Secure Re-keying and Collusion Prevention Using Block Designs

Nathaniel Karst and Stephen B. Wicker, *Senior Member, IEEE*

***Abstract*—We expand and generalize the work of Eltoweissy, Heydari and Morales on secure re-keying in group communications systems. Eltoweissy *et al.* introduced (n, k, m) exclusion basis systems (EBSs) in which k keys are administered to each of n users such that the entire network can be re-keyed with m messages after a single user leaves the system. We generalize their work by proving that an arbitrary simple t-(v, b, r, k, λ) block design generates an EBS. We go on to show that the number re-key messages can be reduced if we restrict our attention to square 2-(v, k, λ) designs. Furthermore, we show that square 2-(v, k, λ) designs generate EBSs in which fewer than $\lceil k/\lambda \rceil$ users cannot successfully collude to illicitly decipher transmissions.**

I. Introduction

The problem of key distribution has become increasingly important with the advent of remote data collection systems that rely on the Internet and other forms of public data transport. Examples of such systems include SCADA remote terminal units, "smart" residential electric meters and home health monitoring systems. When network membership is allowed to be fluid, the problem of effective and efficient key distribution becomes more complex. After a user leaves, the base station must securely distribute replacement keys to the remaining privileged members. In systems with high membership turnover or limited bandwidth, minimizing the number of re-key messages becomes a primary concern.

Eltoweissy, Heydari and Morales introduced the concept of an exclusion basis system as a re-keying mechanism in hierarchical networks [1]. Here, the base station administers key chains of size k to n users such that at most m messages are needed to distribute replacement keys if a user leaves the system. Eltoweissy *et al.* show that a necessary and sufficient condition for the existence of an (n, k, m) EBS is $\binom{k+m}{k} \geq n$ and provide a construction procedure for so-called *canonical* EBSs that support this upper bound on number of users. Redwine introduced mechanisms for ejecting multiple users based on combinatorial properties of canonical EBSs, and gave a brief collusion analysis for canonical EBSs [8].

In this paper, we generalize the results of Eltoweissy *et al.* by constructing EBSs from combinatorial block designs. The modern discussion of block designs stems from the study of experimental design in the middle of the last century, though questions concerning block designs can be found as early as the 1800s (*e.g.*, Kirkman's school girl problem and others [3, Section 1.2]). In Section II, we formally define EBSs and block designs, and go on to prove that any simple block design generates an EBS. In Section III, we show that generating EBSs from carefully selected families of block designs can provide some collusion prevention at the cost of the number of users supported. Section IV features a sample application and reviews its performance compared to appropriate EBSs. In Section V, we review our results and point to potentially interesting areas of further research.

N. Karst is a graduate student in the Center for Applied Mathematics at Cornell University. Email: njk46@cornell.edu.

S. Wicker is a professor in the Department of Electrical and Computer Engineering at Cornell University. Email: wicker@ece.cornell.edu.

The authors are supported in part by the National Science Foundation TRUST Science and Technology Center.

II. Block Designs as Exclusion Basis Systems

Definition 2.1: A collection Γ of subsets of a set $Y = \{1, 2, \ldots, n\}$ is a (n, k, m) *exclusion basis system* (EBS) if for all $d \in Y$, (1) d appears in at most k subsets in Γ; (2) there exist exactly m subsets $A_{d,1}, A_{d,2}, \ldots, A_{d,m} \in \Gamma$ such that $\cup_{i=1}^{m} A_{d,i} = Y \setminus \{d\}$.

In this formulation each element in Y corresponds to a user, and each subset in Γ corresponds to a key, with incidence indicating that a user has access to a key. So (1) implies that each user has at most k keys; part (2) implies that for any user d, the network can be re-keyed to exclude user d with exactly m messages by broadcasting replacement keys encrypted with the keys represented by $A_{d,i}, i = 1, 2, \ldots, m$.

Example 2.1: Let $Y = \{1, 2, \ldots, 10\}$ and suppose Γ is the collection of subsets

$$\Gamma = \{\{1,2,3,4,5,6\}, \{1,2,3,7,8,9\}, \{1,4,5,7,8,10\}, \{2,4,6,7,9,10\}, \{3,5,6,8,9,10\}\}.$$

One can verify that Γ is a canonical (10,3,2) EBS.

Definition 2.2: An ordered pair $D = (X, \mathcal{B})$ is said to be a t-(v, b, r, k, λ) *design* if (1) $\mathcal{B}$ consists of b subsets of X, each having cardinality k; (2) every element and t-subset of the v-set X occurs in r and λ subsets in $\mathcal{B}$, respectively. We call the elements of X *points* and the subsets in $\mathcal{B}$ *blocks*.

Designs containing no repeated blocks are called *simple*. We will see that it is convenient to associate the blocks of a block design with key chains in an EBS. Hence, to avoid confusion arising from distinct users have identical key chains, we restrict our consideration to simple designs.

Example 2.2: Let $X = \{1, 2, \ldots, 7\}$ and suppose that $\mathcal{B}$ is the collection of sets

$$\mathcal{B} = \{\{1,2,4\}, \{2,3,5\}, \{3,4,6\}, \{4,5,7\}, \{5,6,1\}, \{6,7,2\}, \{7,1,3\}\}.$$

By inspection, $(X, \mathcal{B})$ is a 2-(7,7,3,3,1) design.

Block designs can be described by an *incidence matrix*. Let $D = (X, \mathcal{B})$ be a t-(v, b, r, k, λ) design with $X = \{x_i\}_{1 \le i \le v}$ and $\mathcal{B} = \{B_j\}_{1 \le j \le b}$. The incidence matrix of D is a $v \times b$ matrix $M = (m_{ij})$ with $m_{ij} = 1$ if $x_i \in B_j$ and $m_{ij} = 0$ otherwise.

Theorem 2.1: A t-(v, b, r, k, λ) design $D = (X, \mathcal{B})$ forms a $(b, k, v - k)$ EBS.

Proof: Let $M = (m_{ij})$ be the incidence matrix of the design D on $X = \{1, 2, \ldots, v\}$. To recast this design as a $(b, k, v-k)$ EBS $\Gamma = \{A_\ell\}_{1 \le \ell \le v}$, we say that a subset $A_i \in \Gamma$ contains an element $j \in \{1, 2, \ldots, b\}$ if and only if $m_{ij} = 1$. Each block $B_j \in \mathcal{B}$ is incident with exactly k points in X, so in the derived EBS element $j \in \{1, 2, \ldots, b\}$ is contained in k subsets of Γ. Hence, EBS definition part (1) is satisfied.

Now, fix a block B_j. Since D has no repeated blocks and since all blocks have equal cardinality, every block except B_j contains at least one point in $X \setminus B_j$. Then in the derived EBS, any $\ell \in \{1, 2, \ldots, b\}, \ell \neq j$, is in A_i for some $i \in X \setminus B_j$, and $j \notin A_i$ for any $i \in X \setminus B_j$. Hence, it suffices to take the union $\cup_{(X \setminus B_j)} A_i$ to obtain $\{1, \ldots, b\} \setminus \{j\}$. Since B_j was chosen arbitrarily, EBS definition part (2) is satisfied with $m = |X \setminus B_j| = v - k$. We conclude that a t-(v, b, r, k, λ) design forms a $(b, k, v - k)$ EBS. ∎

Corollary 2.1: The set $X = \{1, 2, \ldots, v\}$ together with the collection $\mathcal{B}$ of all k-subsets of X, known as the (v, k) *complete design*, generates the canonical EBS with parameters $\left(\binom{v}{k}, k, v - k\right)$.

The class of block designs in which $t = 2$, $b = v$ and $k = r$ is of special interest both in theory and in applications. We call such designs *square* 2-(v, k, λ) designs. With these parameters, we have the dual of the incidence requirement:

Lemma 2.1 ([3]): In a square 2-(v, k, λ) design $D = (X, \mathcal{B})$, any distinct blocks B_i and B_j satisfy $|B_i \cap B_j| = \lambda$.

Restricting our attention to EBSs generated by square 2-(v, k, λ) designs can provide a reduction in re-key messages.

Theorem 2.2: A square 2-(v, k, λ) design forms a $(v, k, 2k - 3)$ EBS when $\lambda = 1$ and a $(v, k, k - 2)$ EBS when $\lambda > 1$.

Proof: Suppose the user with key chain $B_i \in \mathcal{B}$ leaves the network. Let B_j be any other key chain. Since D is a square 2-design, Lemma 2.1 implies that block B_j will share λ common points with all other blocks in $\mathcal{B}$. Then $(B_j \setminus B_i) \cap B_k$ is non-empty except for those blocks B_k for which $B_k \cap B_j = B_i \cap B_j$. There are exactly $\Lambda_\lambda = \lambda \binom{v-\lambda}{2-\lambda} / \binom{k-\lambda}{2-\lambda}$ such blocks if $\lambda \le 2$ and at most $\Lambda_2 = \lambda$ such blocks if $\lambda > 2$ [3, Theorem 3.3]. So by encrypting replacement keys with the $k - \lambda$ keys in $B_j \setminus B_i$, we can securely re-key all but $\Lambda_1 = k$ users if $\lambda = 1$ and all but at most $\Lambda_2 = \lambda$ users if $\lambda > 1$.

In both cases, two of the remaining blocks are B_i and B_j themselves. Since any two blocks share exactly λ common points, and since the any pair of remaining blocks already have these λ points accounted for in $B_i \cap B_j$, it is necessary and sufficient to include one point from each of the remaining blocks except B_i and B_j. Hence, if $\lambda = 1$, then

$$(|B_j| - |B_i \cap B_j|) + (\Lambda_1 - 2) = (k - 1) + (k - 2) = 2k - 3$$

points from X cover all blocks in $\mathcal{B}$ not equal to B_i, and if $\lambda > 1$, then

$$(|B_j| - |B_i \cap B_j|) + (\Lambda_2 - 2) = (k - \lambda) + (\lambda - 2) = k - 2$$

points suffice. ∎

Potentially interesting candidates for application are square 2-designs generated by difference sets (*e.g.*, see Wallis [7, Chapter 5]). Here the blocks are cyclic shifts of one another. Using such a design would help reduce computation and storage in the base station. These considerations are particularly important if the base station is a wireless device.

III. Collusion Prevention

In many secure group communication systems, preventing users from colluding to illicitly decipher messages is a priority. One important metric is the minimum number of colluders necessary to impersonate either a privileged user or the broadcaster itself. Redwine surveys this and other security questions in the context of EBSs [8].

The key chains generated by the (n, k, m) EBS construction algorithm of Eltoweissy *et al.* consist all possible combinations of k keys from a key pool of size $k + m$ [1]. While affording an optimal number of users for fixed key chain size and re-key messages, this scheme is highly susceptible to collusion, exactly because any k keys form a valid key chain.

Including fewer k-subsets as possible key chains would naturally make it harder for colluders to succeed. However, arbitrarily choosing k-subsets for removal would make quantifying collusion prevention in the resulting network difficult. Instead, we form a collection valid key chains using the blocks of a block design. The regular structure of these designs allows us to make guarantees as to their collusion prevention and to concretely compare the results with the canonical EBSs.

Theorem 3.1: In a EBS generated from a square 2-(v, k, λ) design $D = (X, \mathcal{B})$, $c \ge \lceil k/\lambda \rceil$ colluders are required to forge a valid key chain.

Proof: Suppose some number of colluders wish to forge key chain $B_i \in \mathcal{B}$. By Lemma 2.1, we have $|B_i \cap B_j| = \lambda$ for all $B_j \in \mathcal{B}$ not identical to B_i. Suppose there are c colluders with key chains $B_1, B_2, \ldots, B_c \in \mathcal{B}$. If these colluders can pool their keys to form B_i, then

$$k = |B_i| \le \left| \bigcup_{j=1}^{c} (B_i \cap B_j) \right| \le \sum_{j=1}^{c} |B_i \cap B_j| = c\lambda,$$

where the right-most inequality follows from the union bound. Since B_i was chosen arbitrarily, at least $c = \lceil k/\lambda \rceil$ colluders are required to successfully impersonate any user. ∎

Another property of block designs that may be useful in applications is resolvability. A t-(v, b, r, k, λ) block design is *resolvable* if its b blocks can be partitioned into classes of size v/k such that every point is contained in exactly one block in every class. If in addition any two blocks from different classes have constant intersection cardinality $\mu = k^2/v$, then we call the block design *affine*. EBSs generated from affine block designs provide collusion prevention.

Theorem 3.2: In an EBS derived from an affine t-(v, b, r, k, λ) design, v/k colluders are required to forge a valid key chain.

Proof: Fix B_i as the key chain that the colluders wish to impersonate, and let $\{B_j\}_{1\leq j\leq c} \subseteq \mathcal{B}$ represent the collection of colluders. In the case of successful collusion,

$$k = |B_i| \leq \left| \bigcup_{j=1}^{c} (B_i \cap B_j) \right| \leq \sum_{j=1}^{c} |B_i \cap B_j| \leq c\mu,$$

where the middle inequality follows from the union bound, and the right-most inequality follows from the fact that some colluders may be found in the same class as the user with key chain B_i. Hence, the bound can be met with equality if and only if all colluders are in a single given class other than that of B_i. In this case, $c = k/\mu = v/k$. ∎

IV. Application and Performance

Suppose we wish to design a secure re-keying solution for a wireless sensor network in which the base station can store only 25 keys, and every other node can store a key chain containing only 5 keys. Theorem 2.2 and Theorem 3.1 claim that if we can produce a square 2-(v, k, λ) design with $v \leq 25$ and $k \leq 5$, then we can construct an appropriate EBS that provides some collusion protection. To find such a design, we turn two classes of well-studied block designs.

An *affine plane* of order n is a 2-$(n^2, n^2 + n, n + 1, n, 1)$ design, and a *projective plane* of order n is a square 2-$(n^2 + n + 1, n + 1, 1)$ design [3]. To date, construction procedures exist only for affine and projective planes with n any prime power. These designs are particularly useful for applications, because they admit a simple construction algorithm:

(1) For fixed prime power n, generate $n - 1$ mutually orthogonal Latin squares of side n.
(2) Use the Latin squares from (1) to generate an affine plane.
(3) Perform a simple adjunction procedure to transform the affine plane from (2) into a projective plane.

For a complete presentation of these results and constructions, see Stinson [6, Section 6.4.1]. We note that the affine plane manufactured in (2) is in the class of affine resolvable designs mentioned in Section III. So according to Theorem 3.2, we could stop at (2) and construct an EBS in which at least n colluders are needed to forge a key chain. However, since affine planes are not square, we have much weaker results concerning sufficient re-key messages. Hence, we will generate an EBS from the (square) projective plane formed in (3).

For digital applications it is natural to choose to work over the binary field $GF(2)$. Computation shows that taking $n = 4$ satisfies the chosen constraints; the resulting projective plane EBS supports $b = 21$ users with $k = 5$ keys each. Theorem 2.2 claims at most $2k - 3 = 7$ messages are sufficient to re-key the network after ejecting a single member. Theorem 3.1 guarantees that $c = \lceil k/\lambda \rceil = n+1 = 5$ colluders are required to successfully impersonate another user.

Many canonical EBSs satisfy the proposed specifications, and comparing a projective plane EBS to any one canonical EBS is insufficient. For instance, a canonical EBS satisfying identical constraints can support as many as $\binom{25}{5} = 53130$ users and require as many as 20 re-key messages. Matching the key chain size and number of re-key messages to that of the projective plane EBS gives a canonical EBS that supports $\binom{12}{5} = 792$ users. Yet another canonical solution prescribes a key pool of size 7 and key chains of size 3, leading to $\binom{7}{3} = 35$ supported users and 4 re-key messages. Note that in this last case, the canonical solution both supports more users and requires fewer re-key messages than an identically constrained EBS constructed from a projective plane.

The stark and universal difference between an EBS constructed from a projective plane and equivalent canonical solutions is collusion protection. In the projective plane EBS, $n + 1 = 5$ colluders are required before successful collusion can take place. In contrast, it is possible that 5 colluders in any of the canonical EBSs proposed above could impersonate *any* other user or even the base station itself. While supporting far fewer users, EBS generated from square 2-designs offer a more moderate solution: relatively few re-key messages and relatively strong collusion protection.

V. Discussion and Conclusion

Eltoweissy *et al.* present a family of EBSs which support a maximum number of users for a fixed key chain size and re-key messages. We provide a combinatorial construction procedure based on block designs. We proved that an EBS generated from an arbitrary t-(v, b, r, k, λ) design is in the worst case equivalent to those presented by Eltoweissy *et al.* with respect to re-key messages for fixed key chain and fixed key pool sizes. In addition, we showed that EBSs generated from square and affine designs provide some collusion prevention at the cost of the number of users supported. This procedure can be viewed as removing blocks from an appropriately sized complete design; this framework may be useful in further research on the subject of collusion prevention in EBSs.

We note that while we have proved that square and resolvable block designs generate EBSs with interesting and useful properties, other combinatorial aspects of block designs, such as group-divisibility, might be exploited in applications. And while we have shown that square designs admit a reduction in re-key messages in some cases, further combinatorial insight may also provide improved bounds on re-key messages for more general classes of block designs.

References

[1] M. Eltoweissy, M. Heydari, L. Morales. "Combinatorial optimization of group key management," *J. of Network and Systems Management*, vol. 12, no. 1, pp. 33–50, 2004.

[2] S. Çamtepe, B. Yener. "Combinatorial design of key distribution mechanisms for wireless sensor networks," *IEEE/ACM Trans. on Networking*, vol. 15, no. 2, pp. 346–358, 2007.

[3] C. Colbourn, J. Dinitz, editors. "The CRC handbook of combinatorial designs," CRC Press, Boca Raton 1996.

[4] J. Lee, D. Stinson. "On the Construction of Practical Key Predistribution Schemes for Distributed Sensor Networks Using Combinatorial Designs," *ACM Trans. on Info. and Sys. Security*, vol. 11, no. 2, pp. 1–35, 2008.

[5] J. Lee, D. Stinson. "A combinatorial approach to key predistribution for distributed sensor networks," *IEEE Wireless Communications and Networking Conference*, vol. 2, pp. 1200-1205, 2005.

[6] D. Stinson. "Combinatorial designs: constructions and analysis," *Springer-Verlag*, New York, 2004.

[7] W. Wallis. "Introduction to combinatorial designs," *Champham & Hall/CRC*, Boca Raton, 2007.

[8] S. Redwine, Jr. "A logic for exclusion basis systems," *Proc. of the 37th Hawaii Inter. Conf. on Sys. Sciences*, 2004.