

Ultra-Lightweight Key Predistribution in Wireless Sensor Networks for Monitoring Linear Infrastructure

Keith M. Martin and Maura B. Paterson*

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

Abstract. One-dimensional wireless sensor networks are important for such security-critical applications as pipeline monitoring and perimeter surveillance. When considering the distribution of symmetric keys to secure the communication in such networks, the specific topology leads to security and performance requirements that are markedly distinct from those of the more widely-studied case of a planar network. We consider these requirements in detail, proposing a new measure for connectivity in one-dimensional environments. We show that, surprisingly, optimal results may be obtained through the use of extremely lightweight key predistribution schemes.

Keywords: wireless sensor networks; key management; one-dimensional networks

1 Introduction

The classical view of a wireless sensor network (WSN) is one of thousands of sensor nodes with a random physical distribution, such as would result from the nodes being scattered from an aeroplane. In practice, however, the specific sensing requirements of a given application impose a topology on the network that may differ significantly from this standard picture [12]. One example is the case of a one-dimensional network, in which the sensors are arranged in a line or ring. Topologies of this sort arise naturally from applications such as pipeline monitoring or perimeter surveillance, where it is necessary to take measurements at regular intervals along a lengthy piece of infrastructure. Security is important for such applications, hence it is desirable to use cryptographic techniques to secure the wireless communication. Symmetric primitives provide a less

* This author was supported by EPSRC grant EP/D053285/1.

computationally intensive solution than public-key techniques, but they require the sensor nodes to share common keys.

Key predistribution schemes (KPSs) are a widely-studied means of providing shared keys to the nodes [1, 8, 15]. However, most widely studied KPSs such as that of Eschenauer and Gligor [2] are designed for planar networks, whose properties are very different to those of a network with a one-dimensional topology. Nodes in a one-dimensional network are likely to have fewer neighbours (that is, nodes that lie within communication range of them) on average and the pattern of communication within the network will be quite different. Also, it is reasonable to suppose that the location of the sensors will be known, at least up to the order in which they occur along the network. In turn, the security requirements of a KPS for such a network are not the same as in the planar case. In particular, maintaining network connectivity is critical in one-dimensional networks, as we will demonstrate.

The main contribution of this paper is the identification of the precise requirements a KPS for a one-dimensional WSN, and the observation that it is possible to achieve these properties through the use of schemes with very low storage requirements. In Sect. 2 we discuss applications for one-dimensional WSNs. In Sect. 3 we consider the properties of such networks in detail, and in Sect. 4 we see that this leads to security and performance requirements that are markedly different from those of the classical scenario, particularly with respect to the connectivity of the network. We introduce a new measure, the *s-fallibility*, that more closely models the desirable connectivity properties of a KPS for a one-dimensional network. In Sect. 5 we examine how the *s-fallibility* of the network is affected by the communication range of the sensors, and in Sect. 6 we propose a key predistribution scheme that gives optimal *s-fallibility* with very low storage requirements. The advantages of such schemes are summarised in Sect. 7.

2 Applications Requiring the Use of One-Dimensional WSNs

The monitoring of an extended piece of infrastructure, such as a pipe, lends itself to the use of a one-dimensional sensor network. Pipelines carrying oil, gas or water are critical both in terms of their commercial importance and their impact on national security. Reasons for monitoring such pipelines include the detection of leaks, the measurement of seismic activity that has the potential to damage pipes, or the detection of

malicious activities such as sabotage or deliberate theft of pipes or their contents.

In any context there are advantages to using a wireless (as opposed to wired) network of sensors, such as greater ease of deployment and maintenance of the network. In the one-dimensional scenario, however, even stronger arguments for using a wireless network are provided by consideration of the reliability of the network. An attacker who cuts through the wire can entirely disrupt communication in a conventional network with linear topology. The use of multiple wires complicates the design and deployment of a network, and will not necessarily make it harder for an adversary to disconnect it. An appropriately designed wireless network, on the other hand, can withstand the loss of several sensors without losing connectivity. In the literature there are several proposals for the use of WSNs in such a context [3, 4, 10, 11, 13]. Similar considerations also apply to the monitoring of other types of linear infrastructure such as bridges or railway tracks.

A related application that leads naturally to the use of one-dimensional WSNs is that of perimeter surveillance (e.g. [14]). For example, sensors may be deployed for monitoring the condition of a fence, or for intrusion detection on the boundary of an unfenced region. The resulting networks differ from those required for pipeline monitoring in that the sensors are arranged in a ring, rather than in a line. In Sect. 6.1 we consider how this affects considerations of connectivity in such networks.

Various practical aspects of the performance of one-dimensional ad hoc networks have been studied [6, 7, 9, 16, 17]. However, the security of the communications in such networks has received comparatively little attention. In Secs. 3 and 4 we discuss how the properties of these networks lead to quite specific security requirements when considering the performance of key predistribution schemes.

3 Characteristics of One-Dimensional Sensor Networks

The properties of a one-dimensional WSN differ substantially from those of the two-dimensional analogue. Here we examine those properties that are particularly relevant to the design of KPSs.

Restricted number of neighbours: If we assume that each node has a particular communication range r , then for a given density of node deployment, the number of nodes within communication range of a sensor in a planar network is proportional to r^2 . In a one-dimensional network, however, this number is proportional to r .

Location knowledge: It is possible to consider two-dimensional networks in which there are differing degrees of a priori information about the nodes' locations, ranging from none at all, up to complete knowledge (see [8], for example). In the one-dimensional case it is reasonable to assume that the order in which the nodes occur along the network is known, since they are likely to be deployed sequentially along the object being monitored. In particular, this implies that the neighbours of each particular node are known with high probability prior to deployment.

Pattern of communication: In a one-dimensional network, information is constrained to flowing back and forth along the network. This has particular implications for aspects such as the capacity of the network [7], and the design of routing algorithms [17, 16].

Density of node deployment: Depending on the quantities that are being measured by the sensors, it is likely that density of node deployment required to ensure adequate sensing coverage will exceed the density required for the wireless network to be connected.

4 Security Considerations for Key Distribution in One-Dimensional Sensor Networks

When evaluating the performance of a KPS for a WSN, we are generally interested in the trade-off it provides between storage requirements, network connectivity, and resilience in the face of an adversary that can eavesdrop on all network traffic, as well as capture a certain number of nodes and extract any keys they contain. It is common to measure the connectivity in terms of the probability Pr_1 that two neighbouring nodes share a key, and the resilience in terms of the proportion of link keys that are compromised when a given number of nodes are captured at random [2, 5].

In the case of a one-dimensional WSN, the combination of nodes having a small number of neighbours with the fact that the expected neighbours are known prior to deployment suggests that it may be feasible to employ a KPS that assigns a distinct key to each pair of neighbouring nodes; we refer to this as the *local pairwise KPS*. This scheme has the advantage of having an optimal value of Pr_1 (since any two neighbours share a key with probability 1) and optimal resilience (since any key is shared by just two nodes, and thus the compromise of a node does not expose keys that pairs of uncompromised nodes rely on for communication).

Nevertheless, this is not the end of the story for key predistribution in one-dimensional WSNs. The quantity Pr_1 is useful in that it provides

a means of comparing in some sense the relative network connectivity achieved by KPSs in the case of a planar network with no location knowledge, where the formulation of absolute measures of connectivity is problematic. However, in the one-dimensional case, our detailed knowledge of the network topology enables us to quantify more precisely the connectivity behaviour that is desirable from the point of view of network functionality, in terms of a quantity we refer to as the *s-fallibility* of the network. We will see that while the local pairwise KPS performs well with respect to this measure, there exist schemes that perform equally well, yet require less storage.

In what follows, we assume our network consists of n identical nodes arranged in a line at regularly spaced intervals (in Sect. 6.1 we will consider how the situation changes if the nodes lie in a ring.) We suppose each node can communicate with all nodes located within distance r (where the distance between two adjacent nodes is taken to be 1). We restrict our attention to KPSs in which each key is shared by at most two nodes, as this provides optimal resilience. Furthermore, we focus on KPSs that can be used to distribute keys to networks with arbitrarily large values of n : the available location knowledge makes it possible to avoid assigning shared keys to pairs of nodes that are not within communication range, and hence extending a scheme to a large number of nodes does not adversely affect the storage requirements of any given node. Finally, instead of considering only adversaries that capture nodes at random, we analyse resistance against a “worst-case” adversary that chooses which nodes to capture based on the amount of damage it can do to the network.

Sensor nodes are small, cheap devices that are deployed with sufficient density that the performance of the network is not affected if a number of them fail. Thus the loss of a small number of sensors through adversarial capture is not in itself critical. However, it becomes a serious problem if the adversary is able to capture nodes in such a way as to prevent large sections of the network from communicating with each other. We formalise this notion in the following definition.


Definition 1. *Two disjoint sets S_1 and S_2 of nodes in a one-dimensional WSN are isolated from each other if no node in S_1 is in range of and shares an uncompromised key with a node in S_2 .*

Example 1. Consider the network given by the following diagram, in which dots represent nodes and lines connect nodes that share a key:



The capture of the white node and compromise of the corresponding keys isolates the set of black nodes on the left from the set of black nodes on the right, as none of the leftmost three nodes shares a key with any of the rightmost three nodes.

Example 2. Consider a one-dimensional network in which the nodes are labeled $0, 1, 2, \dots$ in turn. If each node with label i shares a key with the nodes labelled $i + 2$ and $i - 2$ then the set of nodes with odd labels is isolated from the set of nodes with even labels.



In a network where each node stores k keys, then an adversary can always isolate a single node Ψ from the rest of the network by capturing up to k other nodes that between them possess all the keys stored by Ψ . However, it could be argued that the adversary could achieve the same effect more easily by simply capturing Ψ directly. Hence the standard graph-theoretic notion of vertex connectivity does not quite serve to measure how well a KPS for a one-dimensional network stands up to node compromise. The exclusion of a small number of nodes from the rest of the network is not a major source of concern, however the partitioning of a network into two halves that cannot communicate with each other is a serious problem. In order to give a more practical measure of the extent to which an adversary can damage a network in which the KPS is deployed, we define the *s-fallibility* of a KPS to be the smallest number of nodes an adversary has to capture in order to cause a catastrophic failure in the connectivity of the network.


Consider a KPS that can be applied to a one-dimensional network of n of nodes (where n can be made arbitrarily large). If the KPS is *s-fallible*, then an adversary who captures at most $s - 1$ nodes can only isolate a small (i.e. constant with respect to n) number of nodes from the rest of the network, whereas an adversary who captures s nodes can partition the network into two (or more) large sets of nodes that are isolated from each other. We formalise the definition of *s-fallibility* as follows:

Definition 2. A KPS for a one-dimensional network consisting of a set \mathcal{N} of n nodes, where n is arbitrary, is *s-fallible* if the following two conditions hold:

1. After the capture of any $s - 1$ nodes, there exists a set $\mathcal{E} \subset \mathcal{N}$ of size at most $O(1)$ such that $\mathcal{N} \setminus \mathcal{E}$ is connected.

Theorem 2. *Suppose a KPS that yields a connected network assigns keys to nodes such that the largest distance between two nodes that share a key is b . Then it is s -fallible for some $1 \leq s \leq b$.*

Proof. As above, consider an adversary that captures b adjacent nodes from the centre of the network. By construction, no node in the “left-hand half” of the network shares a key with any node in the “right-hand half”, as shown below.



Thus we see that a KPS can be at most r -fallible, and that for this to be the case the KPS must assign shared keys to pairs of nodes at distance r . There do exist r -fallible schemes: one example is the local pairwise KPS¹. This scheme requires each node to store $2r$ keys. However, we will show in Sect. 6 that r -fallibility can in fact be achieved using storage that is independent of both r and n .


6 An Ultra-Lightweight KPS Providing Optimal s -Fallibility

By an extension of the proof of Theorem 2, we can see that for a KPS to achieve r -fallibility, it is necessary for each pair of nodes at distance r to share a key (except perhaps for a constant (with respect to n) number of pairs at either end of the network). However, this alone does not provide r -fallibility, since this distribution results in a network consisting of r isolated sets of (n/r) nodes (as was the case in Example 2 for $r = 2$). In order to achieve r -fallibility, it is necessary to introduce more keys into the network. The following construction leads to the surprising result that it is possible to obtain r -fallibility when each node stores only four keys.

Construction 1 *Assign keys to the nodes of a one-dimensional network such that each node shares unique pairwise keys with each of the nodes at distance r and 1.*

¹ The addition of extra pairwise shared keys to a KPS will evidently not decrease its s -fallibility, and may even increase it. As the local pairwise scheme ensures that every pair of nodes within communication range shares a key, you would thus expect its s -fallibility to be as high, or higher, than that of any other scheme. The fact that it is indeed r -fallible is a direct consequence of our proof of Theorem 3.

Example 3. If Construction 1 is applied to a one-dimensional network with twenty nodes for which $r = 3$, then the nodes share keys as illustrated by the following diagram.



Theorem 3. *The KPS of Construction 1 is r -fallible.*

Proof. Let the nodes of the network be labeled sequentially by the integers $0, 1, \dots, n-1$. Suppose an adversary captures $r-1$ nodes from the network and extracts the keys they contain.

Let Ψ_1 and Ψ_2 be two nodes that occur at distance at least $r+1$ from any captured node; without loss of generality we suppose the label of Ψ_1 is greater than that of Ψ_2 . None of their keys are known to the adversary. As the number of captured nodes is $r-1$, there must be some integer x with $0 \leq x \leq r-1$ such that no node with a label equivalent to $x \pmod{r}$ has been captured.

Then a secure path from Ψ_1 to Ψ_2 can be found as follows:


1. Take hops of length one from Ψ_1 towards Ψ_2 until a node whose label is equivalent to $x \pmod{r}$ is reached. This requires at most $r-1$ hops, hence each of these hops is secure by the assumption that no captured node lies within distance $r+1$ of Ψ_1 .
2. Take hops of length r towards Ψ_2 until a node at distance less than r from Ψ_2 is reached. As the keys required for these hops all belong to nodes whose label is equivalent to $x \pmod{r}$, none of them is known to the adversary.
3. Finally, complete the path by hops of length one until Ψ_2 is reached. The fact that no captured node lies within distance $r+1$ of Ψ_2 implies these are also secure.

The number of nodes within distance $r+1$ of a captured node is at most $2(r+1)(r-1)$; this does not depend on n . Hence we have shown that no matter which $r-1$ nodes are compromised, it is possible to exclude a constant (with respect to n) number of nodes such that the remaining nodes form a connected network, and thus the s -fallibility of this scheme is at least r . Together with Theorem 1, this shows that this KPS is in fact r -fallible.

6.1 Lightweight Key Predistribution for Ring Topologies

The scheme of Construction 1 can also be applied directly to one-dimensional networks where the sensors are arranged in a ring. For instance, for a

network of twenty nodes with $r = 3$, the pattern of key sharing can be depicted as follows:



The behaviour of a ring network with respect to s -fallibility is slightly different to a linear network. An adversary has to “cut” the network twice in order to disconnect it; the first cut essentially turns it from a ring network to a linear network. It is debatable as to what is the most useful definition of fallibility in this scenario, since for a large ring making even a single “cut” greatly increases the average number of hops required for two nodes to communicate. In any case, the proof of Theorem 3 can be adapted for the case of a ring network, hence we can argue that the security of the KPS of Construction 1 is at least as strong when it is applied to a ring network as when it is used in a linear network. Hence this scheme is also suitable for networks designed for perimeter surveillance and other such application.

7 Conclusion

We have seen that the topology of a one-dimensional network affects its potential connectivity, due to the restricted number of neighbours possessed by each node. This leads to the unexpected result that the KPS given in Construction 1 performs as well as the local pairwise KPS in terms of both connectivity and resilience, despite requiring each node to store only four keys. To summarise, the advantages of this scheme include the following:

- optimal resilience;
- optimal s -fallibility;
- very low storage.

In short, this scheme is thus ideal for use in WSNs for monitoring linear infrastructure, no matter how constrained the memories of the nodes.

Acknowledgements.

We would like to thank the referees for their helpful suggestions for improving the clarity of this paper.

References

1. S. A. Çamtepe and B. Yener. Key distribution mechanisms for wireless sensor networks: a survey. Rensselaer Polytechnic Institute, Computer Science Department, Technical Report TR-05-07, March 2005.
2. L. Eschenauer and V. D. Gligor. A key-management scheme for distributed sensor networks. In *CCS '02: Proceedings of the 9th ACM conference on Computer and communications security*, pages 41–47, New York, NY, USA, 2002. ACM Press.
3. I. Jawhar, N. Mohamed, and K. Shuaib. A framework for pipeline infrastructure monitoring using wireless sensor networks. In *Wireless Telecommunications Symposium*, pages 1–7, 2007.
4. I. Jawhar, N. Mohamed, K. Shuaib, and N. Kesserwan. Monitoring linear infrastructures using wireless sensor networks. In *The 2008 IFIP Conference on Wireless Sensor and Actor Networks (WSAN08)*, volume 264, pages 185–196. Springer, 2008.
5. J. Lee and D. R. Stinson. On the construction of practical key predistribution schemes for distributed sensor networks using combinatorial designs. *ACM Transactions on Information and System Security*, 11(2):1–35, 2008.
6. O. Lévêque and E. Preissmann. Scaling Laws for One-Dimensional Ad Hoc Wireless Networks. *IEEE Transactions on Information Theory*, 51(11):3987–3991, 2005.
7. B. Liu, P. Thiran, and D. Towsley. Capacity of a wireless ad hoc network with infrastructure. In *MobiHoc '07: Proceedings of the 8th ACM international symposium on Mobile ad hoc networking and computing*, pages 239–246, New York, NY, USA, 2007. ACM.
8. K. M. Martin and M. B. Paterson. An application-oriented framework for wireless sensor network key establishment. *Electronic Notes in Theoretical Computer Science*, 192(2):31–41, 2008.
9. D. Miorandi and E. Altman. Connectivity in one-dimensional ad hoc networks: a queueing theoretical approach. *Wireless Networks*, 12(5):573–587, 2006.
10. N. Mohamed and I. Jawhar. A fault tolerant wired/wireless sensor network architecture for monitoring pipeline infrastructures. *Sensor Technologies and Applications, International Conference on*, 0:179–184, 2008.
11. N. Mohamed, I. Jawhar, and K. Shuaib. Reliability challenges and enhancement approaches for pipeline sensor and actor networks. In H. R. Arabnia and V. A. Clincy, editors, *ICWN*, pages 46–51. CSREA Press, 2008.
12. K. Römer and F. Mattern. The design space of wireless sensor networks. *IEEE Wireless Communications Magazine*, 11(6):54–61, 2004.
13. I. Stoianov. Pipenet: A wireless sensor network for pipeline monitoring. In *IPSN07*, pages 264–273. ACM, 2007.
14. G. Wittenburg, K. Terfloth, F. L. Villafuerte, T. Naumowicz, H. Ritter, and J. H. Schiller. Fence monitoring - experimental evaluation of a use case for wireless sensor networks. In K. Langendoen and T. Voigt, editors, *EWSN*, volume 4373 of *Lecture Notes in Computer Science*, pages 163–178. Springer, 2007.

15. Y. Xiao, V. K. Rayi, B. Sun, X. Du, F. Hu, and M. Galloway. A survey of key management schemes in wireless sensor networks. *Computer Communications*, 30(11-12):2314–2341, 2007.
16. M. Zimmerling. An energy-efficient routing protocol for linear wireless sensor networks. In *Lecture Notes in Informatics*, pages 239–242, March 2008.
17. M. Zimmerling, W. Dargie, and J. M. Reason. Localized power-aware routing in linear wireless sensor networks. In *CASEMANS '08: Proceedings of the 2nd ACM international conference on Context-awareness for self-managing systems*, pages 24–33, New York, NY, USA, 2008. ACM.