Key Exchange in a Broadcast Environment via a Public-Key-Cryptosystem

Conf. Proc., Loughborough, U.K. 1985

R.M.F. Goodman, and A.J. McAuley *

Summary

In this paper we consider the problem of securely exchanging a key in an electronic network environment. We use a public key cryptosystem (PKC) in order to solve the key distribution problem. In particular we consider the problem of securely exchanging a session key in a broadcast situation. That is, where the broadcaster wishes to send an identical message to a number of users. In the paper we present solutions that are network specific, as well as a general method that effectively trades security for redundancy. An example using the trapdoor knapsack PKC is also given..

1 Introduction

In this paper we consider the problem of how to securely set up a broadcast transmission in an electronic network environment. Such networks include broadcast satellite and packet switched digital data services. We assume that any user may take the role of broadcaster , and that the broadcaster wishes to send an identical message to any subset of other users. In the case of a packet switched network it has been shown that considerable savings in transmission time (that is packet hops) can be achieved by using multi-address packets to broadcast to a set of users, as opposed to using individually addressed packets (Ref 1.). We therefore assume that the message is to be encrypted under a session key, and that any network user in possession of the session key will be able to decrypt the message. The message can then be sent in one broadcast to all the intended recipients, via multi-addressing. The problem then comes down to one of distributing the session key securely. Finally, we assume that no collusion exists between legitimate and illigitimate recipients. That is, the legal recipients of the broadcast have a vested interest in keeping the session key secret amongst themselves, and will not distribute it to unauthorised users.

2 Key Exchange

Diffie and Hellman (Ref 2.) have described a secure key exchange method based on the discrete logarithm problem. The method involves setting up a key by means of a two-way exchange between the two users. As such the method do not fit into the context of a passive broadcast. The two way exchange can be eliminated by using a true public key cryptosystem. In this case each user has a public encryption key which is stored in a public directory, and a private decryption key which is kept secret. The session key could then be encrypted under each recipient's key and sent via a separate packet. Alternatively, a multi-addressed packet could be sent with each user only able to decrypt their part of the ciphertext block. Either way, there is a considerable time overhead involved because we are essentially sending a separate cryptogram to each user. This is particularly true if the number of recipients (n) is large, and if the actual message is short. The session key setup could in the latter case be many hundreds of times the message length, and this of course translates directly into cost in a packet network.

An alternative solution is to trade complexity in terms of the number of public keys in the system rather than in distribution time. In this case the directory contains a pair of keys for each possible combination of broadcasters and receivers. This technique clearly suffers from an exponential explosion in number of keys, and would be impractical for more than ten destinations. We now consider ways in which the growth in numbers of keys, and in distribution time may be limited.

3 Exploiting Network Structure

If the network has the form of a ring, such as in a local area network, we can consider the following distribution method. The broadcaster enciphers the session key and a list of station addresses under the public key of the first receiver in say a clockwise direction. This receiver decrypts the information, and then re-encrypts under the public key of the next intended receiver.

* Department of Electronic Engineering
University of Hull
Hull UK

The information is thus passed from user to user round the ring with only the intended subset of users able to decrypt the session key and addresses. Additionally, each intended receiver can 'sign' the re-encrypted information using the signature property of a PKC before passing the packet on. The last legitimate user in the ring passes the information on and back to the broadcaster, who can then check the signatures to ensure that no user has been missed out of the chain, perhaps intentionally, or that no errors have occurred.

An alternative method of exploiting structure exists if the network or users have a fixed group structure. If the users form a layered hierarchical structure then security can be arranged in layers with a pair of keys for each level in the hierarchy. Thus the level 1 key would only allow level 1 users to decrypt the session key. The level 2 key would allow both level 2 and level 1 users to decrypt, and so forth. Any particular level therefore possesses the keys to its own level and to those below it.

Savings are also possible if the group of users are in the form of a tree structure. In this case the users are at the leaves of the tree, and the nodes define a 'pathway' of keys with which to access any subset of users. For example, consider a binary tree with 16 users. The root node is a key which all users possess, and therefore all leaves can be accessed with this key. There are two nodes that stem from the root, and these enable the two halves of the tree to be accessed separately. There are 4 nodes at the next level, and 8 nodes at the level before the leaves. Each user possesses a number of keys equal to the depth of the tree. Thus one key can cover 1,2,4,8, and 16 users in one key transmission provided the users are in the binary group covered by the key. If they are not then multiple key transmissions are required. There is therefore a structure-time type of trade-off with this organisation.

4 A Broadcast Cryptosystem

Each users public key essentially defines a one to one mapping between a particular message block and a particular cryptogram. The secret decryption key defines the inverse mapping. Ideally we may define a broadcast cryptosystem as one in which a particular cryptogram, when decrypted by each intended user's secret key, gives the same message. The advantage of this is that the single cryptogram needed to encipher the message (session key) would fit into the call set up packet and therefore require no overhead. There are, however, a number of drawbacks. Firstly, no such session key may exist between any two given users, or the choice of keys may be too low for security. Secondly, the choice of session key may be decipherable by others. And thirdly, it may be very difficult to find such keys in a practical system.

A possible solution to the first problem is to make use of redundancy. That is, the number of bits in the session key is less than the number of bits in the message. A given cryptogram does not now have to decrypt to exactly the same message for each intended user, but rather the subset of bits that forms the session key does. There will now be many more possible session keys. In fact we can show that there is a trade off between the probable number of session keys, and the redundancy used . This is given by:

$$R = b/n(k-1) + (k-2)/(k-1)$$

where R is the redundancy, that is, the number of bits used for the session key divided by the total number of message bits n, and k is the number of intended users. The parameter b is log to the base 2 of the number of probable keys. A number of points emerge from this relationship. Firstly, the minimum redundancy occurs when b=0, that is, the lowest security situation in which only one probable key exists, and depends only on the number of users. This minimum is 50% for three users and rises to 90% for 11 users. Secondly, for a fixed number of users and number of probable keys, the redundancy can be reduced by increasing the message length n. Finally, if the redundancy and message size is fixed, the security in terms of possible keys decreases with increasing users, as would be expected. The security can be assessed by comparing the difference in number of possible keys for k and k+1 users. This is approximately b/k , where b is the larger number of keys. A practical example is therefore n=2000 bits with a redundancy of 0.95 giving a session key length of 100 bits. The number of probable keys is 2 to the 2700, and the security level is 2 to the 270.

Now consider an small implementation based on the trapdoor knapsack PKC. We wish to distribute a 3 bit key to 2 users whose public keys consist of 8-bit knapsack vectors
a=(KN3,92,233,61,11,188,103,134)
b=(132,210,177,50,201,107,88,54)
If we chose the message vectors as:
ax=1,0,1,1,1,0,0,1,0 and bx=1,0,1,1,0,1,0,1
then both ciphertext subset sums compute to 520, and both users will therefore correctly decipher the session key as 101.

5 References

1.Quillon,J.M.,'Enhanced Message Addressing for Computer Networks',Proc.IEEE,Vol.66,1978.
2.Diffie,W.,and Hellman,M.E.,'New directions in Cryptography',Proc.IEEE,Vol.IT22,1976.