 19-23, 2001, Proceedings*, volume 2139 of *Lecture Notes in Computer Science*, pages 310–331. Springer, 2001. (Cited on page 6).
- [18] Subhamoy Maitra, Goutam Paul, and Sourav Sengupta. Attack on broadcast RC4 revisited. In Antoine Joux, editor, *Fast Software Encryption - 18th International Workshop, FSE 2011, Lyngby, Denmark, February 13-16, 2011, Revised Selected Papers*, volume 6733 of *Lecture Notes in Computer Science*, pages 199–217. Springer, 2011. (Cited on page 7).
- [19] Itsik Mantin and Adi Shamir. A practical attack on broadcast RC4. In Mitsuru Matsui, editor, *Fast Software Encryption, 8th International Workshop, FSE 2001 Yokohama, Japan, April 2-4, 2001, Revised Papers*, volume 2355 of *Lecture Notes in Computer Science*, pages 152–164. Springer, 2001. (Cited on page 7).
- [20] David A. McGrew and John Viega. The Galois/Counter Mode of operation (GCM), 2004. Submission to NIST Modes of Operation process. (Cited on page 16).
- [21] Microsoft. Microsoft Security Advisory 2868725, 2013. Published 12 November 2013. <https://technet.microsoft.com/en-us/library/security/2868725.aspx> and <http://blogs.technet.com/b/srd/archive/2013/11/12/security-advisory-2868725-recommendation-to-disable-rc4.aspx>. (Cited on page 7).
- [22] Toshihiro Ohigashi, Takanori Isobe, Yuhei Watanabe, and Masakatu Morii. How to recover any byte of plaintext on RC4. In Tanja Lange, Kristin E. Lauter, and Petr Lisoněk, editors, *Selected Areas in Cryptography - SAC 2013 - 20th International Conference, Burnaby, BC, Canada, August 14-16, 2013, Revised Selected Papers*, volume 8282 of *Lecture Notes in Computer Science*, pages 155–173. Springer, 2013. (Cited on page 8).
- [23] Kenneth G. Paterson and Mario Strefer. A practical attack against the use of RC4 in the HIVE hidden volume encryption system. In Feng Bao, Steven Miller, Jianying Zhou, and Gail-Joon Ahn, editors, *Proceedings of the 10th ACM Symposium on Information, Computer and Communications Security, ASIA CCS '15, Singapore, April 14-17, 2015*, pages 475–482. ACM, 2015. (Cited on pages 7 and 9).
- [24] Andrew Roos. A Class of Weak Keys in the RC4 Stream Cipher, 1995. Two posts in sci.crypt, message-id 43u1eh\$1j3@hermes.is.co.za and 44ebge\$11f@hermes.is.co.za. (Cited on pages 7 and 8).
- [25] Pouyan Sepehrdad, Serge Vaudenay, and Martin Vuagnoux. Discovery and exploitation of new biases in RC4. In Alex Biryukov, Guang Gong, and Douglas R. Stinson, editors,

- Selected Areas in Cryptography - 17th International Workshop, SAC 2010, Waterloo, Ontario, Canada, August 12-13, 2010, Revised Selected Papers*, volume 6544 of *Lecture Notes in Computer Science*, pages 74–91. Springer, 2010. (Cited on page 7).
- [26] Pouyan Sepehrdad, Serge Vaudenay, and Martin Vuagnoux. Statistical Attack on RC4 - Distinguishing WPA. In Kenneth G. Paterson, editor, *Advances in Cryptology - EURO-CRYPT 2011 - 30th Annual International Conference on the Theory and Applications of Cryptographic Techniques, Tallinn, Estonia, May 15-19, 2011. Proceedings*, volume 6632 of *Lecture Notes in Computer Science*, pages 343–363. Springer, 2011. (Cited on page 7).
- [27] Erik Tews and Martin Beck. Practical attacks against WEP and WPA. In David A. Basin, Srdjan Capkun, and Wenke Lee, editors, *Proceedings of the Second ACM Conference on Wireless Network Security, WISEC 2009, Zurich, Switzerland, March 16-19, 2009*, pages 79–86. ACM, 2009. (Cited on page 13).
- [28] Erik Tews, Ralf-Philipp Weinmann, and Andrei Pyshkin. Breaking 104 bit WEP in less than 60 seconds. In Sehun Kim, Moti Yung, and Hyung-Woo Lee, editors, *Information Security Applications, 8th International Workshop, WISA 2007, Jeju Island, Korea, August 27-29, 2007, Revised Selected Papers*, volume 4867 of *Lecture Notes in Computer Science*, pages 188–202. Springer, 2007. (Cited on pages 2 and 7).
- [29] Serge Vaudenay and Martin Vuagnoux. Passive-only key recovery attacks on RC4. In Carlisle M. Adams, Ali Miri, and Michael J. Wiener, editors, *Selected Areas in Cryptography, 14th International Workshop, SAC 2007, Ottawa, Canada, August 16-17, 2007, Revised Selected Papers*, volume 4876 of *Lecture Notes in Computer Science*, pages 344–359. Springer, 2007. (Cited on page 7).