J. Cryptology (1999) 12: 227–239

Journal of
CRYPTOLOGY

Bounds for Key Distribution Patterns

Kathleen A. S. Quinn
Department of Pure Mathematics, The Open University,
Walton Hall, Milton Keynes MK7 6AA, England

Communicated by Douglas R. Stinson

Received 29 January 1996 and revised 12 September 1997

Abstract. This paper is concerned with the problem of distributing pieces of information to nodes in a network in such a way that any pair of nodes can compute a secure common key but the amount of information stored at each node is small. It has been proposed that a special type of finite incidence structure, called a *key distribution pattern* (*KDP*), might provide a good solution to this problem. We give various lower bounds on the information storage of KDPs. Our main result shows that in general KDP schemes necessarily have greater information storage at the nodes than the minimum possible. This minimum is achieved by a scheme not based on KDPs.

Key words. Key distribution patterns, Secure networks.

1. Introduction

Consider a network of v nodes, each of which must be able to communicate with each other node, using a symmetric cryptosystem. Then each pair of nodes requires a cryptographic key available to them but to no other node. To provide sufficient security, each of these keys has to contain a certain amount of information. Throughout this paper we use n to denote the number of bits of information required in each key. (It is convenient from Section 4 onwards to measure information in bits, and so we do so throughout.)

A *key distribution scheme* (KDS) is a method of distributing secret pieces of information to nodes in the network in such a way that any pair of nodes can compute a secure common key. The nodes compute the keys without further secure communication with the server which initially distributes the secret information.

The obvious (*trivial*) KDS would be for every node in the network to be provided with a separate key for use with each other node. This would require each node to store $v - 1$ keys, each of n bits, and the server to generate, and probably store, $\binom{v}{2}$ keys, each of n bits. The disadvantage of this scheme is the large amount of information storage required.

Various KDSs have been proposed which significantly reduce the amount of information storage over that required by this trivial scheme, but this can be done only at a cost. We say that a KDS is w-*secure* if, given any pair of nodes, any set of w or fewer other nodes may pool their information and still have no better chance of correctly guessing the key of the pair than an outsider of the network. Clearly, the trivial KDS is $(v-2)$-secure, but all the proposed schemes are w-secure only for much smaller values of w.

Blom [1] has shown that with any w-secure KDS, each node must store at least $(w+1)n$ bits of information. This bound is tight: Blom gives a construction for a class of w-secure KDSs achieving it. The total amount of secret information generated and stored by the server in Blom's scheme is $\frac{1}{2}(w+1)(w+2)n$ bits.

In [7] Mitchell and Piper show how design theory provides a source of KDSs. They define a certain special type of finite incidence structure, which they call a *key distribution pattern* (KDP). Essentially the idea is as follows. The server generates a ground set of *subkeys*, each of which consists of independent secret information, and distributes a different subset of the ground set to each node. Information about which subkeys each node has is public knowledge, using reference numbers for the subkeys. The key to be used by a pair of nodes is made up by combining those subkeys which the pair of nodes have in common. The combining should be done using a publicly known function f, which takes a number of subkeys as argument and yields a key containing n bits of information. The sets of subkeys distributed to the nodes have to be specially chosen to ensure that the system is w-secure for a specified value of w, and this is what the structure of a KDP achieves. We give a formal definition of a KDP in Section 2. Several classes of KDPs are described in [4] and [6]–[9]. In [3] KDPs are constructed using probabilistic methods.

For our purposes in this paper the nature of the function f does not matter. If the subkeys and keys are bit-strings, we can use concatenation, but the resulting long strings of bits may be unsuitable in some applications. If the subkeys and keys are all bit-strings of the same length n, we can simply add the subkeys in $\mathrm{GF}(2)^n$ to obtain the keys. A more general, and flexible, approach is the use of so-called *resilient functions*; this is discussed in [10].

In this paper we are concerned with the question of the extent to which the key storage requirement of a network can be reduced by the use of a KDP.

Section 2 covers the basic definitions relating to KDPs. In Section 3 we briefly review some bounds on the number of subkeys at each node and on the total number of subkeys which other authors have published, and also prove some new bounds.

In order to consider properly the question of the extent to which the key storage requirement of a network can be reduced by the use of a KDP, we should seek bounds not only on the number of subkeys, but also on the total information which the subkeys contain. The definitions relevant to this are introduced in Section 4.

In Section 5 we prove a lower bound on the information storage at each node which can be achieved using a KDP. This bound shows that KDPs will not in general yield KDSs with node storage as good as that achieved by Blom's KDSs. The difference between the bound and the node storage achieved by Blom's KDSs is negligible if w (the number of colluders protected against) is small compared with n (the number of bits of information required in each key), but as w increases, the difference becomes significant.

2. Key Distribution Patterns

KDPs may be thought of either as set systems, as in [3], or as incidence structures, as in [7], the original paper on the subject. The main reason for thinking of them as incidence structures, which is slightly more complicated, is that design theory provides many interesting examples of KDPs [6]–[9]. We use the incidence structure representation, for consistency with earlier work.

All references to designs in this paper are brief and non-central. Definitions of any design theory terms and notation which we do not define can be found in [5].

An *incidence structure* (or simply *structure*) is a triple $\mathcal{S} = (\mathcal{P}, \mathcal{B}, I)$, where $\mathcal{P}$ and $\mathcal{B}$ are non-empty sets of objects, called *points* and *blocks*, respectively, and $I \subseteq \mathcal{P} \times \mathcal{B}$. If $\mathcal{P}$ and $\mathcal{B}$ are finite, then we say that $\mathcal{S}$ is a *finite* incidence structure. We call $\mathcal{P}$ the *point set* of $\mathcal{S}$ and $\mathcal{B}$ the *block set* of $\mathcal{S}$. We usually denote $|\mathcal{P}|$ by v and $|\mathcal{B}|$ by b. If point P and block x are such that $(P, x) \in I$, then we say that P and x are *incident* with each other. The set of blocks incident with a point P is usually denoted by (P), and analogously the set of points incident with a block x is denoted by (x).

A common way to specify an incidence structure is to list the b sets in the multiset $\{(x) \mid x \in \mathcal{B}\}$. For example, the following list specifies an incidence structure with eight points $1, 2, \dots, 8$ and sixteen blocks:

$$\begin{array}{cccc}
\{1,3,5,7\}, & \{1,4,5,8\}, & \{1,3,6,8\}, & \{1,4,6,7\}, \\
\{2,4,6,8\}, & \{2,3,6,7\}, & \{2,4,5,7\}, & \{2,3,5,8\}, \\
\{1,2\}, & \{3,4\}, & \{5,6\}, & \{7,8\}, \\
\{1,2\}, & \{3,4\}, & \{5,6\}, & \{7,8\}.
\end{array}$$

An incidence structure could equally well be specified by listing the v sets in the multiset $\{(P) \mid P \in \mathcal{P}\}$. Although this is non-standard, it is sometimes a convenient way to think of incidence structures when considering KDPs. For example, if we label the blocks of the above structure by $1, 2, \dots, 16$, respectively, then this same structure is specified by the following list of subsets of the set $\mathcal{B} = \{1, 2, \dots, 16\}$:

$$\begin{array}{cc}
\{1,2,3,4,9,13\}, & \{5,6,7,8,9,13\}, \\
\{1,3,6,8,10,14\}, & \{2,4,5,7,10,14\}, \\
\{1,2,7,8,11,15\}, & \{3,4,5,6,11,15\}, \\
\{1,4,6,7,12,16\}, & \{2,3,5,8,12,16\}.
\end{array}$$

A *w-secure key distribution pattern* (*w-KDP*) is a finite incidence structure $\mathcal{K} = (\mathcal{P}, \mathcal{B}, I)$ with at least $w + 2$ points with the property that

$$(P_1) \cap (P_2) \not\subseteq (Q_1) \cup (Q_2) \cup \cdots \cup (Q_w)$$

for all subsets $\{P_1, P_2, Q_1, Q_2, \dots, Q_w\}$ of $w + 2$ points of $\mathcal{P}$, where $w \geq 1$. To use a KDP $\mathcal{K}$ to give a KDS, we identify each node of the network with a point of $\mathcal{K}$ and a subkey with each block of $\mathcal{K}$. Each node P is issued with the subkeys in (P). When discussing KDPs, we use the words *point* and *node* and the words *block* and *subkey* interchangeably. The special property which we require of a w-KDP ensures that the key of a pair of nodes cannot be compromised by any colluding set of w or fewer other nodes, since no set of w other nodes will between them hold all of the subkeys which the pair have in common. Thus, if every subkey contains n bits of information, the same as in the keys, then a w-secure KDP gives a w-secure KDS. Of course, a w-KDP is a w'-KDP for all w' with $1 \leq w' \leq w$. Since, in particular, every w-KDP is a 1-KDP, a 1-KDP will usually just be called a KDP.

For example, the particular structure with eight points specified earlier in this section is a 1-KDP. To see this, it is perhaps easiest to look at the second representation of the structure, the list of the sets in $\{(P) \mid P \in \mathcal{P}\}$: these correspond to the eight sets of subkeys to be issued to the eight nodes, and it can be verified that no two of the sets listed have an intersection which is a subset of a third set in the list. This second representation is essentially the KDP as a set system.

Note that the reason for identifying points with nodes and blocks with subkeys, rather than the other way round—which might initially seem more natural—is that with this definition, many designs, and structures derived from designs, are KDPs. In design theory we specify that two points must have a number of blocks in common, which connects with our requirement that in a KDP two nodes must have some subkeys in common.

The trivial KDS for v nodes, in which each node is provided with a separate key for use with each other node, corresponds to a KDP with v points, $\binom{v}{2}$ blocks and $v - 1$ blocks on each point, which we call the *trivial KDP* on v points. (It is just the trivial 2-$(v, 2, 1)$ design.) For this KDP, subkeys are the same as keys. The trivial KDP with v points is $(v - 2)$-secure.

The trivial KDPs provide a yardstick by which we may judge other KDPs. We denote the number of subkeys of a KDP incident with a node P by r_P and the total number of subkeys of the KDP by b. A KDP which has r_P smaller than $v - 1$ for each node P gives a better node storage than the trivial KDP with v nodes. Similarly, a KDP which has b smaller than $\binom{v}{2}$ gives a better total storage at the server than the trivial KDP with v nodes.

3. Bounds for the Number of Subkeys in a KDP

In this section we give lower bounds for the values b and r_P of a KDP. These can be compared with the values $r_P = v - 1$ and $b = \binom{v}{2}$ for the trivial KDP with v nodes. If we are concerned only with the case in which subkeys consist of n bits (other possibilities are discussed in Section 4), the lower bounds for r_P can be multiplied by n to give a direct comparison with Blom's lower bound of $(w + 1)n$ bits for the node storage in a KDS. We use lg to denote $\log_2$.

The first few bounds in this section are derived from a result of Sperner, which we now state. A set $\mathcal{F}$ of subsets of a finite ground set G is called a *Sperner system* if none of the sets in $\mathcal{F}$ contains another. Examples of large Sperner systems are easily found: if $|G| = g$, then, for any s with $0 \leq s \leq g$, the set of all s-subsets (that is, subsets with s elements) of G is a Sperner system. Of these, the set of all $\lfloor g/2 \rfloor$-subsets of G contains the most sets, and the following theorem of Sperner (1928) states that no Sperner system with ground set G can contain more sets than this. A proof can be found in [2].

Result 3.1. *A Sperner system whose ground set contains g elements consists of at most $\binom{g}{\lfloor g/2 \rfloor}$ sets.*

The following result gives simple lower bounds on r_P and b. It essentially appears in [3] and in a slightly different form in Mitchell and Piper's original paper [7].

Result 3.2. *For a KDP with v nodes,*

$$r_P \geq \lg v \qquad \textit{for any node } P,$$

and

$$b \geq 2\lg(v-1).$$

Proof. Let P be a node of a KDP $\mathcal{K} = (\mathcal{P}, \mathcal{B}, I)$ with v nodes. Notice that (P) must have at least $v-1$ distinct non-empty subsets, namely, the sets in $\{(P) \cap (Q) \mid Q \in \mathcal{P}\backslash\{P\}\}$. Hence $v - 1 \leq 2^{r_P} - 1$, that is, $r_P \geq \lg v$.

For the bound on b, notice that the $\binom{v}{2}$ sets in $\{(P_1) \cap (P_2) \mid \{P_1, P_2\} \subseteq \mathcal{P}\}$ form a Sperner system with ground set $\mathcal{B}$. Therefore, by Result 3.1 (Sperner's theorem),

$$\binom{b}{\lfloor b/2 \rfloor} \geq \binom{v}{2}.$$

It is straightforward to prove by induction that

$$2^{b-1} \geq \binom{b}{\lfloor b/2 \rfloor}$$

for all positive integers b, so we may deduce that $2^{b-1} \geq \frac{1}{2}(v-1)^2$, that is, $b \geq 2\lg(v-1)$. □

Note that the inequality

$$\binom{b}{\lfloor b/2 \rfloor} \geq \binom{v}{2}$$

of the above proof is used in [3] to derive the bound $b \geq 2\lg v$, which is slightly better than the bound we give. However, it appears that a less straightforward derivation is required for a marginal improvement in the bound.

In [3], Dyer et al. use probabilistic methods to construct 1-KDPs with v nodes which have $b = \lceil 13 \lg v \rceil$, and which have a mean value for r_P of approximately $\frac{26}{3} \lg v$. Non-probabilistic constructions have not yielded KDPs with values of b and r_P of this order. Amongst the best of the deterministically constructed families of 1-KDPs are the following. It is easy to show that the biplanes (symmetric 2-designs) are KDPs which achieve $b = v$ and $r_P \approx \sqrt{2v}$ for each node P, but note that the largest known biplane has $v = 79$. An infinite family of 1-KDPs constructed from finite projective planes in [9] has $b = 2v$ and $r_P \approx 2\sqrt{v}$ for each node P.

Result 3.2 above gives bounds for r_P and b for KDPs in general. We would expect to find stronger bounds for KDPs which are w-secure for $w > 1$. In [3] Dyer et al. use Sperner's theorem to derive a bound on b for w-KDPs, as follows.

Result 3.3. [3] *For a w-KDP with v nodes,*

$$b \geq w(2 \lg v - \lg w - 1).$$

We see from this that we must have b growing at least linearly with the number of potential colluders.

Dyer et al., in [3], use probabilistic methods to construct w-KDPs on v nodes which have $b < \lceil 2(w+2)^3 \ln v \rceil$. The multimap scheme of [4] gives an infinite family of KDPs with v nodes and $b = (w+1)v$.

In [3], which is primarily concerned with finding KDPs with a small value of b, no bound better than the simple bound of Result 3.2 is given for r_P in a w-KDP. Here we give two improved bounds.

First, we use the technique of Dyer et al.'s proof of the bound for b in Result 3.3 to derive an analogous bound for r_P.

Theorem 3.4. *For any node P of a w-KDP with v nodes,*

$$r_P \geq w(\lg(v-1) - \lg w).$$

Proof. Let P be a node of a w-KDP with v nodes. Consider the $v-1$ sets in

$$\{(P) \cap (Q) \mid Q \in \mathcal{P} \backslash \{P\}\}.$$

We claim that the $\binom{v-1}{w}$ possible unions of w of the elements of this set form a Sperner system (with ground set (P)). For suppose not. Then for some nodes we would have

$$[(P) \cap (Q_1)] \cup \cdots \cup [(P) \cap (Q_w)] \subseteq [(P) \cap (R_1)] \cup \cdots \cup [(P) \cap (R_w)],$$

where there is a node Q_i on the left which is not one of the nodes $R_1, \ldots, R_w$ on the right. For this node we obtain

$$(P) \cap (Q_i) \subseteq (R_1) \cup \cdots \cup (R_w),$$

which contradicts the assumption that $\mathcal{K}$ is a w-KDP.

Applying Sperner's theorem gives

$$\binom{r_P}{\lfloor r_P/2 \rfloor} \geq \binom{v-1}{w}.$$

Using the fact that

$$2^{r_P - 1} \geq \binom{r_P}{\lfloor r_P/2 \rfloor}$$

we obtain

$$2^{r_P - 1} \geq \prod_{i=0}^{w-1} \frac{v-1-i}{w-i}.$$

It is easy to see that the factors on the right increase with i so we may deduce that $2^{r_P-1} \geq ((v-1)/w)^w$, from which it follows that $r_P - 1 \geq w(\lg(v-1) - \lg w)$. The result follows. □

We see from this that r_P too must grow at least linearly with the number w of potential colluders, and also with the logarithm of the number of nodes in the network.

Dyer et al. in [3], use probabilistic methods to construct w-KDPs on v nodes which have $r_P < \lceil 4(w+2)^2 \ln v \rceil$. In [9] finite affine planes are used to construct w-KDPs on v nodes with $r_P \approx (w+2)\sqrt{v}$, for $w \le \sqrt{v} - 1$.

In the next theorem (Theorem 3.6) we give a second bound for r_P which shows that r_P must grow at least with the square of the number w of potential colluders, until r_P reaches the value for r_P found in the trivial KDP. We begin by establishing a lemma.

Lemma 3.5. *Let P_1 and P_2 be two nodes of a w-KDP $\mathcal{K} = (\mathcal{P}, \mathcal{B}, I)$ such that every subkey in $(P_1) \cap (P_2)$ is held by at least one other node. Then, for any subset S of $\mathcal{P} \backslash \{P_1, P_2\}$ with $0 \le |S| \le w$,*

$$\left| (P_1) \cap (P_2) \backslash \bigcup_{Q \in S} (Q) \right| \ge w + 1 - |S|.$$

Proof. Suppose not. Let S be a maximal subset of $\mathcal{P} \backslash \{P_1, P_2\}$ with $0 \le |S| \le w$ such that

$$\left| (P_1) \cap (P_2) \backslash \bigcup_{Q \in S} (Q) \right| \le w - |S|.$$

Since $\mathcal{K}$ is a w-KDP we know that

$$\left| (P_1) \cap (P_2) \backslash \bigcup_{Q \in S} (Q) \right| \ge 1.$$

It follows from these two inequalities that $|S| \le w - 1$. Also, from the second of these two inequalities and the fact that every subkey in $(P_1) \cap (P_2)$ is held by at least one other node, we can deduce that some node $Q' \in \mathcal{P} \backslash (S \cup \{P_1, P_2\})$ must hold a subkey in $(P_1) \cap (P_2)$. Therefore

$$\left| (P_1) \cap (P_2) \backslash \bigcup_{Q \in S \cup \{Q'\}} (Q) \right| \le w - |S| - 1.$$

We see from this that the set $S \cup \{Q'\}$ has the necessary properties to contradict the maximality of S. This completes the proof. □

Before we prove the next theorem, we point out that it is an immediate corollary of Lemma 3.5 that every pair of nodes of a w-KDP must either have a common subkey which is held by no other node (as they would in the trivial KDP), or must combine at least $w + 1$ subkeys to form their key.

Theorem 3.6. *For any node P of a w-KDP with v nodes,*

$$r_P \ge \min\{v - 1, \tfrac{1}{2}(w+1)(w+2)\}.$$

Proof. Let P be a node of a w-KDP $\mathcal{K} = (\mathcal{P}, \mathcal{B}, I)$ with v nodes. Suppose that $r_P < v - 1$.

Let S be the set of all nodes in $\mathcal{P}\backslash\{P\}$ which are such that every subkey held by the node and by P is also held by a third node. Let S' consist of all other nodes in $\mathcal{P}\backslash\{P\}$, that is, those which hold a subkey held by P which is held by no third node. Then $|S'| = v - 1 - |S|$.

We now show that $|S| \geq w + 2$. Since P has a different subkey in common with every node in S', $r_P \geq v - 1 - |S|$. Hence, since $r_P < v - 1$, S is non-empty. Let $Q \in S$. By Lemma 3.5, $|(P) \cap (Q)| \geq w + 1$. Also P has at least one distinct subkey not held by Q in common with every node in S'. Hence $r_P \geq (w + 1) + (v - 1 - |S|) = w + v - |S|$. However, $r_P < v - 1$, so it follows that $v - 1 > w + v - |S|$, that is, $|S| \geq w + 2$.

Let $\{Q_1, Q_2, \ldots, Q_{w+1}\} \subset S$. By Lemma 3.5, for $0 \leq j \leq w$, $(P) \cap (Q_{j+1})$ contains at least $w + 1 - j$ subkeys not in $\bigcup_{i=1}^{j}(Q_i)$. So

$$\begin{aligned} r_P &\geq (w + 1) + w + (w - 1) + \cdots + 1 \\ &= \tfrac{1}{2}(w + 1)(w + 2) \end{aligned}$$

as stated. □

Unless a w-KDP is known to have further properties of the type discussed in the next section, in an implementation of the KDP the subkeys would have to contain the same number n of bits of information as the keys. Theorem 3.6 tells us that in such an implementation of a w-KDP on v nodes, the information storage at any node P must either be at least $(v - 1)n$ bits, that is, at least what it would be if the trivial KDS were used, or else it must be at least $\frac{1}{2}(w + 1)(w + 2)n$ bits. This compares badly with the optimal value of $(w + 1)n$ bits achieved by Blom's KDSs.

4. KDP Systems in Which Subkeys Can Contain Less Information than Keys

Having fewer subkeys is one way in which a KDP can improve on the information storage of the trivial KDP, but there is another way. Some KDPs have properties which permit the subkeys to contain less information than the keys. This was first pointed out by Mitchell in [6]. We consider this next, but first we take note of an assumption which applies throughout the remainder of this paper and upon which the results in Section 5 are strongly dependent.

In all the published work on KDPs, it is assumed that the subkeys are strings of binary digits, randomly generated by the server. In this paper we are essentially still concerned with this same situation. However, for precision we make a more minimal assumption, namely, that the subkeys each contain at least one bit of information. We assume that this is true of the keys also. Recall that we also assume that the subkeys are independent of each other.

Let $\mathcal{K} = (\mathcal{P}, \mathcal{B}, I)$ be a KDP. For each subkey x of $\mathcal{K}$, let $l(x)$ denote the information content of x in bits, so that l is a mapping from the set $\mathcal{B}$ of subkeys to the real numbers greater than or equal to one. We call any mapping $l{:}\, \mathcal{B} \to [1, \infty)$ an *information map* for $\mathcal{K}$.

We write $(\mathcal{K}, l)$ for the system consisting of a KDP $\mathcal{K}$ together with an information map l for K. If X is a set of subkeys of a system $(\mathcal{K}, l)$, then we denote the total information content in bits of the subkeys in X by $\|X\|$, that is, $\|X\| = \sum_{x \in X} l(x)$.

Suppose that $\mathcal{K}$ is a w-KDP for some $w \geq 1$. Bearing in mind that we want a w-secure KDS, we consider the property which we would require an information map l of $\mathcal{K}$ to possess. Let $\{P_1, P_2, Q_1, Q_2, \ldots, Q_w\}$ be any set of $w + 2$ distinct nodes of $\mathcal{K}$. Since we must ensure that, even if $Q_1, Q_2, \ldots, Q_w$ pool their subkey sets, their chance of guessing the n-bit key of the pair P_1, P_2 is no greater than that of an outsider who knows none of the subkeys in $(P_1) \cap (P_2)$, we must ensure that at least n bits of information not from the subkeys in $\bigcup_{i=1}^{w} Q_i$ contribute to the key of the pair P_1, P_2. That is, we require

$$\left\| (P_1) \cap (P_2) \backslash \bigcup_{i=1}^{w} (Q_i) \right\| \geq n. \tag{1}$$

We require this to hold for all sets $\{P_1, P_2, Q_1, Q_2, \ldots, Q_w\}$ of $w + 2$ distinct nodes of $(\mathcal{K}, l)$. We thus make the following definition.

Let $\mathcal{K} = (\mathcal{P}, \mathcal{B}, I)$ be a KDP. We call $l : \mathcal{B} \to \mathbb{N}$ a *w-secure information map* for $\mathcal{K}$ if $(\mathcal{K}, l)$ satisfies (1) for every set $\{P_1, P_2, Q_1, Q_2, \ldots, Q_w\}$ of $w + 2$ nodes of $\mathcal{K}$. Notice that $\mathcal{K}$ must be a w-KDP for such a mapping l to exist, and that if l is w-secure for $\mathcal{K}$, then l is w'-secure for $\mathcal{K}$ for every w' with $1 \leq w' \leq w$. Clearly, if $\mathcal{K}$ is any w-KDP, the information map which sets the information content of every subkey to be n bits, the same as the information content of the keys, is a w-secure information map for K. We call this map L_n. However, many w-KDPs admit a better information map.

As an example of this consider the following. There is a particular KDP $\mathcal{K} = (P, \mathcal{B}, I)$ with 10 nodes and 30 subkeys which is w-secure for $w \leq 3$ and has the property that any three nodes have exactly one subkey in common and any two nodes have exactly four subkeys in common. (A 3-(10, 4, 1) inversive plane is such a KDP.) This means that, considering $\mathcal{K}$ as a 1-KDP, if we take any set $\{P_1, P_2, Q\}$ of three nodes, then $(P_1) \cap (P_2) \backslash (Q)$ will contain exactly three subkeys. Therefore subkeys need only contain $n/3$ bits of information. Similarly, considering it as a 2-KDP, if we take any set $\{P_1, P_2, Q_1, Q_2\}$ of four nodes, then $(P_1) \cap (P_2) \backslash (Q_1) \cup (Q_2)$ will contain at least (in fact, exactly) two subkeys. Therefore subkeys need only contain $n/2$ bits of information. If we consider it as a 3-KDP, then subkeys have to contain n bits of information.

This particular KDP is of course too small to be of any practical use but larger inversive planes are useful KDPs with similar properties. A fuller discussion of the inversive planes in this context can be found in [6].

Let $\mathcal{K}$ be a KDP and let l be an information map for $\mathcal{K}$. We write β for $\|\mathcal{B}\|$ and, for any node P of $\mathcal{K}$, we write ρ_P for $\|(P)\|$. So ρ_P represents the total number of bits of information contained in the subkeys in (P), and β represents the total number of bits of information generated, and perhaps stored, by the server.

A good KDP is one which allows ρ_P, for each node P, to be reduced to significantly less than the corresponding value for the trivial KDP on the same number v of nodes, that is, to significantly less than $(v - 1)n$, or which allows β to be reduced to significantly less than the corresponding value for the trivial KDP on the same number of nodes, that is, to significantly less than $\binom{v}{2}n$. Whether we are interested in a small value of ρ_P, or a small value of β, or both, will depend on the application. Usually we are interested in a small information storage at the nodes, so ρ_P is likely to be important.

As we have seen, there are two ways in which a w-KDP can achieve a small information storage at the nodes: by having a small value of r_P for each node P, or by admitting a "good" w-secure information map. The biplanes have only the first of these properties, and the inversive planes have only the second, but some KDPs have both properties: for example, a family of such KDPs is constructed in [9].

Similarly, a w-KDP can achieve a small total information storage by having a small value of b or by admitting a good w-secure information map.

Of course, allowing the subkeys to contain less information than the keys can reduce the information storage at the nodes and the total information storage of a KDP to at best $1/n$ times what these values would be otherwise, and this would involve a large number of subkeys each containing only a small amount of information (just one bit for maximum effect), which may in itself have disadvantages.

Finally, in this section, it is possibly worth pointing out that Theorem 3.6 gives us some insight into the maximum number of nodes whose collusion a KDP can protect against. Consider a network of v nodes. First suppose that we use a KDP on v nodes in which the number of subkeys at a node P is at least the same as the number of subkeys at P in the trivial KDP, that is, at least $v - 1$. Then the information storage at P must be at least $v - 1$ bits, that is, at least $1/n$ times the node storage of the trivial KDP. (We assume throughout that we use the information map L_n for the trivial KDP: clearly, it does not admit a better information map.) Suppose on the other hand that we use a KDP in which the number of subkeys at a node P is less than the number at P in the trivial KDP, that is, less than $v - 1$. Then, by Theorem 3.6, we must have $\frac{1}{2}(w + 1)(w + 2) < v - 1$, showing that the maximum number of colluders that the KDP can protect against is at most approximately $\sqrt{2v}$.

5. Bounds for the Information Storage in a KDP System in Which Subkeys Can Contain Less Information than Keys

Note first that each of the bounds for r_P and b in Section 3 immediately gives a corresponding bound for ρ_P or β. Under the assumptions at the beginning of Section 4, for any KDP with any 1-secure information map, we have that $\rho_P \geq r_P$ for any node P, and $\beta \geq b$. Therefore we may replace r_P by ρ_P and b by β in all the lower bounds of Section 3 to yield valid bounds for ρ_P and β.

However, in this section we adapt the proofs of Lemma 3.5 and Theorem 3.6, under these same assumptions, to yield a better lower bound for ρ_P. This lower bound gives the comparison with Blom's optimal schemes to which we referred in the introduction. We begin with the adaptation of Lemma 3.5.

Lemma 5.1. *Let $\mathcal{K}$ be a w-KDP and let l be a w-secure information map for $\mathcal{K}$. Let P_1 and P_2 be any two nodes of $(\mathcal{K}, l)$ such that the total amount of information contained in subkeys held by only P_1 and P_2 is less than n bits. Then, for any subset S of $\mathcal{P}\backslash\{P_1, P_2\}$*

such that $0 \le |S| \le w$,

$$\left\| (P_1) \cap (P_2) \setminus \bigcup_{Q \in S} (Q) \right\| \ge w + n - |S|.$$

Proof. Suppose not. Let S be a maximal subset of $\mathcal{P} \setminus \{P_1, P_2\}$ with $0 \le |S| \le w$ such that

$$\left\| (P_1) \cap (P_2) \setminus \bigcup_{Q \in S} (Q) \right\| < w + n - |S|.$$

Since l is w-secure for $\mathcal{K}$ we know that

$$\left\| (P_1) \cap (P_2) \setminus \bigcup_{Q \in S} (Q) \right\| \ge n.$$

It follows from these two inequalities that $|S| < w$. Also, from the second of these two inequalities and the fact that the total amount of information contained in subkeys held only by P_1 and P_2 is less than n bits, we can deduce that some node $Q' \in \mathcal{P} \setminus (S \cup \{P_1, P_2\})$ must hold a subkey in $(P_1) \cap (P_2)$. This subkey contains at least one bit of information and therefore

$$\left\| (P_1) \cap (P_2) \setminus \bigcup_{Q \in S \cup \{Q'\}} (Q) \right\| < w + n - |S| - 1.$$

We see from this that the set $S \cup \{Q'\}$ has the necessary properties to contradict the maximality of S. This completes the proof. □

The next theorem states the main result.

Theorem 5.2. *Let $\mathcal{K}$ be a w-KDP and let l be any w-secure information map for $\mathcal{K}$. Then, for any node P of $(\mathcal{K}, l)$,*

$$\rho_P \ge \min \left\{ v - 1, \ (w + 1)\left(n + \frac{w}{2} \right) \right\}.$$

Proof. Let P be a node of $\mathcal{K}$. Suppose that $\rho_P < v - 1$. Let S be the set consisting of all nodes in $\mathcal{P} \setminus \{P\}$ which are such that the total amount of information contained in subkeys which are held by the node and by P but by no third node is less than n bits. Let S' consist of all other nodes in $\mathcal{P} \setminus \{P\}$, that is, those which are such that the the total amount of information contained in subkeys which are held by them and by P but by no third node is at least n bits. Then $|S'| = v - 1 - |S|$. We now show that $|S| \ge w + 2$. Since P must have a different subkey in common with every node in S', $r_P \ge v - 1 - |S|$, and hence since subkeys must contain at least one bit of information, $\rho_P \ge v - 1 - |S|$. However, $\rho_P < v - 1$ so $v - 1 > v - 1 - |S|$, that is, $|S| \ge 1$. Let $Q \in S$. By Lemma 5.1, $\|(P) \cap (Q)\| \ge w + n \ge w + 1$. Also P has at least one distinct subkey not held by Q in common with every node in S', and each of these must contain at least one bit of information. Hence $\rho_P \geq (w+1) + (v-1-|S|) = w+v-|S|$. However, $\rho_P < v-1$ so it follows that $v-1 > w+v-|S|$, that is, $|S| \geq w+2$.

Let $\{Q_1, Q_2, \ldots, Q_{w+1}\} \subset S$. By Lemma 5.1, for $0 \leq j \leq w$, the subkeys in $(P) \cap (Q_{j+1})$ contain at least $w+n-j$ bits of information not contained in subkeys in $\bigcup_{i=1}^{j}(Q_i)$. So

$$\begin{aligned}\rho_P &\geq (w+n) + (w+n-1) + (w+n-2) + \cdots + n \\ &= (w+1)\left(n+\frac{w}{2}\right)\end{aligned}$$

as stated. □

The first bound in Theorem 5.2 above says that ρ_P is greater than or equal to $1/n$ times the node storage $(v-1)n$ of the trivial KDP: this is poor in comparison with the value $\rho_P = (w+1)n$ attained by Blom's KDSs. The second possibility $\rho_P \geq (w+1)(n+w/2)$ is also inferior to the value of ρ_P attained by Blom's KDSs. If the value of w (the number of colluders protected against) is relatively small compared with n (the number of bits of information required in each key), the difference between $(w+1)(n+w/2)$ and $(w+1)n$ is neglible, but as w increases, the difference becomes significant. Note also that in order to approach the bound $\rho_P \geq (w+1)(n+w/2)$ we would have to use a KDP system involving a large number of subkeys each containing only a small amount of information (approaching one bit), which may in itself have disadvantages.

There is an infinite family of KDPs, with node storage better than that of the corresponding trivial KDP, attaining the first bound $\rho_P \geq v-1$, as follows. Let q be any prime power. Taking an external structure of a 3-$(q^2+1, q+1, 1)$ inversive plane yields a 2-design. This design is a KDP on q^2 nodes, and for w satisfying $1 \leq w \leq q-1$, the information map which assigns $\lceil n/(q-w) \rceil$ bits of information to each subkey is w-secure for $\mathcal{K}$. With this information map, $\rho_P = (q^2-1)\lceil n/(q-w) \rceil$ for each node P. If n and w are chosen appropriately, then $\lceil n/(q-w) \rceil$ becomes equal to one and the bound is met.

It seems reasonable to conjecture that the second bound $\rho_P \geq (w+1)(n+w/2)$ of Theorem 5.2 could be improved upon: however, it is good enough to show that in general the node storage of a KDP will be inferior to that of the KDSs constructed by Blom, under the assumptions at the beginning of Section 4.

References

[1] R. Blom, An optimal class of symmetric key generation systems, in *Advances in Cryptology*: *Proceedings of Eurocrypt* 84, vol. 209 of Lecture Notes in Computer Science, Springer-Verlag, Berlin, 1985, pp. 335–338.

[2] B. Bollobás, *Combinatorics*, Cambridge University Press, Cambridge, 1986.

[3] M. Dyer, T. Fenner, A. Frieze, and A. Thomason, On key storage in secure networks, *Journal of Cryptology*, vol. 8 (1995), pp. 189–200.

[4] L. Gong and D.J. Wheeler, A matrix key distribution scheme, *Journal of Cryptology*, vol. 2 (1990), pp. 51–59.

[5] D.R. Hughes and F.C. Piper, *Design Theory*, Cambridge University Press, Cambridge, 1988.

[6] C.J. Mitchell, Combinatorial techniques for key storage reduction in secure networks, Technical memo, Hewlett Packard Laboratories, Bristol, 1988.

[7] C.J. Mitchell and F.C. Piper, Key storage in secure networks, *Discrete Applied Mathematics*, vol. 21 (1988), pp. 215–228.

[8] C.M. O'Keefe, Key distribution patterns using Minkowski planes, *Designs, Codes and Cryptography*, vol. 5 (1995), pp. 261–267.

[9] K.A.S. Quinn, Some constructions for key distribution patterns, *Designs, Codes and Cryptography*, vol. 4 (1994), pp. 177–191.

[10] D.R. Stinson, On some methods for unconditionally secure key distribution and broadcast encryption, *Designs, Codes and Cryptography*, vol. 12, no. 3 (1997), pp. 215–243.