The Rise and Fall and Rise of Combinatorial Key Predistribution

Keith M. Martin

Information Security Group
Royal Holloway, University of London
Egham, Surrey TW20 0EX, U.K.
keith.martin@rhul.ac.uk

Abstract. There are many applications of symmetric cryptography where the only realistic option is to predistribute key material in advance of deployment, rather than provide online key distribution. The problem of how most effectively to predistribute keys is inherently combinatorial. We revisit some early combinatorial key predistribution shemes and discuss their limitations. We then explain why this problem is back "in fashion" after a period of limited attention by the research community. We consider the appropriateness of combinatorial techniques for key distribution and identify potential areas of further research.

Keywords: Key predistribution, combinatorial designs, sensor networks.

1 Introduction

Key management is a vital, and often overlooked, component of any cryptosystem. One of the most challenging phases of the cryptographic key life cycle is key establishment. This is particularly so in symmetric cryptosystems, where keys need to be established using some form of secure channel prior to use. In the following discussion we will only consider fully symmetric cryptosystems.

One of the main options for conducting symmetric key establishment is for one entity, which could be a trusted third party, to generate a key and then distribute it to those entities that require it. This process is often referred to as *key distribution*. Many application environments in which symmetric cryptography is deployed cannot rely on a key distribution service being regularly available whenever keys are required. Indeed in many applications such a service is impossible to provide after the network entities (which we will refer to as *nodes*) have been deployed. In such cases the only realistic option is for a trusted third party (which we will subsequently refer to as a *key centre*) to *predistribute* keys prior to deployment as part of a secure initialisation process. After deployment of the network, the key centre plays no further role in key establishment. Two nodes who require a common key must now try to derive one from the keys that they were each equipped with by the key centre prior to deployment. For this approach to be effective, the precise allocation of keys to nodes during initialisation is critical. This allocation is often termed a *key predistribution scheme*.

A. Biryukov, G. Gong, and D.R. Stinson (Eds.): SAC 2010, LNCS 6544, pp. 92–98, 2011.

A potential advantage of key predistribution is that the establishment of keys must happen at the key centre, which should be a controlled environment. However a significant disadvantage is that later stages of the key life cycle become more challenging to manage. Nonetheless, key predistribution is a popular approach to key establishment in real applications, especially those whose network topology is essentially "star-shaped", in the sense that communication only takes place between node and a central authority (*hub*) of some sort. In such cases it generally suffices to predistribute a unique key to each node, which its shares with the hub.

2 The Rise and Fall of Combinatorial Key Predistribution

A more interesting question is how to design a predistribution scheme for more general network topologies. Early research on key predistribution schemes focussed on the case where the network topology is the complete graph. The goal of such a key predistribution scheme is thus to enable any pair of nodes to share a predistributed key. One trivial solution is to predistribute a single key to all nodes, which results in minimal storage requirements for each node but has severe consequences if any node is compromised. At the other extreme, predistributing a unique key to every pair of nodes results in optimal resilience against node compromise but results in excessive storage requirements ($n-1$ keys for each node in a network of size n).

An interesting compromise between these two trivial solutions is the idea of a *w-key distribution pattern* (*KDP*) [12]. A w-KDP is an allocation of keys to nodes with the property that:

1. any pair of nodes (N_1, N_2) have some keys in common;
2. any w nodes other than N_1 and N_2 do not collectively have all the keys that are shared by N_1 and N_2.

Thus if a w-KDP is used as the basis for a key predistribution scheme, any pair of nodes have at least one predistributed key that is not known by an adversary who has compromised up to w other nodes in the network. Some (or all) of these keys than then be used to derive a key that N_1 and N_2 can use to secure their communication. We describe the resulting key predistribution schemes as *combinatorial* because most of the known techniques for constructing w-KDPs rely on combinatorial mathematics. Various generalisations of the idea of a w-KDP are possible and have been studied.

An alternative approach to designing a key predistribution scheme for networks based on the complete graph is to use symmetric polynomials. In *Blom's* key predistribution scheme [2] a polynomial $P(x,y) \in \mathrm{GF}(q)[x,y]$ with the property that $P(i,j) = P(j,i)$ for all $i,j \in \mathrm{GF}(q)$. The elegantly simple idea is that:

- Node N_i stores the univariate polynomial $f_i(y) = P(N_i, y)$;
- In order to establish a common key with N_j, node N_i computes $K_{ij} = f_i(N_j) = f_j(N_i)$.

Similarly to a w-KDP, this scheme is secure against an adversary who can compromise at most w nodes. A significant advantage is that each node is only required to store $w + 1$ polynomial coefficients, which is less information than most w-KDPs. The main related cost is that each node is not actually storing predistributed keys, but rather information that can be used to derive them. For many applications this tradeoff is likely to favour the Blom scheme.

The Blom key predistribution scheme easily generalises to key predistribution applications where groups of t nodes require common keys [3]. It was further shown that this approach is optimal with respect to node key storage. These observations lie behind my assertion that research combinatorial key predistribution underwent a rise and fall. The "rise" was the discovery of some very elegant key predistribution schemes based on combinatorial mathematics. The "fall" was a period of inactivity in this area, perhaps due to an impression that that the interesting questions had all been answered.

3 Evolving Network Security

There has been a significant increase in interest in key predistribution schemes in recent years. The main motivation is evolution of networking technology, with a trend towards distributed, dynamic, wireless networks consisting of lightweight devices of limited capability. These can manifest themselves in various different guises, including examples of mobile ad-hoc networks, tactical networks, ambient networks, vehicular networks and sensor networks. What is of most interest for a key management perspective is the following two common properties of such networks:

1. a lack of centralised post-deployment infrastructure;
2. the reliance on *hop-based* communication between nodes, where nodes are expected to act both as end points and routers of communication.

The first of these properties favours the use of key predistribution for key establishment. The second of these implies that it is not necessary for *every* pair of nodes to share a predistributed key, thus motivating the study of key predistribution schemes for more "relaxed" network topologies than the complete graph. Indeed, in many cases it suffices that nodes share keys with a small number of immediate neighbour nodes.

The lightweight nature of nodes in such networks has additional implications for key predistribution scheme design:

- limited memory may constrain the number of key that a node can store;
- limited power may constrain the computations and communications that a node can perform;
- fragility of nodes increase the risk of node compromise.

Thus the requirements for a particular application will almost always necessitate a tradeoff between contradictory requirements. For example it may be desirable to predistribute a large number of keys to each node from a connectivity perspective, since it increases the chances of two nodes sharing a key. However, it may also be desirable to limit the number of keys that each node stores due to memory constraints and a desire to reduce the impact of node compromise.

4 A Key Establishment Framework

In order to capture the different requirements placed on a key predistribution scheme by a potential application, a basic framework was proposed in [9]. The significant factors that influence the design of a key predistribution scheme are:

- *Homogeneity of nodes.* This determines whether all the nodes in the network have the same capabilities. The most common assumptions are that a network is either *homogeneous* (all nodes have the same capabilities) or *hierarchical* (there exists a hierarchy of capabilities, with nodes at higher levels having increased capabilities).
- *Deployment location control.* This categorises the extent to which the location of a node within the network is known prior to deployment, at the time that the key centre initalises it with predistributed keys. Clearly, location information is likely to help in the design of a suitable key distribution scheme. One extreme is *full location control*, where the precise location is known prior to deployment. In particular this means that the network neighbours of a node are predetermined. At the other extreme is *no location control*, where there is no information about the node location prior to deployment. Interestingly, there is potential for *partial location control*, where some location information may be known, for example that a certain group of nodes will be deployed in close proximity. Also of relevance is whether nodes are *static* or *mobile*.
- *Communication structure.* This determines what the desired communication structure of the network is. For example, are all nodes expected to directly communicate with one another, if possible, or are they only expected to communicate with near neighbours? Are group keys required as well as pairwise keys?

A particular key distribution scheme designed for a specific set of requirements within this framework can then be assessed in terms of the relevant metrics, for example storage requirements, energy requirements, efficiency of secure path establishment, etc.

5 The Second Rise of Combinatorial Key Predistribution

There has been a resurgence of interest in key predistribution schemes since Eschenauer and Gligor proposed the *random key predistribution scheme* [5], in which the key centre allocates keys to a node uniformly without replacement from a finite pool of keys. Schemes with different properties can be designed based on the size of the key pool and the number of keys allocated to each node, but they are all probabilistic, since it can no longer be guaranteed that a specific pair of nodes share a key.

This idea has been the basis for a large number of key predistribution scheme proposals, not all of which have been well motivated or analysed. While a substantial number of these proposals have been extensions of the random key predistribution scheme, one interesting avenue of research has focussed on the design of deterministic key predistribution schemes. These have certain potential advantages:

- by being deterministic, certain properties are guaranteed;
- analysis of deterministic key predistribution schemes is often simpler;
- an amount of established research has already been conducted on deterministic key predistribution schemes (the "first rise");
- some deterministic key predistribution schemes have useful algebraic structure (for example, they allow efficient shared key discovery).

A natural place to look for ideas for constructing deterministic key predistribution schemes is combinatorial mathematics. The focus of the earliest research in this "second rise" was to look at classical combinatorial structures, such as *projective planes*, most of which still provided full connectivity, in the sense that they were designed to facilitate a shared key between any pair of nodes. While some of these schemes offer interesting tradeoffs between the important parameters, they tend to be too restrictive and have high storage and resilience costs. More flexibility can be obtained by basing key predistribution schemes on combinatorial structures that are not fully connected. Indeed, several entirely new combinatorial structures of this type, for example *common intersection designs* [7], have been proposed and investigated specifically for adoption as key predistribution schemes for evolving networks.

However, given that design requirements of a key predistribution scheme often involve tradeoffs between competing parameters, a more natural role for combinatorial structures is to provide components from which more complex key predistribution scheme scan be built (for example [6]). A range of techniques for building key predistribution schemes in this way has been explored. Some of these build deterministic key predistribution schemes from deterministic components, while others use both deterministic and probabilistic components (for example. There would seem potential for further development of these *combinatorial engineering* approaches.

6 Research Directions

There have been a large number of recent proposals for key predistribution schemes, mostly explicitly targeted at wireless sensor network applications ([4] provides a good survey from 2005, but much has happened since). A substantial number of these consider the case of homogeneous, static nodes which are deployed with no location control. Many proposals seem rather ad hoc and are only compared against a limited number of previous proposals, largely using simulations to support claims about their worth. It is far from clear that such ad hoc proposals necessarily add much to the knowledge base concerning the design of key predistribution schemes. That said, there has also been some very interesting research conducted in this area and there is plenty more to do. We suggest the following guidance on future research direction:

1. *Deeper exploration of construction techniques.* It is relatively easy to propose a new construction technique for a key predistribution scheme. What is sometime harder, but is equally important, is to explore why the resulting properties arise. Simply showing that something works can suffice in some engineering disciplines, but we should be aiming higher in the study of key predistribution schemes.
2. *Better understanding of tradeoffs.* The tradeoffs between the desirable properties of a key predistribution scheme mean that, in theory, there are many different notions of "desirable tradeoff" amongst the potential properties of a key predistribution scheme. While this does suggest that there is a need for different design approaches, it is important to also develop a better general understanding of how different properties trade off against one another. In particular, the tradeoff between notions of connectivity and resilience seems deep and intriguing.
3. *Meaningful and well-motivated scenarios.* The diversity of potential evolving network application scenarios have the potential to motivate a number of quite distinct types of key predistribution scheme. In particular, consideration of degrees of location control present interesting variations of the more established problem (existing work on this includes our own treatment of linear networks [10], grids [1] and group-based deployment [11]). What is important is that particular application models are well-motivated and assessed in a meaningful way.
4. *Greater consideration of the key lifecycle.* Key establishment is just one phase of the wider key lifecycle. The use of key predistribution is mainly due to an assumption that the key centre is not generally available to maintain keys after deployment. However, this does not prevent the nodes in the network from jointly conducting some key management activities, such as network optimisation, key refreshment and key change, perhaps with occasional external assistance. Further research on post-deployment key management issues is merited.

The supporting theory behind combinatorial techniques has already played a significant role in helping to form the basis for a deeper understanding of how to build desirable key predistribution schemes for a wide range of different types of application (a survey of the role of combinatorics in key predistribution scheme design can be found in [8]). While not all of the above research will rely entirely on combinatorial approaches to key predistribution, there is no doubt that such approaches have a significant role to play.

References

1. Blackburn, S.R., Etzion, T., Martin, K.M., Paterson, M.B.: Distinct-Difference Configurations: Multihop Paths and Key Predistribution in Sensor Networks. IEEE Transactions in Information Theory 56(8), 3961–3972 (2010)
2. Blom, R.: An optimal class of symmetric key generation systems. In: Beth, T., Cot, N., Ingemarsson, I. (eds.) EUROCRYPT 1984. LNCS, vol. 209, pp. 335–338. Springer, Heidelberg (1985)
3. Blundo, C., De Santis, A., Herzberg, A., Kutten, S., Vaccaro, U., Yung, M.: Perfectly-secure key distribution for dynamic conferences. In: Brickell, E.F. (ed.) CRYPTO 1992. LNCS, vol. 740, pp. 471–486. Springer, Heidelberg (1993)
4. Çamtepe, S.A., Yener, B.: Key distribution mechanisms for wireless sensor networks: a survey. Rensselaer Polytechnic Institute, Computer Science Department, Technical Report TR-05-07 (March 2005)
5. Eschenauer, L., Gligor, V.: A key management scheme for distributed sensor networks. In: Proceedings of 9th ACM Conference on Computer and Communication Security (November 2002)
6. Lee, J., Stinson, D.R.: Deterministic key predistribution schemes for distributed sensor networks. In: Handschuh, H., Hasan, M.A. (eds.) SAC 2004. LNCS, vol. 3357, pp. 294–307. Springer, Heidelberg (2004)
7. Lee, J., Stinson, D.R.: Common intersection designs. Journal of Combinatorial Designs 14(4), 251–269 (2009)
8. Martin, K.M.: On the applicability of combinatorial designs to key predistribution for wireless sensor networks. In: Chee, Y.M., Li, C., Ling, S., Wang, H., Xing, C. (eds.) IWCC 2009. LNCS, vol. 5557, pp. 124–145. Springer, Heidelberg (2009)
9. Martin, K.M., Paterson, M.B.: An application-oriented framework for wireless sensor network key establishment. Electron. Notes Theor. Comput. Sci. 192(2), 31–41 (2008)
10. Martin, K.M., Paterson, M.B.: Ultra-lightweight key predistribution in wireless sensor networks for monitoring linear infrastructure. In: Markowitch, O., Bilas, A., Hoepman, J.-H., Mitchell, C.J., Quisquater, J.-J. (eds.) Information Security Theory and Practice. LNCS, vol. 5746, pp. 143–152. Springer, Heidelberg (2009)
11. Martin, K.M., Paterson, M.B., Stinson, D.R.: Key predistribution for homogeneous wireless sensor networks with group deployment of nodes. ACM Transactions in Sensor Networks, 7(2), article No. 11 (2010)
12. Mitchell, C.J., Piper, F.C.: Key storage in secure networks. Discrete Applied Mathematics 21, 215–228 (1988)