Combinatorial Aspects of Key Distribution for Sensor Networks

Douglas R. Stinson

David R. Cheriton School of Computer Science
University of Waterloo

CanaDAM 2013

Monday, June 10, 2013

This talk is based on joint work with Kevin Henry, Jooyoung Lee and Maura Paterson.

Wireless Sensor Networks

- sensor nodes have limited computation and communication capabilities
- a network of 1000 – 10000 sensor nodes is distributed in a random way in a possibly hostile physical environment
- the sensor nodes operate unattended for extended periods of time
- the sensor nodes have no external power supply, so they should consume as little battery power as possible
- usually, the sensor nodes communicate using secret key cryptography
- a set of secret keys is installed in each node, before the sensor nodes are deployed, using a suitable key predistribution scheme (or KPS)
- nodes may be stolen by an adversary (this is called node compromise)

Two Trivial Schemes

1. If every node is given the same secret master key, then memory costs are low. However, this situation is unsuitable because the compromise of a single node would render the network completely insecure.
2. For every pair of nodes, there could be a secret pairwise key given only to these two nodes. This scheme would have optimal resilience to node compromise, but memory costs would be prohibitively expensive for large networks because every node would have to store $n - 1$ keys, where n is the number of nodes in the WSN.

Eschenauer-Gligor and Related Schemes

- In 2002, Eschenauer and Gligor proposed a probabilistic approach to key predistribution for sensor networks. For a suitable value of k, every node is assigned a random k-subset of keys chosen from a given pool of secret keys.
- In 2003, Chan, Perrig and Song suggested that two nodes should compute a pairwise key only if they share at least η common keys, where the integer $\eta \geq 1$ is a pre-specified intersection threshold. Such a pair of nodes is termed a link.
- Suppose that U_i and U_j have exactly $\ell \geq \eta$ common keys, say $\mathbf{key}_{a_1}, \ldots, \mathbf{key}_{a_\ell}$, where $a_1 < a_2 < \cdots < a_\ell$. Then they can each compute the same pairwise secret key,

$$K_{i,j} = h(\mathbf{key}_{a_1} \parallel \ldots \parallel \mathbf{key}_{a_\ell} \parallel i \parallel j),$$

using a key derivation function h that is constructed from a secure public hash function, e.g., SHA-1.

Attack Model

- The most studied adversarial model in WSNs is random node compromise.
- An adversary compromises a fixed number of randomly chosen nodes in the network and extracts the keys stored in them.
- Any links involving the compromised nodes are broken.
- However, this can also cause other links to be broken that do not directly involve the compromised nodes.
- A link formed by two nodes A_1, A_2, where $|A_1 \cap A_2| \geq \eta$, will be broken if a node $B \notin \{A_1, A_2\}$ is compromised, provided that $A_1 \cap A_2 \subseteq B$.
- If s nodes, say $B_1, \ldots, B_s$, are compromised, then a link A_1, A_2 will be broken whenever

$$A_1 \cap A_2 \subseteq \bigcup_{i=1}^{s} B_i.$$

Important Metrics

Storage requirements

The number of keys stored in each node, which is denoted by k, should be "small" (e.g., at most 100).

Network connectivity

The probability that a randomly chosen pair of nodes can compute a common key is denoted by $\mathbf{Pr_1}$. $\mathbf{Pr_1}$ should be "large" (e.g., at least 0.5).

Network resilience

The probability that a random link is broken by the compromise of s randomly chosen nodes not in the link is denoted by $\mathbf{fail(s)}$. We want $\mathbf{fail(s)}$ to be small: high resilience corresponds to a small value for $\mathbf{fail(s)}$. In this talk we consider $\mathbf{fail(1)}$.

Remark: As η is increased, $\mathbf{Pr_1}$ and $\mathbf{fail(1)}$ both decrease.

Deterministic Schemes

- In 2004, deterministic KPS were proposed independently by Camtepe and Yener; by Lee and Stinson; and by Wei and Wu.
- A suitable set system is chosen, and each block is assigned to a node in the WSN (the design and the correspondence of nodes to blocks is public).
- The points in the block are the indices of the keys given to the corresponding node.
- Probabilistic schemes are analyzed using random graph theory, and desirable properties hold with high probability.
- Deterministic schemes can be proven to have desirable properties, and they have more efficient algorithms for shared-key discovery than probabilistic schemes.

Combinatorial Set Systems (aka Designs)

- A set system is a pair $(X, \mathcal{A})$, where the elements of X are called points and $\mathcal{A}$ is a set of subsets of X, called blocks.
- We pair up the blocks of the set system with the nodes in the WSN.
- The points in the block are the key identifiers of the keys given to the corresponding node.
- The degree of a point $x \in X$ is the number of blocks containing x
- $(X, \mathcal{A})$ is regular (of degree r) if all points have the same degree, r; then each key occurs in r nodes in the WSN.
- If all blocks have size k, then $(X, \mathcal{A})$ is said to be uniform (of rank k); then each node is assigned k keys.
- A (v, b, r, k)-configuration is a set system $(X, \mathcal{A})$ where $|X| = v$ and $|\mathcal{A}| = b$, that is uniform of rank k and regular of degree r, such that every pair of points occurs in at most one block.
- In a configuration, it holds that $vr = bk$.

Toy Example

We list the blocks in a $(7, 7, 3, 3)$-configuration (a projective plane of order 2) and the keys in a corresponding KPS:

node	block	key assignment
N_1	$\{1, 2, 4\}$	k_1, k_2, k_4
N_2	$\{2, 3, 5\}$	k_2, k_3, k_5
N_3	$\{3, 4, 6\}$	k_3, k_4, k_6
N_4	$\{4, 5, 7\}$	k_4, k_5, k_7
N_5	$\{1, 5, 6\}$	k_1, k_5, k_6
N_6	$\{2, 6, 7\}$	k_2, k_6, k_7
N_7	$\{1, 3, 7\}$	k_1, k_3, k_7

The actual values of keys are secret, but the lists of key identifiers (i.e., the blocks) are public.
In this example, $\mathbf{Pr_1} = 1$ and $\mathbf{fail(1)} = 1/5$.

Properties of Configuration-based KPS

- For a configuration-based KPS, we take $\eta = 1$.
- Every block intersects $k(r-1)$ blocks in one point and is disjoint from all the other blocks.
- Therefore
$$\mathbf{Pr_1} = \frac{k(r-1)}{b-1}.$$
- A link L is defined by two blocks that intersect in one point, say x.
- There are $r-2$ other blocks that contain x; the corresponding nodes will compromise the link L.
- Therefore,
$$\mathbf{fail(1)} = \frac{r-2}{b-2}.$$
- There is a tradeoff between $\mathbf{Pr_1}$ and $\mathbf{fail(1)}$, which is quantified by computing the ratio $\rho = \mathbf{Pr_1}/\mathbf{fail(1)}$:
$$\rho = \frac{k(b-2)(r-1)}{(b-1)(r-2)} \approx k.$$

Transversal Designs

- Lee and Stinson (2005) proposed using transversal designs to construct KPS.
- Let n, k and t be positive integers
- A transversal design TD(t, k, n) is a triple $(X, \mathcal{H}, \mathcal{A})$, where X is a finite set of cardinality kn, $\mathcal{H}$ is a partition of X into k parts (called groups) of size n, and $\mathcal{A}$ is a set of k-subsets of X (called blocks), which satisfy the following properties:
  1. $|H \cap A| = 1$ for every $H \in \mathcal{H}$ and every $A \in \mathcal{A}$, and
  2. every t elements of X from different groups occurs in exactly one block in $\mathcal{A}$.
- Bose-Bush bound: When $t = 2, 3$, a TD(t, k, n) exists only if $k \leq n + t - 1$.

An Easy Construction for Transversal Designs

- Suppose that p is prime and $t \leq k \leq p$.
- Define
$$X = \{0, \ldots, k-1\} \times \mathbb{Z}_p.$$
- For every ordered t-subset $\mathbf{c} = (c_0, \ldots, c_{t-1}) \in (\mathbb{Z}_p)^t$, define a block
$$A_{\mathbf{c}} = \left\{ \left(x, \sum_{i=0}^{t-1} c_i x^i \right) : 0 \leq x \leq k-1 \right\}.$$
- Let
$$\mathcal{A} = \{A_{\mathbf{c}} : \mathbf{c} \in (\mathbb{Z}_p)^t\}.$$
- Then $(X, \mathcal{A})$ is a TD(t, k, p).
- The construction can be adapted to any finite field $\mathbb{F}_q$, where q is a prime power.
- These transversal designs are equivalent to Reed-Solomon codes.

Properties of KPS from TDs with $t = 2$

- A TD$(2, k, n)$ is an (nk, n^2, n, k)-configuration.
- Therefore

$$\mathbf{Pr_1} = \frac{k(n-1)}{n^2-1} = \frac{k}{n+1} \quad \text{and} \quad \mathbf{fail(1)} = \frac{n-2}{n^2-2}.$$

- Since the set system is a configuration, we have $\rho \approx k$.
- Benefit: We can make $\mathbf{Pr_1}$ arbitrarily close to 1.
- Benefit: Shared-key discovery is very efficient, due to the underlying algebraic structure of the TDs.
- Drawback: The network size is n^2, which may not be large enough for "reasonable" values of n.
- Drawback: The ratio $\rho \approx k$ is a bit small for many applications (this applies to any configuration-based KPS).

Properties of KPS from TDs with $t = 3$, $\eta = 2$

- We can base a KPS on a TD$(3, k, n)$ with $\eta = 1$ or 2.
- When $\eta = 2$, we have

$$\mathbf{Pr_1} = \frac{k(k-1)}{2(n^2+n+1)} \quad \text{and} \quad \mathbf{fail(1)} = \frac{n-2}{n^3-2}.$$

- Drawback: The maximum value of $\mathbf{Pr_1}$ is about $1/2$.
- Drawback: Shared-key discovery is less efficient (but still reasonable).
- Benefit: The network size is n^3, which is quite large, even for "reasonable" values of n.
- Benefit: The ratio $\rho \approx k^2/2$ is now considerably larger.

Flexibility of Parameters

- The network size for a TD-based KPS is n^2 when $t = 2$ and n^3 when $t = 3$.
- For the "easy" constructions, we want n to be a prime power.
- The traditional viewpoint with respect to combinatorial KPS is that if a specific network size m is desired, then it suffices to choose parameters to give a scheme for a network of size greater than m and simply discard excess nodes.
- Bose, Dey and Mukerjee (2013) disagree with this viewpoint, saying "if we then discard the unnecessary node allocations to get the final scheme for use, this final scheme will not preserve the $\mathbf{Pr_1}$ and $\mathbf{fail(s)}$ values of the original scheme and hence the properties of the final scheme in this regard can become quite erratic".
- We dispute this statement, and we have two ways to counter their argument.

Flexible KPS from TDs with $t = 2$

- When n is a prime power, the "easy" TD$(2, k, n)$ can be resolved into n parallel classes, each containing n blocks.
- Suppose we take ℓ of the n parallel classes.
- We obtain an $(nk, n\ell, \ell, k)$-configuration.
- Therefore

$$\mathbf{Pr_1} = \frac{k(\ell - 1)}{\ell n - 1} \quad \text{and} \quad \mathbf{fail(1)} = \frac{\ell - 2}{\ell n - 2}.$$

- As long as ℓ is not very small, we have a KPS whose values of $\mathbf{Pr_1}$, $\mathbf{fail(1)}$ and ρ are similar to what they were before; the value of k is unchanged.
- But we can now accommodate many possible network sizes for a given value of n: any multiple of n from $2n$ to n^2.

Flexible KPS from TDs with $t = 3$

- When n is a prime power, the "easy" TD$(3, k, n)$ can be resolved into n TD$(2, k, n)$'s, each containing n^2 blocks.
- Suppose we take ℓ of these n TD$(2, k, n)$'s.
- When $\eta = 2$, we have

$$\mathbf{Pr_1} = \frac{k(k-1)(\ell-1)}{2(\ell n^2 - 1)} \quad \text{and} \quad \mathbf{fail(1)} = \frac{\ell - 2}{\ell n^2 - 2}.$$

- Again, as long as ℓ is not very small, we have a KPS whose values of $\mathbf{Pr_1}$, $\mathbf{fail(1)}$ and ρ are similar to what they were before; the value of k is unchanged.
- We can now accommodate many possible network sizes for a given value of n: any multiple of n^2 from $2n^2$ to n^3.

Random Deletion of Nodes from a KPS

- Suppose we randomly delete nodes from a combinatorial KPS.
- Question: How are the values of $\mathbf{Pr_1}$ and $\mathbf{fail}(\mathbf{1})$ affected?
- Answer: Hardly at all. The concerns of Bose *et al.* seem to be unfounded.
- We did large numbers of experiments which showed unequivocally that the "random deletion" approach works very well in practice.
- There is some variation in the values of $\mathbf{Pr_1}$ and $\mathbf{fail}(\mathbf{1})$, but the standard deviation is very small.

Example: Connectivity of KPS derived from $\mathbf{TD}(2, 20, 109)$

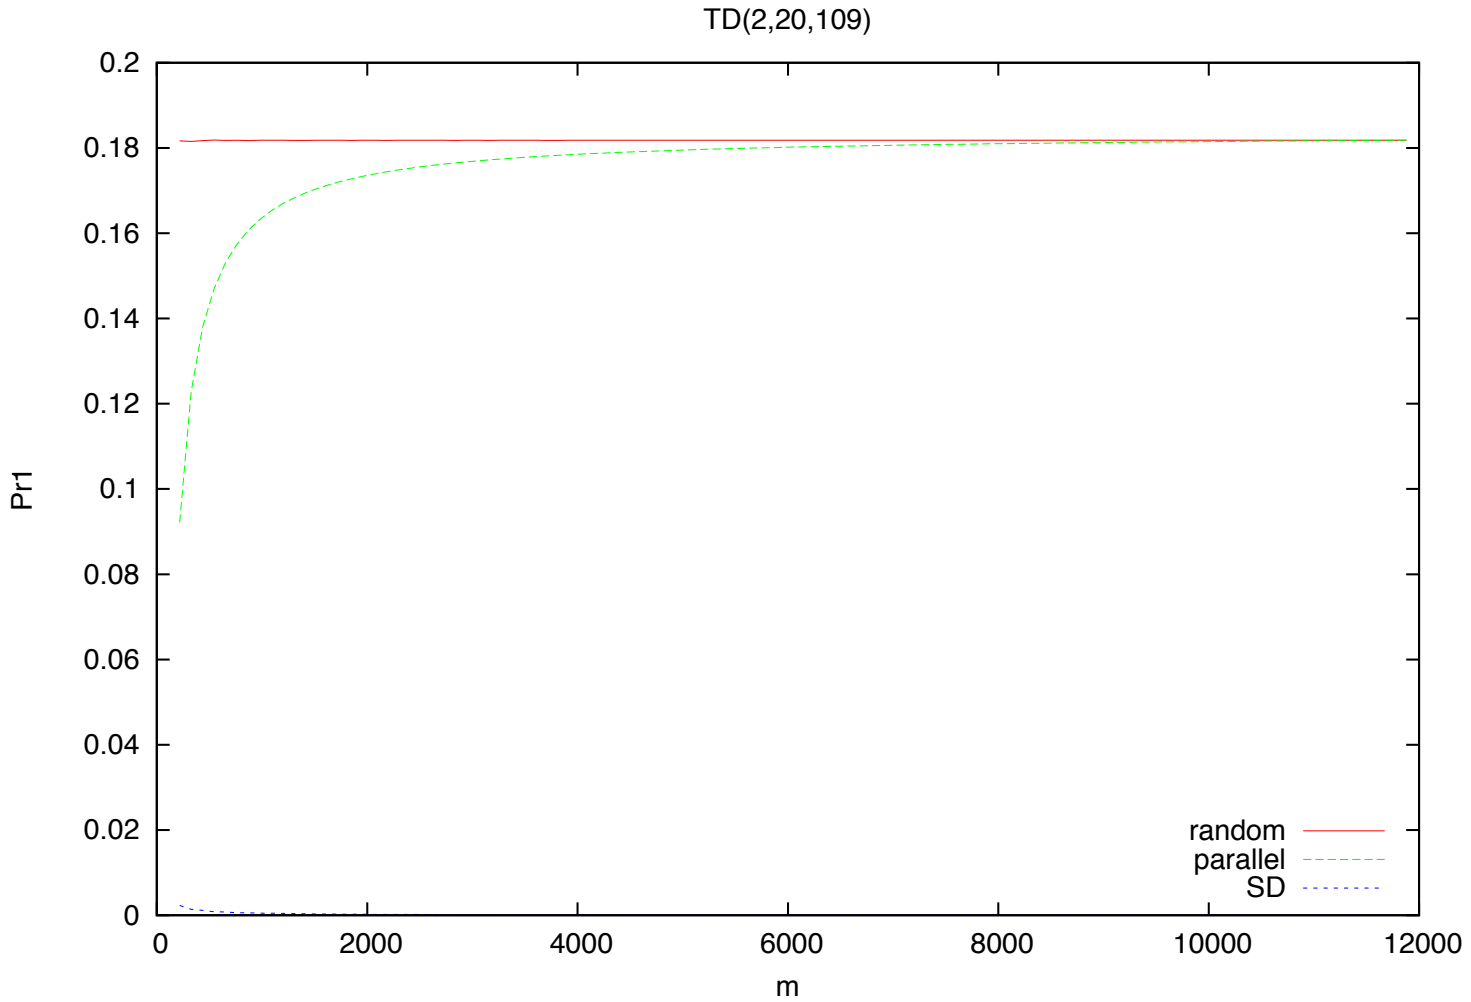

Example: Resilience of KPS derived from $\mathbf{TD}(2, 20, 109)$

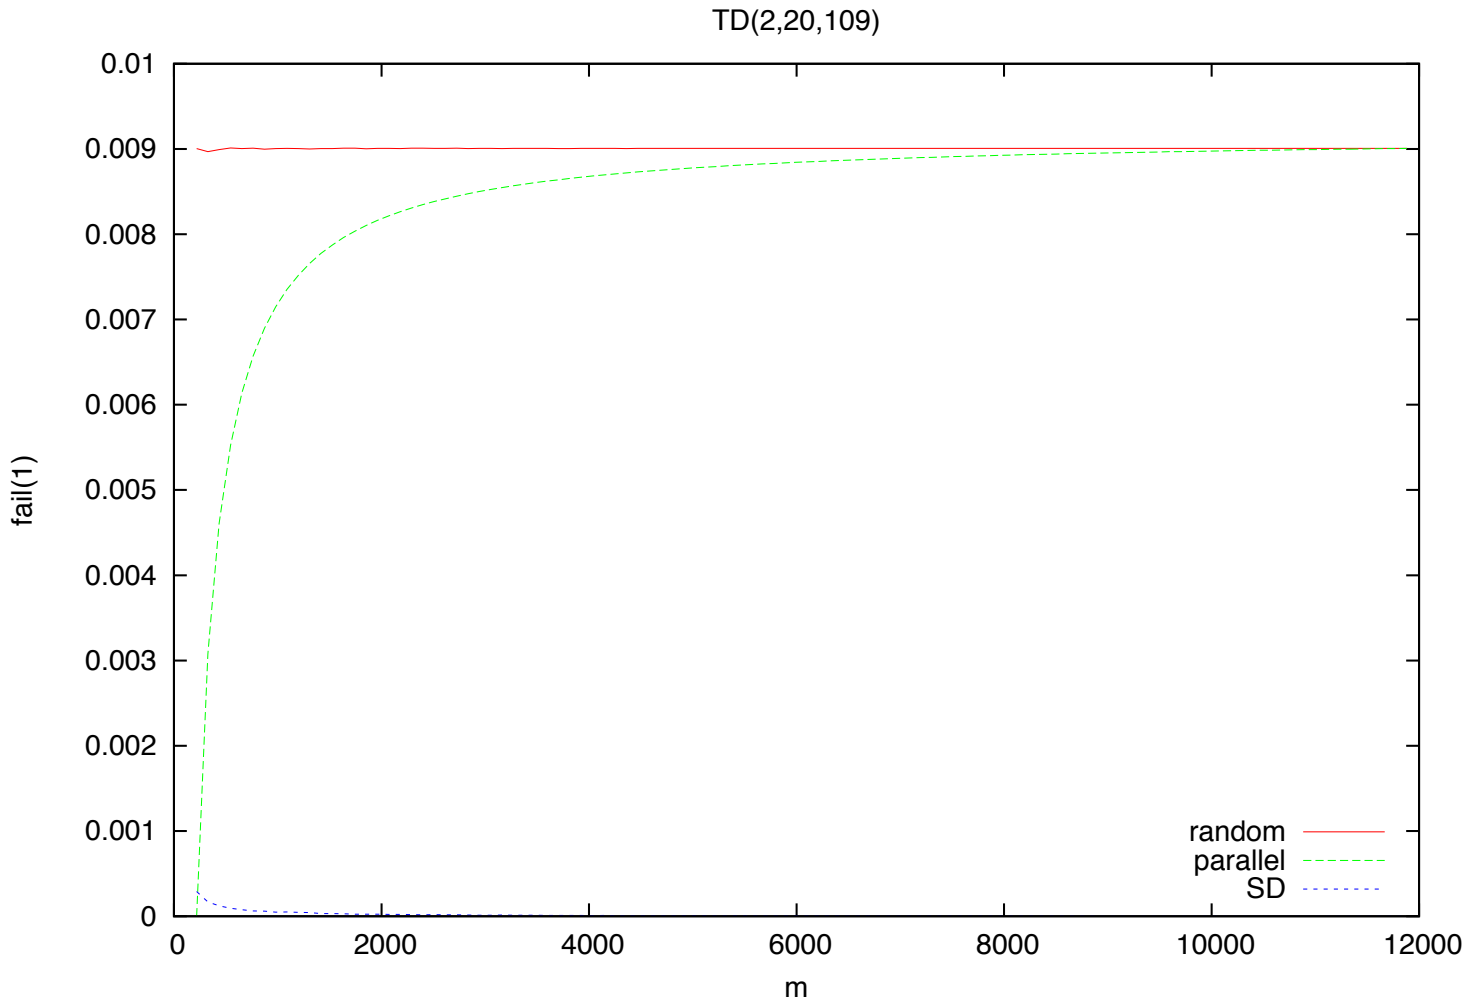

Example: Connectivity of KPS derived from $\mathbf{TD}(3, 20, 23)$

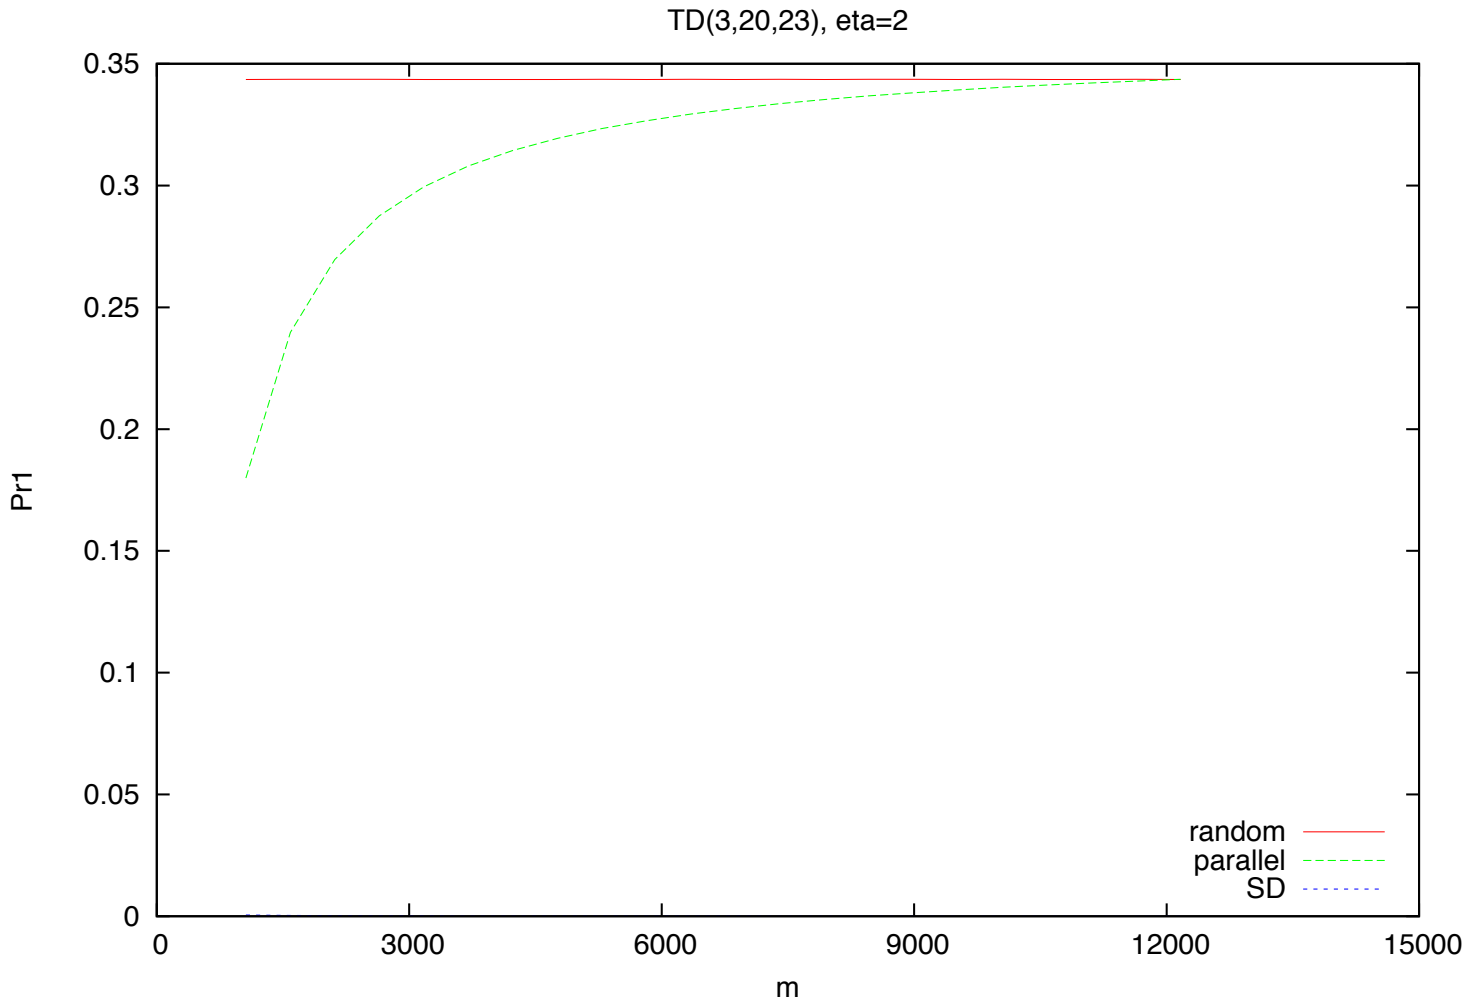

Example: Resilience of KPS derived from $\mathbf{TD}(3, 20, 23)$

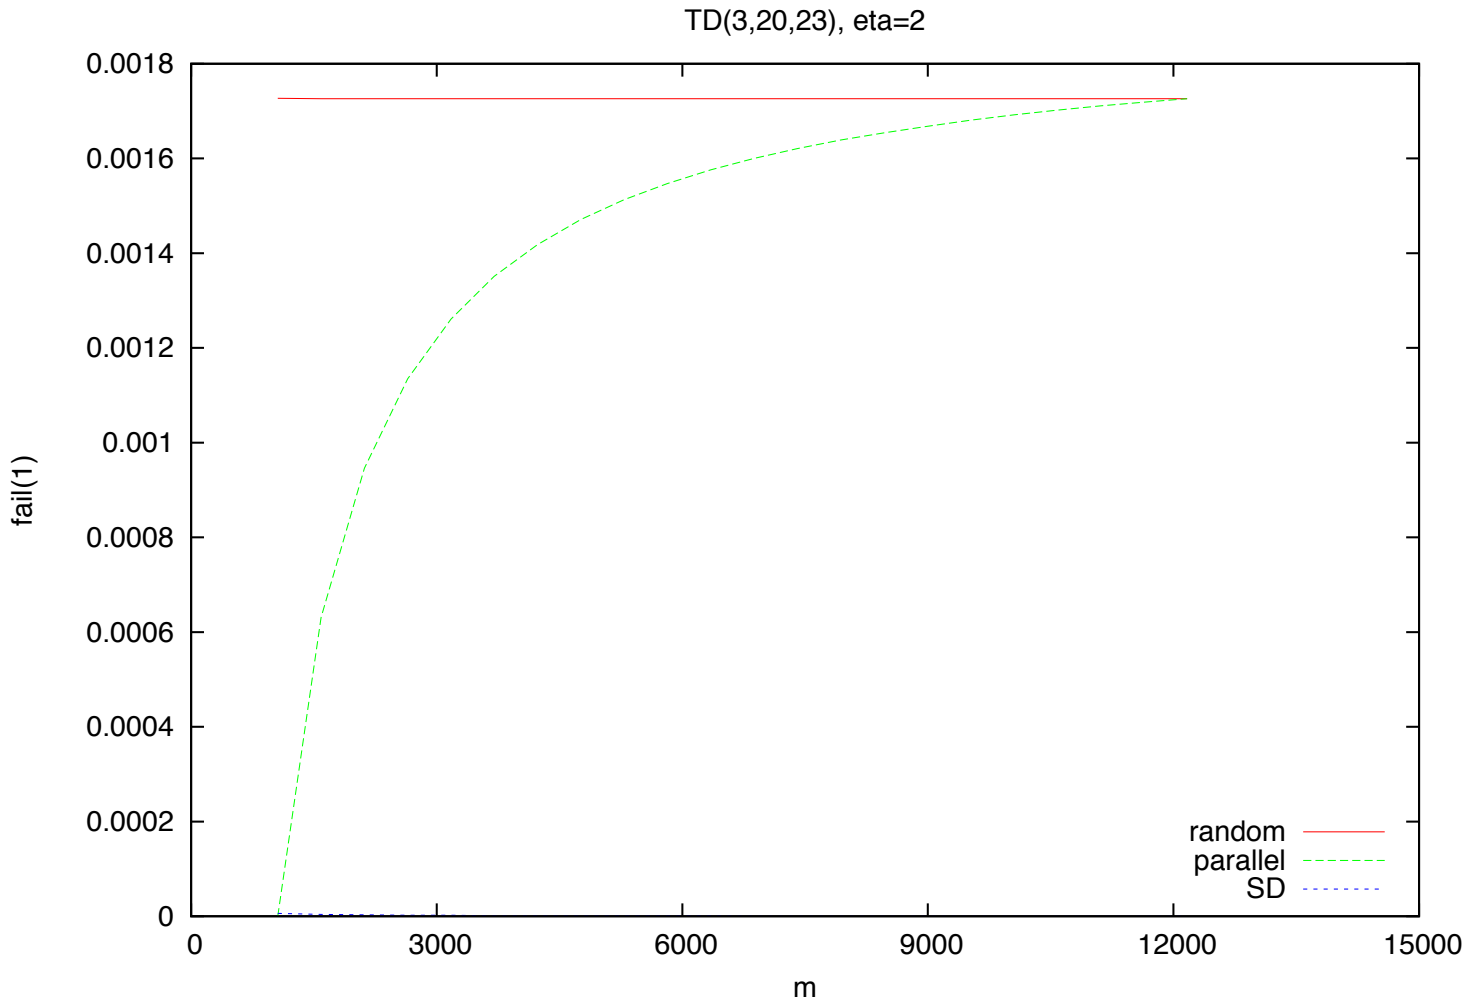

Using Less Regular Set Systems

- We have been employing schemes based on combinatorial structures (TDs, especially).
- Question: Could there be any advantage in using less "regular" structures to construct KPS?
- Suppose we use a set system with block size k where the maximum intersection of two blocks equals 1.
- We do not require that every point occurs in the same number of blocks.
- So we are relaxing the requirements of a configuration.
- Suppose that point i occurs in r_i blocks, for $1 \leq i \leq v$.
- Then $\sum r_i = bk$.

Properties of the Resulting KPS

- We have

$$\mathbf{Pr_1} = \frac{\sum_{i=1}^{v} r_i(r_i-1)}{b(b-1)}$$

and

$$\mathbf{fail(1)} = \frac{\sum_{i=1}^{v} r_i(r_i-1)(r_i-2)}{(b-2)\sum_{i=1}^{v} r_i(r_i-1)}.$$

- Therefore,

$$\rho = \frac{(b-2)\left(\sum_{i=1}^{v} r_i(r_i-1)\right)^2}{b(b-1)\sum_{i=1}^{v} r_i(r_i-1)(r_i-2)}.$$

- Conjecture (?) Assuming that $\sum_{i=1}^{v} r_i = bk$ is fixed, value of ρ is maximized when $r_1 = \cdots = r_v = bk/v$.

References

[1] M. Bose, A. Dey and R. Mukerjee. Key predistribution schemes for distributed sensor networks via block designs. *Designs, Codes and Cryptography* **67** (2013), 111–136.

[2] K. Henry, M. B. Paterson and D. R. Stinson. Practical approaches to varying network size in combinatorial key predistribution schemes. Preprint.

[3] J. Lee and D. R. Stinson. A combinatorial approach to key predistribution for distributed sensor networks. *IEEE Wireless Communications and Networking Conference (WCNC 2005)*, vol. 2, pp. 1200–1205.

[4] M. B. Paterson and D. R. Stinson. A Unified Approach to Combinatorial Key Predistribution Schemes for Sensor Networks. *Designs, Codes and Cryptography*, to appear.

thank you for your attention!