Lightweight Key Establishment for Distributed Networking Environments

Keith Martin

Information Security Group, Royal Holloway, University of London

COSIC Seminar 2007

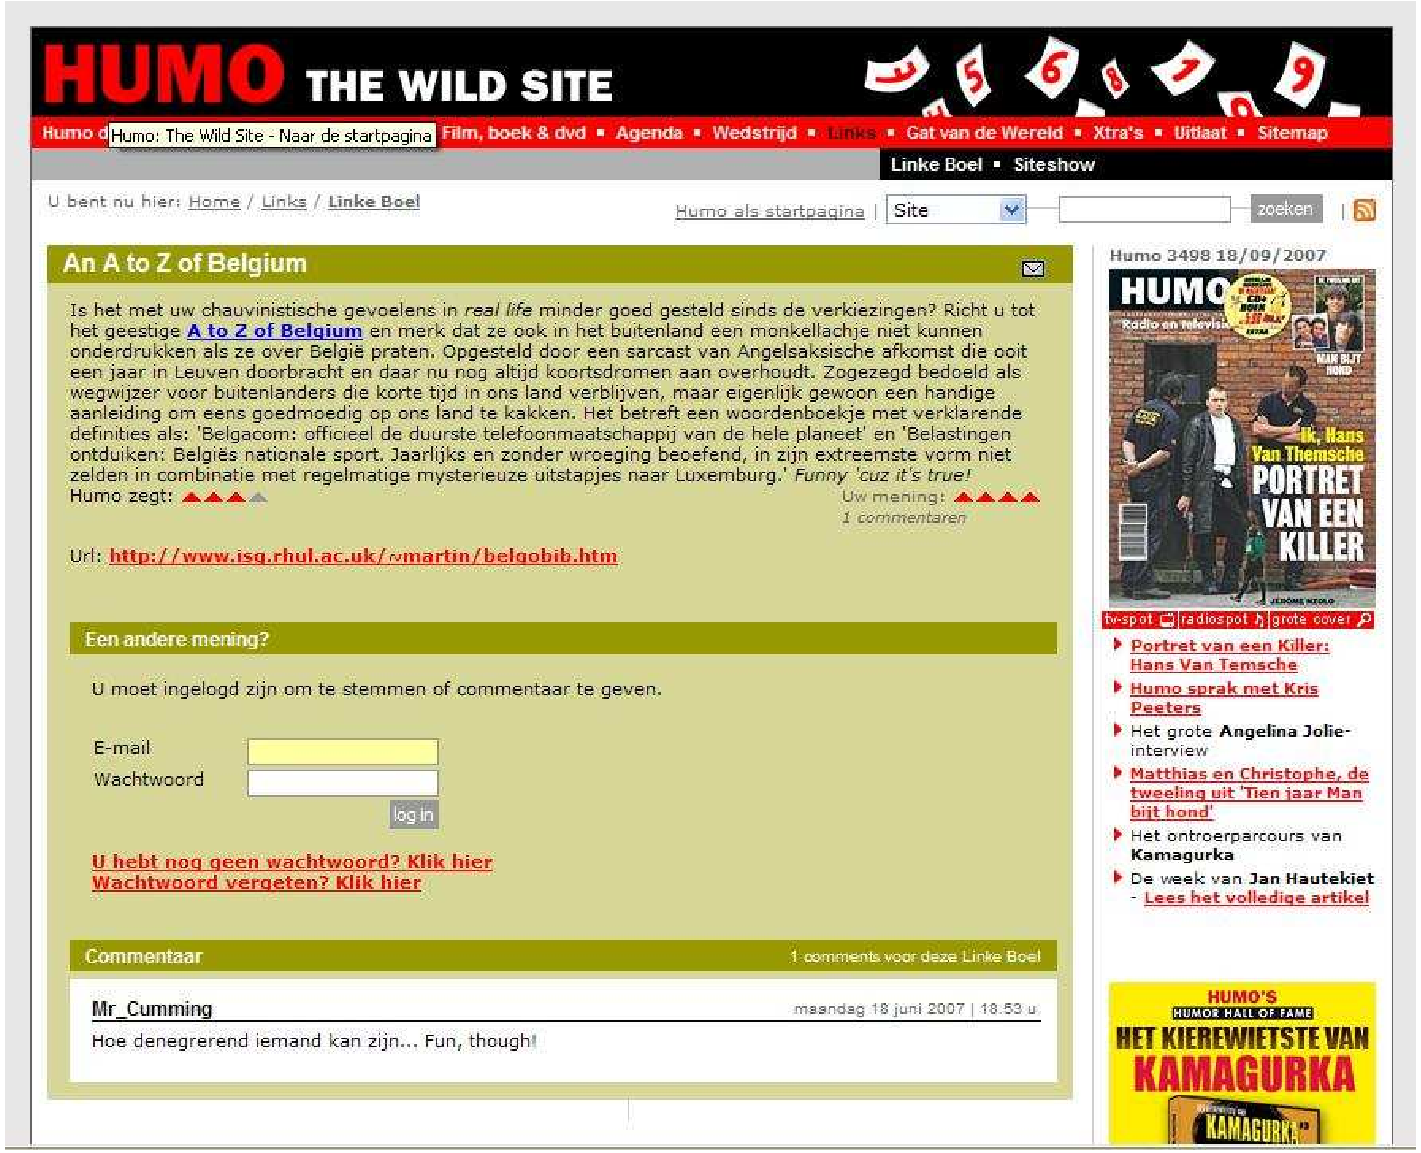
HUMO THE WILD SITE
Humo d Humo: The Wild Site - Naar de startpagina Film, boek & dvd ▪ Agenda ▪ Wedstrijd ▪ Links ▪ Gat van de Wereld ▪ Xtra's ▪ Uitlaat ▪ Sitemap
Linke Boel ▪ Siteshow
Humo als startpagina | Site
zoeken
tv-spot | radiospot | grote cover
log in
HUMO'S
HUMOR HALL OF FAME
HET KIEREWIETSTE VAN
KAMAGURKA

U bent nu hier: Home / Links / Linke Boel

An A to Z of Belgium

Is het met uw chauvinistische gevoelens in *real life* minder goed gesteld sinds de verkiezingen? Richt u tot het geestige **A to Z of Belgium** en merk dat ze ook in het buitenland een monkellachje niet kunnen onderdrukken als ze over België praten. Opgesteld door een sarcast van Angelsaksische afkomst die ooit een jaar in Leuven doorbracht en daar nu nog altijd koortsdromen aan overhoudt. Zogezegd bedoeld als wegwijzer voor buitenlanders die korte tijd in ons land verblijven, maar eigenlijk gewoon een handige aanleiding om eens goedmoedig op ons land te kakken. Het betreft een woordenboekje met verklarende definities als: 'Belgacom: officieel de duurste telefoonmaatschappij van de hele planeet' en 'Belastingen ontduiken: Belgiës nationale sport. Jaarlijks en zonder wroeging beoefend, in zijn extreemste vorm niet zelden in combinatie met regelmatige mysterieuze uitstapjes naar Luxemburg.' *Funny 'cuz it's true!*

Humo zegt: ▲▲▲△ Uw mening: ▲▲▲▲

1 commentaren

Url: **http://www.isg.rhul.ac.uk/~martin/belgobib.htm**

Een andere mening?

U moet ingelogd zijn om te stemmen of commentaar te geven.

E-mail

Wachtwoord

U hebt nog geen wachtwoord? Klik hier
Wachtwoord vergeten? Klik hier

Commentaar

1 comments voor deze Linke Boel

Mr_Cumming maandag 18 juni 2007 | 18.53 u

Hoe denegrerend iemand kan zijn... Fun, though!

Humo 3498 18/09/2007

- **Portret van een Killer: Hans Van Temsche**
- **Humo sprak met Kris Peeters**
- Het grote **Angelina Jolie**-interview
- **Matthias en Christophe, de tweeling uit 'Tien jaar Man bijt hond'**
- Het ontroerparcours van **Kamagurka**
- De week van **Jan Hautekiet** - **Lees het volledige artikel**

The plan

1. **Wireless sensor networks**
2. **A key establishment framework**
3. **Key establishment for grids**

Wireless sensor networks

Wireless sensor networks

- consists of small(ish) **sensors** forming an **ad-hoc network**
- sensors may lack an independent power supply and hence have limited:
  - storage
  - communication bandwidth
  - computational power
- sensors are vulnerable to compromise through capture
- sensors are vulnerable to failure (or may regularly go down by design)

A “great” debate... ?

Can you perform public-key cryptographic operations on a sensor?

Maybe you can, maybe you can't...

We will assume for the rest of this talk that it is preferred to design a fully symmetric solution, because *even if you can do public-key cryptography on a sensor*:

- it might be preferable to minimise its use to rarely conducted operations
- it isn't necessarily the right choice for implementation
- the same question is going to arise for the next generation of even smaller "sensors"
- key establishment using public-key cryptography is practically challenging but mathematically dull!

The devil's security advocate...

Wireless sensor networks are a solution awaiting a problem (*Gollman*)?

Applications

- Environmental monitoring
- Wildlife monitoring
- Disaster response
- Precision agriculture
- Surveillance
- Health care
- Process flow monitoring
- ...

The “classical” scenario

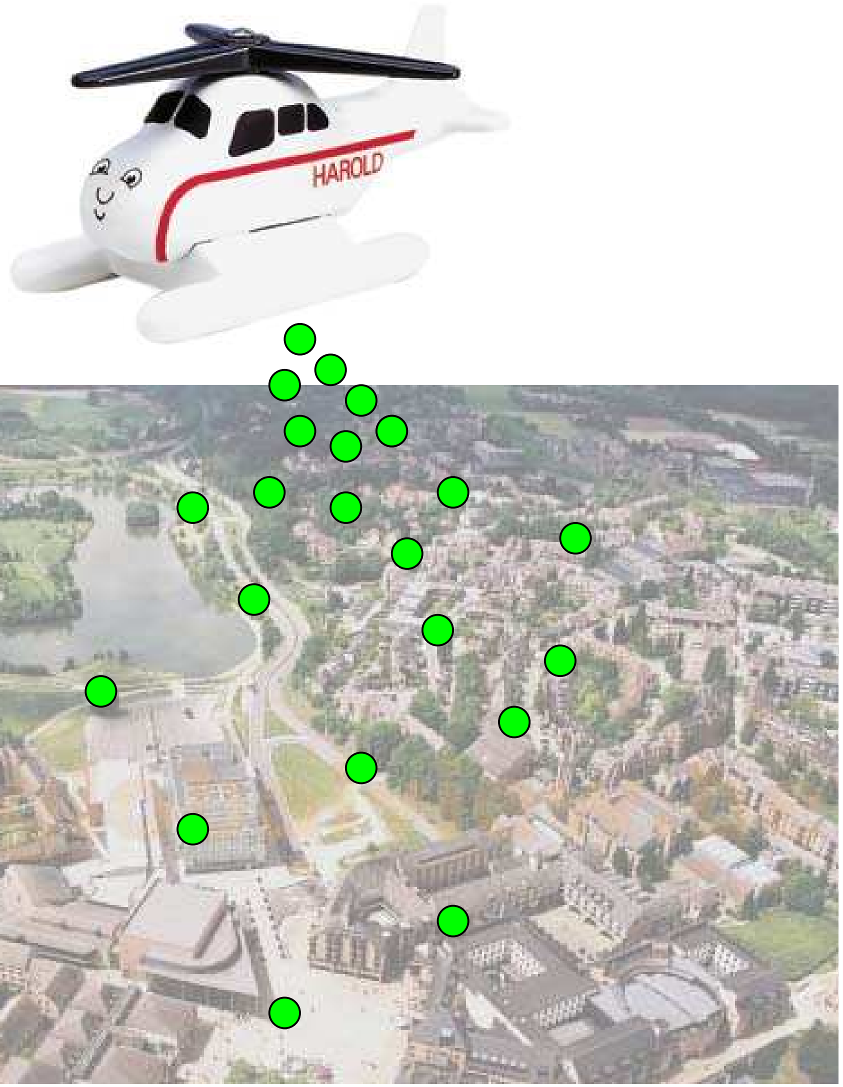

ZebraNet

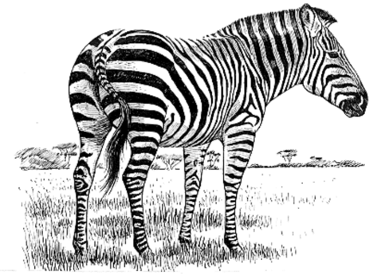

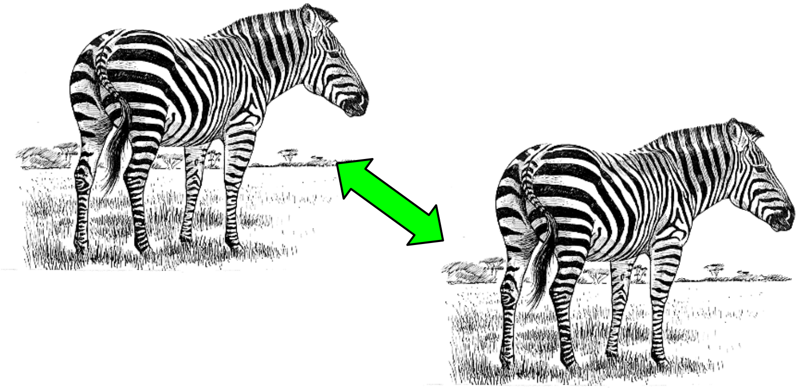

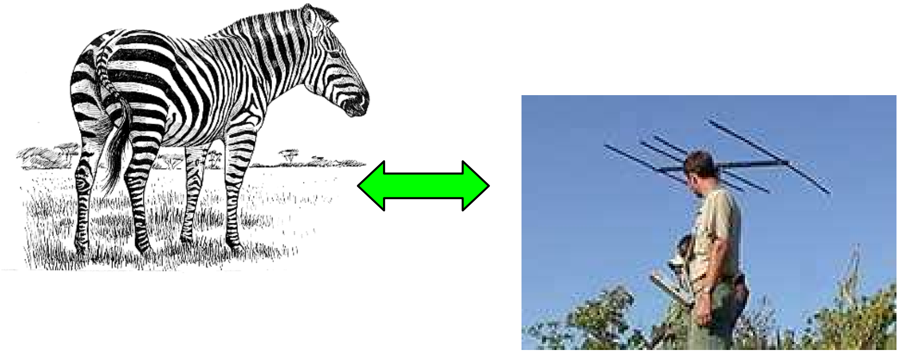

Smart vineyards

Three-stage key establishment

1. **Key predistribution**: Due to sensor compromise risk we cannot use a fixed key across all sensors. Thus each sensor assigned a set of keys chosen from a **key ring**. *How should this key ring be chosen?*

2. **Shared key discovery**: Two sensors can only communicate if they are in close physical proximity, and communicate **securely** if they also share a common key. *How should they determine whether they have any keys in common?*

3. **Path-key establishment**: If two nodes cannot communicate securely directly then they must establish a secure **multi-hop path** that utilises other sensors in the network. *How should a secure multi-hop path be determined?*

Communication structures

- **Ideal communication structure**: the groups of sensors for whom we (ideally) wish to establish common (group) keys

- **Network communication structure**: the groups of sensors who share predistributed keys

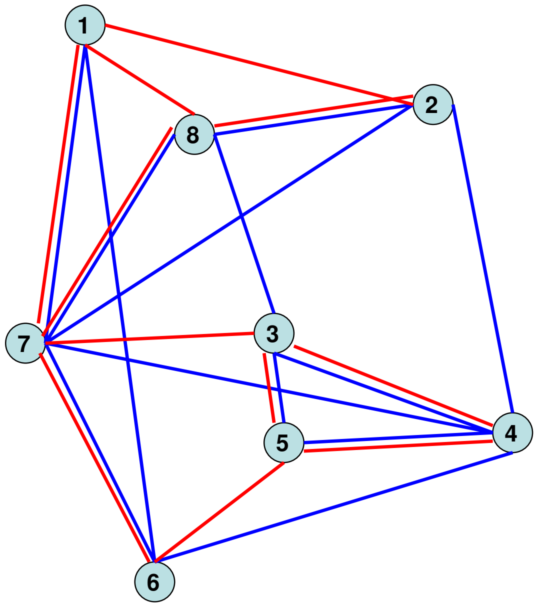

An application-oriented key establishment framework

Key establishment framework

1. Categories of sensor networks that significantly affect key establishment design.
2. Relevant variable parameters that determine instances within each of the above defined categories.
3. Performance indicators that can be used to assess specific key establishment schemes.

1A - Homogeneity

Sensor networks tend to fall into one of two classes:

1. **Homogeneous**: all sensors have the same capabilities.
2. **Hierarchical**: there is a natural hierarchy of sensors with respect to their capabilities (with fewer sensors at higher, more "powerful" levels).

The most common hierarchical networks are **two-level**, where there are two classes of sensor. Note that "powerful" could relate to issues such as amount of key storage, computational capability or degree of mobility.

1B - Deployment location control

Five classes of sensor network can be identified:

1. **Fixed, full control**: *the precise location of sensors is known before deployment*. Applications where sensors may then undertake strictly limited mobility (for example monitoring points on a glacier) can be placed within this class for the purposes of key management.

2. **Fixed, partial control**: *some information about the location of sensors is known before deployment*. This class includes applications where clusters of sensors are dropped from the air over fixed locations.

3. **Fixed, no control**: *the location of sensors cannot be predicted before deployment*. This class includes applications where sensors are randomly scattered over a monitoring area.

4. **Locally mobile**: *sensors are mobile within a controlled locality.* In this class, sensors can be assumed to be free to move to any location within a strictly defined local area, but cannot stray out of this area.

5. **Fully mobile**: *sensors are mobile.* In this class, sensors are free to move anywhere within the network environment.

1C - Ideal communication structure

Three important classes of ideal communication structure are:

1. ***t*-complete**: all subsets of sensors of size t. Most commonly **pairwise complete** (2-complete).
2. **locally *t*-complete**: all local subsets of sensor of size t, where *local* generally refers to sensors who are neighbours of one another in some sense. Most commonly **pairwise locally complete**, which arises in applications where the most commonly required communication flow is between a (mobile) external **sink** and any sensor.
3. **regionally *t*-complete**: all subsets of sensors of size t within a specified region. (Differs from locally t-complete as sensors belonging to same "region" are required to share key associations.)

2 - Variable network parameters

The following parameters can be set to define a specific scheme.

- **Storage**: *The storage capability of a sensor.* This is perhaps the most significant parameter in terms of its direct limiting effect on key establishment scheme design.
- **Energy**: *The energy available for a sensor to conduct computations and communications.* It is generally considered that the energy requirements for communication far outweigh those of computation.
- **Range**: *The communication range over which a sensor can contact other sensors.* This is also related to the energy capability since greater communication ranges tend to consume more power.

3 - Performance indicators

- **Connectivity**: Measures how closely the network communication structure matches the ideal communication structure.
- **Scalability**: Measures the feasibility of use with large network sizes. It essentially reflects the storage requirements relative to the number of nodes in the network.
- **Resilience**: Indicates the proportion of established keys that become compromised once the adversary has access to the secret data from a small proportion of the nodes.
- **Computation/Communication overheads**: Measure the precise costs of a particular solution.

Research snapshot

	2-c/ t-c	locally 2-c/ locally t-c	regionally 2-c/ regionally t-c	hierarchical- 2-level
fixed, full control				
fixed, partial control		1	9	
fixed, no control	23			1

Observations

- In many applications there is a degree of control over sensor location. Knowledge of the network topology and location of sensors is likely to be exploitable in the design of key establishment schemes that are more efficient than those defined for the default scenario.
- The majority of applications have no apparent need for a pairwise ideal communication structure. Since applications of this type only really need local communication between sensors in order to securely relay information, partial control over sensor location should lead to more efficient schemes.

Key establishment for grids

(locally 2-complete schemes for a network with fixed sensors and full location control)

Key establishment for grids

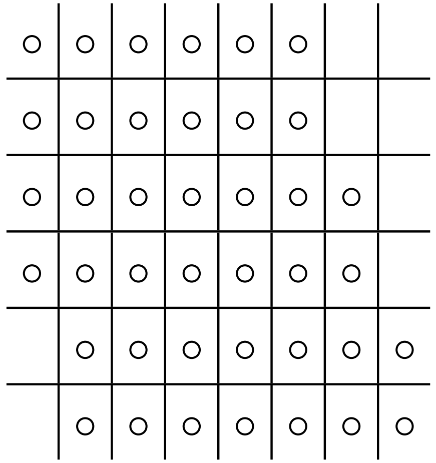

- monitoring vines in a vineyard or trees in a commercial plantation
- studying traffic or pollution levels on city streets
- measuring humidity and temperature on library shelves
- performing acoustic testing at each of the seats in a theatre

Lee spheres

A **Lee sphere of radius** r centred at a given square consists of the set of squares that lie at (Manhattan) distance at most r from that square.

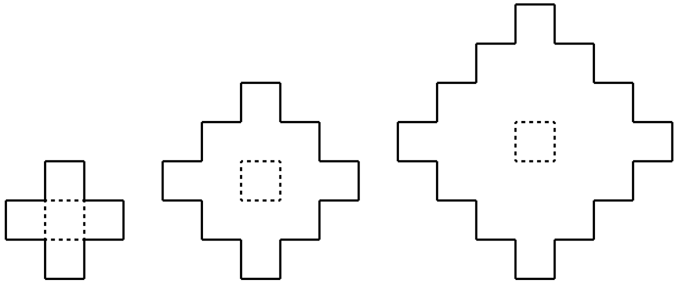

Lee spheres of radii 1, 2 and 3

We want to design a schemes in which each node shares a key with as many nodes as possible in the Lee sphere of radius r *surrounding it.*

Definitions

- **$[m, \alpha]$-KPS**: each sensor stores at most m keys and each key is shared by at most α sensors.

- **(λ, r)-coverage**: the expected proportion of sensors within the Lee sphere of radius r centred at some sensor Ψ that are within λ-hop distance of Ψ.

- **$(1, r)$-coverage**: the expected number of nodes within the Lee sphere that share keys with Ψ. Note that in an $[m, \alpha]$-KPS the $(1, r)$-coverage is less than or equal to
$$\frac{m(\alpha - 1)}{2r(r+1)}.$$

- **tight $(1, r)$-coverage**: a scheme meeting the above bound.

Tight schemes

A scheme achieves tight $(1, r)$-coverage when the following conditions are met:

1. Each node stores exactly m keys and each key is shared by exactly α nodes.
2. No pair of nodes shares two or more keys.
3. The Manhattan distance between any two nodes sharing a key is at most r.

Costas arrays

A **Costas array** of order n is an $n \times n$ matrix with the following properties:

- each position is either blank or contains a dot,
- each row and each column contains exactly one dot,
- all $\binom{n}{2}$ vectors connecting pairs of dots are all distinct as vectors (any two vectors are different in either magnitude or orientation).

$$\begin{bmatrix} & \bullet & \\ & & \bullet \\ \bullet & & \end{bmatrix}$$

Costas array scheme

Given an $n \times n$ Costas array, distribute keys to sensors as follows:

- place the Costas array in every possible position in the square grid,
- associate some key k with each positioning of the Costas array,
- store k in the nodes corresponding to the dots of the array.

Suppose sensor Ψ stores keys labelled a, b and c. The other nodes that share these keys are:

	·	·	b	
·	a	·	·	b
·	·	Ψ	·	·
a	·	·	c	·
	c	·	·	

Properties of Costas array scheme

1. Each sensor has n different keys.
2. Each key is assigned to n sensors.
3. Any two sensors have at most one key in common.
4. The (Manhattan) distance between two sensors which have a common key is at most $2(n-1)$.

This is an $[n, n]$*-KPS with tight* $(1, 2(n-1))$*-coverage.*

Distinct difference arrays

The Costas array construction relies on the property that the vectors connecting pairs of dots in a Costas array are pairwise distinct. *We do not, however, make use of the the requirement that each row and column have exactly one dot.*

A **distinct-difference configuration** $DD(m, r)$ consists of a set of m dots placed in a square grid such that:

- any two dots are Manhattan distance at most r apart,
- all $\binom{m}{2}$ differences between pairs of dots are distinct as vectors.

Distinct difference array scheme

A $DD(m, r)$ generates a KPS with the properties:

1. Each sensor has m different keys.
2. Each key is assigned to m sensors.
3. Any two sensors have at most one key in common.
4. The (Manhattan) distance between two sensors which have a common key is at most r.

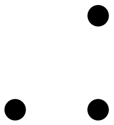

		$\cdot$		
	$\cdot$	a	b	
$\cdot$	a	Ψ	b	$\cdot$
	c	c	$\cdot$	
		$\cdot$		

Balancing storage, resilience and coverage

- Low number of keys per sensor is:
  - **good** for memory
  - **good** for resilience
  - **bad** for coverage.
- A compromise is to require **complete** $(2, r)$**-coverage**, and seek a $DD(m, r)$ with the minimum value of m.

$DD(m,r)$'s with minimal m

$r=1,2$ $m=3$	$r=3$ $m=4$	$r=4$ $m=5$	$r=5$ $m=5$	$r=6$ $m=6$	$r=7$ $m=6$
● ▪ ● ●	● ● ▪ ▪ ▪ ● ● ▪ ▪	● ▪ ▪ ● ▪ ▪ ▪ ● ● ▪ ▪ ▪ ● ▪ ▪	● ▪ ▪ ▪ ● ▪ ▪ ▪ ▪ ▪ ● ● ▪ ▪ ▪ ▪ ● ▪ ▪ ▪	● ▪ ▪ ▪ ● ▪ ▪ ▪ ▪ ▪ ● ▪ ▪ ▪ ▪ ● ▪ ▪ ● ▪ ▪ ▪ ▪ ▪ ● ▪ ▪ ▪	▪ ● ▪ ▪ ▪ ● ▪ ● ▪ ▪ ● ▪ ▪ ▪ ▪ ▪ ▪ ▪ ▪ ● ▪ ▪ ▪ ▪ ▪ ▪ ▪ ▪ ▪ ▪ ● ▪ ▪ ▪ ▪

Scheme from $DD(5,5)$

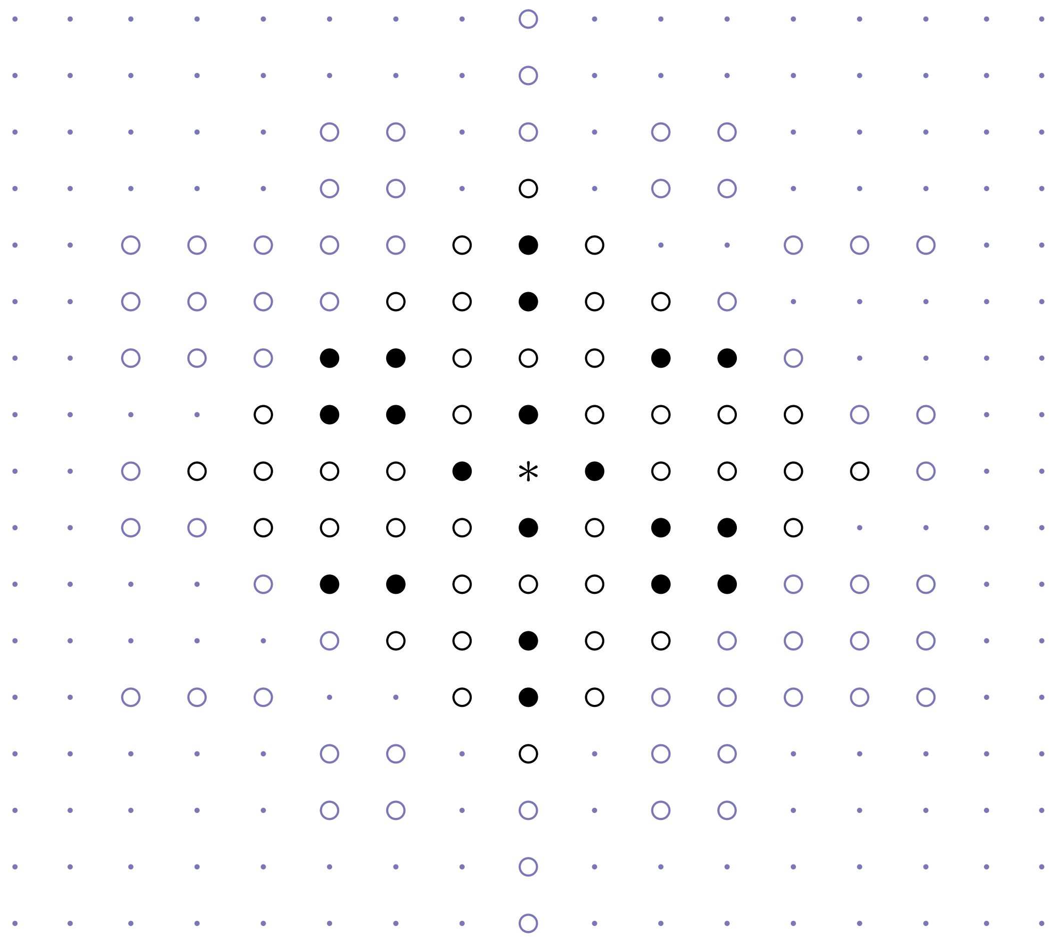

Concluding remarks

- There is plenty scope for investigating lightweight key establishment in sparsely populated regions of the framework grid
- Worth investigating other fixed network topologies
- Distinct difference arrays are interesting objects in their own right and little known about them

For more details...

- K.M. Martin, M.B. Paterson, *An application-oriented framework for wireless sensor network key establishment*, Proceedings of WCAN'07, to appear in ENTCS.
- S. Blackburn, T. Etzion, K.M. Martin, M.B. Paterson, *Efficient Key Predistribution for Grid-Based Wireless Sensor Networks*, Preprint.