Dynamic access policies for unconditionally secure secret sharing schemes

Dr Keith Martin

Information Security Group

Royal Holloway, University of London, U.K.

keith.martin@rhul.ac.uk

Ueli's orders

- Choose a topic related to information theoretic security
- Provide an overview
- Identify some open problems
- Give a "nice talk"

A little bit of history…

- First manifestation of this problem posed by by Agnes Chan in 1991 (and then at Crypto 92)
- Slow trickle of research contributions over the last 14 years:
  - some directly looking at part of the problem
  - some indirectly looking at part of the problem
  - almost all considering a slightly different version of the problem
  - not all looking at the information theoretically secure model

Plan for the “nice talk”

- What’s the problem?
- The different environments in which the problem can be tackled
- Some results and open research questions

A secret sharing scheme is…

- … a cryptographic primitive protecting a **secret**…
- … by distributing **shares** of the secret amongst a set of **shareholders**…
- … in such a way that only certain predetermined sets of shareholders (specified by the **access structure** or **access policy**) can recover the secret from their collective shares.

If the access structure consists of all sets of at least k shareholders (out of a total on n shareholders) then we call this a (k,n)-threshold scheme.

Information theoretically speaking…

A secret sharing scheme for protecting secret S using an access policy Γ defined on shareholder set P can be expressed as a set of probability distributions such that for any subset A of P,

$$H(S \mid A) = \begin{cases} 0 & \text{if } A \in \Gamma \\ H(S) & \text{if } A \notin \Gamma \end{cases}$$

Schemes of this type are often referred to as being **perfect**. They have the property that for every shareholder X we have

$$H(X) \geq H(S)$$

Sample applications

- Protection of master keys
- Protection of private root certificate keys
- Key recovery systems
- Shared generation of cryptographic secrets
- Key management in distributed environments
- Multiparty computation

Two operational phases

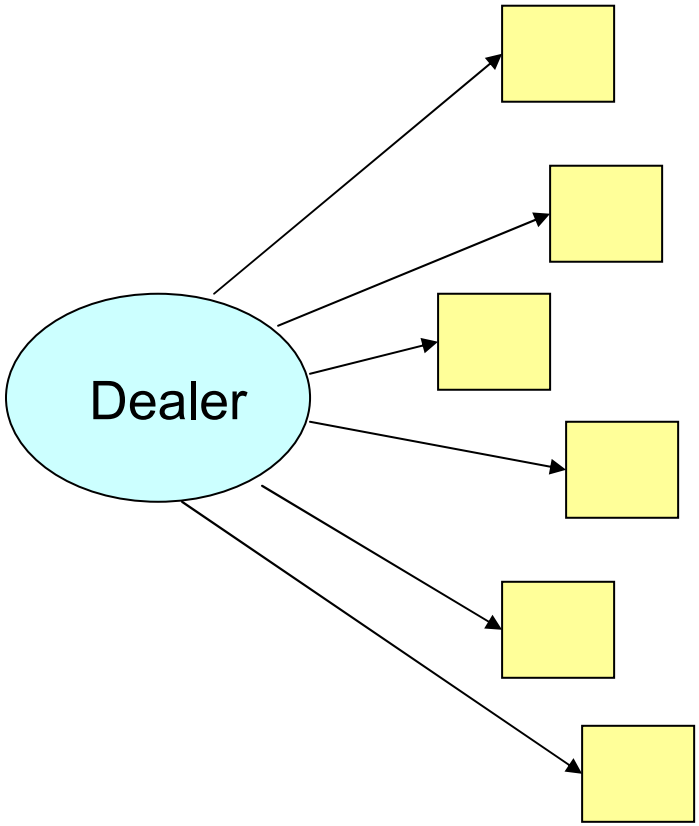

Initialisation phase

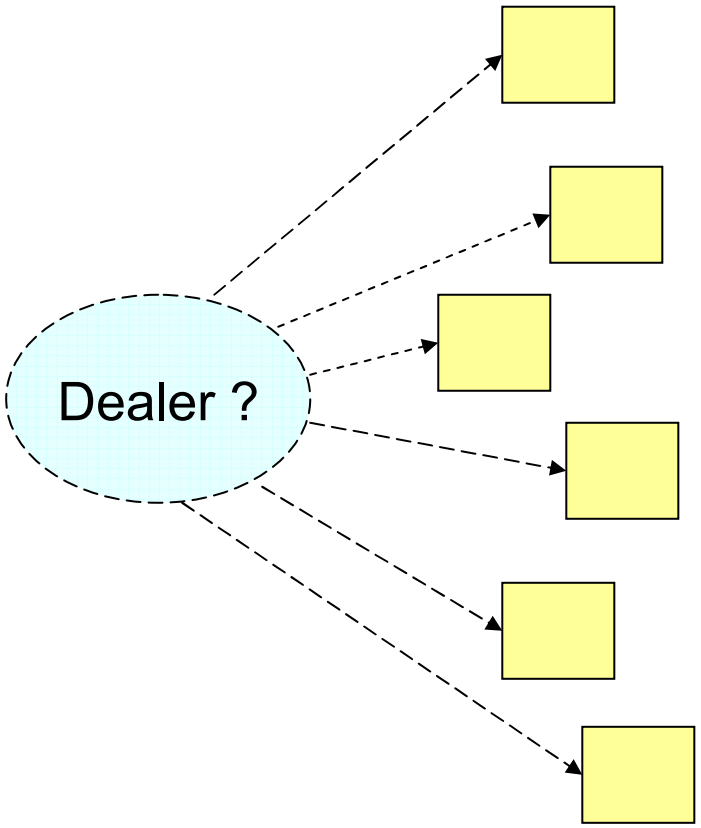

Running phase

Dynamic access policies

- During the running phase we might want to…
  - Remove a (corrupted) shareholder (**disenrollment**)
  - Add a new shareholder (**enrollment**)
  - Strengthen a component of the access policy
  - Weaken a component of the access policy

- …without repeating the initialisation phase because:
  - Initialisation is expensive
  - Communication conditions during the running phase may have changed. For example:
    - The dealer may no longer be present
    - Secure links may no longer be possible between dealer and shareholders

Dynamic access policies

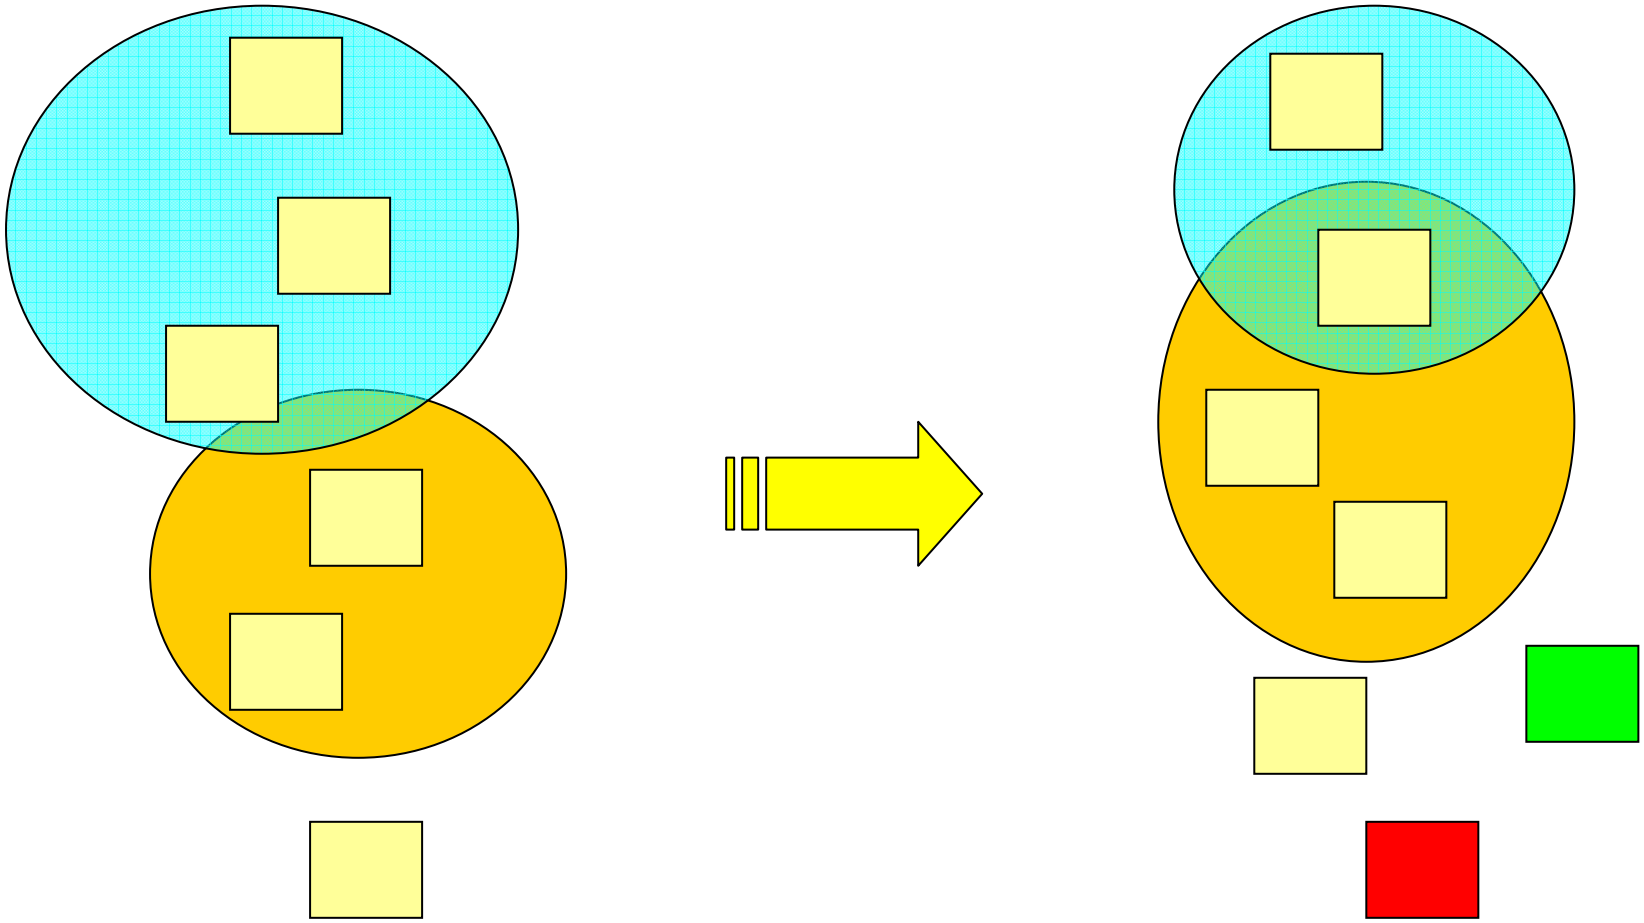

Before **After**

Management models

Model	Initialisation Phase			Running Phase		
	D	D - P	P - P	D	D - P	P - P
1	X	S	-	X	S	S
2	X	S	-	X	S	B
3	X	S	-	X	S	-
4	X	S	-	X	B	S
5	X	S	-	X	B	B
6	X	S	-	X	B	-
7	X	S	-			S
8	X	S	-			B
9	X	S	-			-
10			S			S
11			S			B
12			S			-

Type classification

Type	Knowledge of dynamic access policy changes at initialisation phase
A	None
B	Changes known to be necessary but not nature of changes
C	Some information known about the nature of changes that will be necessary

Type A schemes the most flexible

Type B schemes typically involve a bound on the number of permitted changes

Type C schemes usually dedicated to particular classes of change (such as disenrollments)

Robustness issues

- Most schemes:
  - are robust against coalitions of unauthorised shareholders attempting to reconstruct secrets that they are not permitted to.
  - assume that shareholders leaving the scheme are no longer trusted.
  - assume that shareholders follow protocols honestly and correctly
- There are some variations on this amongst the different security models

Comment on applicability

In information theoretically secure models if there exists a set of shareholders $A \in \Gamma$ such that after an access policy change $A \notin \Gamma^*$ then the secret must also change.

- A (k,n)-threshold scheme can be updated to a (k-1,n-1)-threshold scheme in Model 6.
- Method: broadcast the share of the shareholder being disenrolled.
- Secret does **not** change.

- A (k,n)-threshold scheme can be updated to a (k,n-1)-threshold scheme in several models.
- Method: any.
- Secret **must** change.

Possible measures of efficiency

- The amount of secret information that shareholders need to store
- The amount of secret information that shareholders need to communicate to facilitate an access policy update
- The amount of information that needs to be broadcast on public channels to facilitate an access policy update.
- Computational requirements of all participating entities.

Case 1: Fully functional dealer

Model	Initialisation Phase			Running Phase		
	D	D - P	P - P	D	D - P	P - P
1	X	S	-	X	S	S
2	X	S	-	X	S	B
3	X	S	-	X	S	-

Received almost no attention because the running phase does not differ significantly from the initialisation phase.

Assumption has always been that the dealer can simply reinitialise the scheme when an access policy update is required.

Not quite the end of the story…

Open Problem 1

Given access policies Γ and Γ^*, what is the most efficient Type A technique for updating a secret sharing scheme for Γ to a secret sharing scheme for Γ^* in any of Models 1, 2 or 3.

Case 2: Broadcast capable dealer

Model	Initialisation Phase			Running Phase		
	D	D - P	P - P	D	D - P	P - P
4	X	S	-	X	B	S
5	X	S	-	X	B	B
6	X	S	-	X	B	-

Most appropriate for applications where shareholders are initialised via a local trusted centre but then during the running phase are located within a distributed network with limited infrastructure (for example Internet applications or limited-area communications environments such as video conferencing, broadcast services etc).

Two basic techniques (Model 6)

Advance share technique

- During the initialisation phase, equip each shareholder with one share for every access policy that they might need to change to in the future.
- Access policy change requires a simple instruction
- Type C scheme
- Potentially enormous shares

Advance key technique

- During the initialisation phase, equip each shareholder with a share of the original scheme and L unique shareholder round keys (where L is the maximum anticipated access policy changes)
- Access policy change requires new shares to be broadcast by the dealer, each protected by the appropriate shareholder round key
- Type B scheme
- Potentially large broadcast information

Schemes with L-fold disenrollment

Model 6 Type B

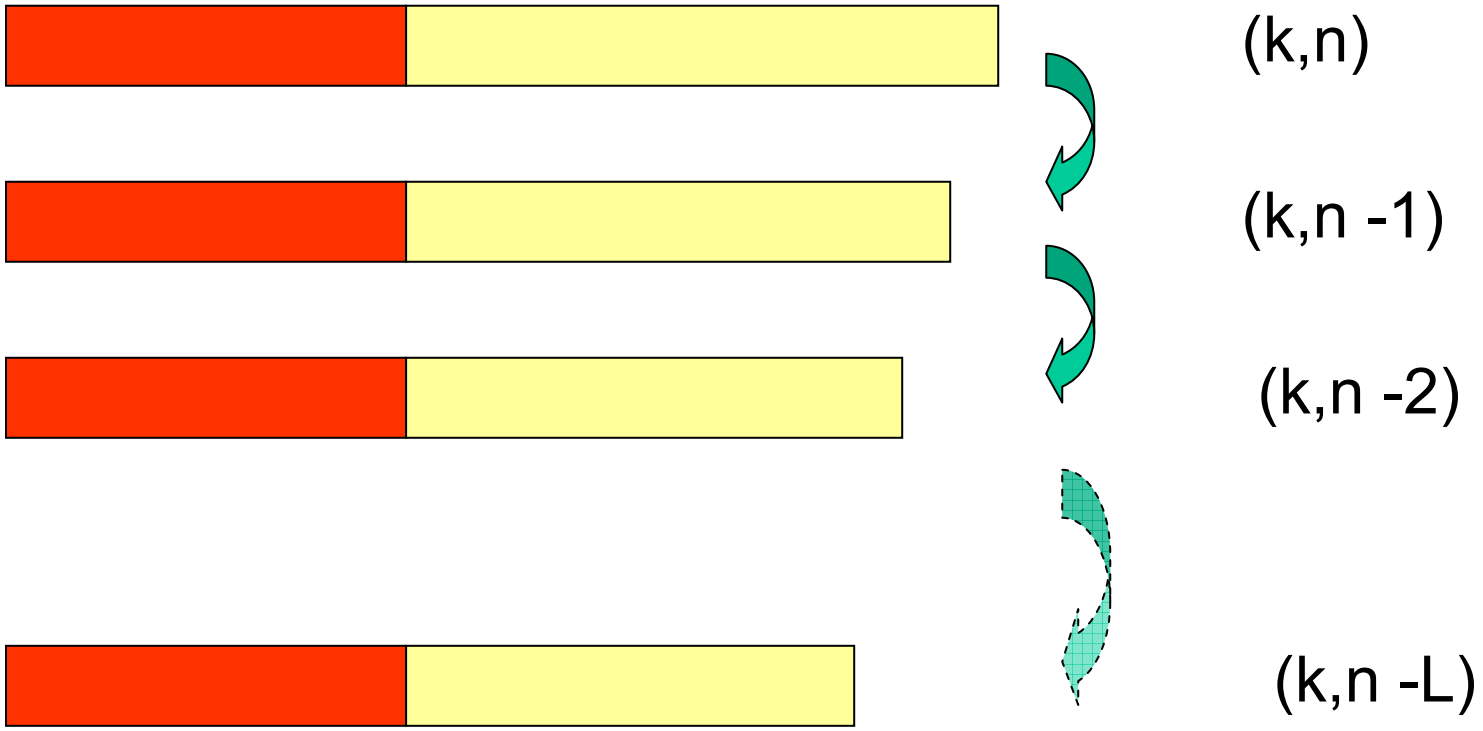

BBCM threshold scheme

Blakley et al (Crypto 92) proved that a (k,n)-threshold scheme with L-fold disenrollment satisfies

$$H(X) \geq (L+1)H(S)$$

- During the initialisation phase, equip each shareholder with one share for each of L independent threshold schemes with parameters (k,n), (k+1,n), … (k+L,n).
- Access policy change number r requires:
  - **An instruction to move onto the share set (r+1), which corresponds to the (k+r,n) threshold scheme**
  - **The dealer broadcasts the r shares corresponding to the disenrolled shareholders**
  - **Result is a (k,n-r) threshold scheme**

Other disenrollment schemes

- BBCM is a variant of a more general L-fold disenrollment scheme for general access policies (Martin 91).
- BBCM conjectured a bound on the size of broadcast information in any (k,n)-threshold with L-fold disenrollment.
- This conjecture is false and the correct bound and scheme with minimal share and broadcast size was shown in Barwick et al (ACISP 02).
- Schemes with more powerful notion of collusion security and a degree of forward secrecy were studied by Li and Poovendran (SAC 03).

Arbitrary updates to access policies

- **Dynamic secret sharing** was proposed by Blundo et al (Crypto 93):
  - Provided a formal information theoretically secure model.
  - Lower bounds for share size of such dynamic schemes were determined (essentially generalising BBCM bound to the more complex general case).
- **Threshold changeable schemes** were studied by Barwick et al (IEEE Trans Inf Theory 2005):
  - Bounds were established on the size of broadcast information needed to make one update to a Type C scheme with minimal share size for every possible range of update parameters.
  - Optimal schemes were constructed that met these bounds.

Threshold changeable schemes

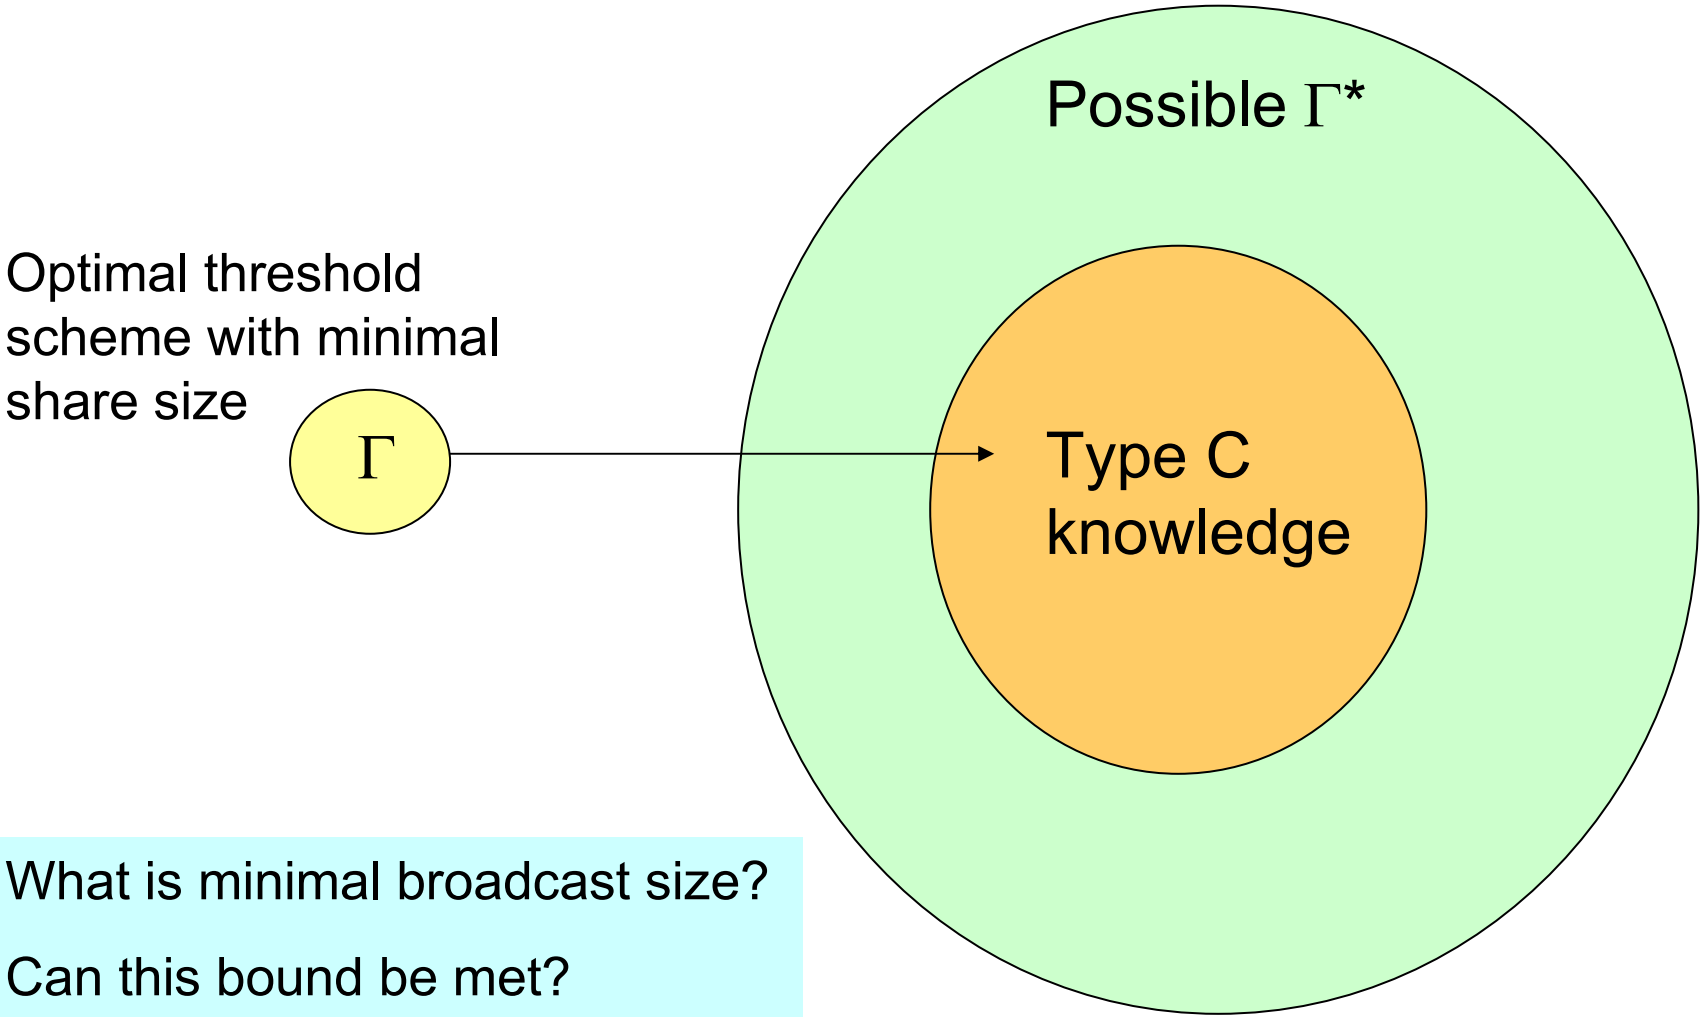

What is minimal broadcast size?

Can this bound be met?

Example bounds w.r.t. Type C knowledge

Case	Area defined by			Minimum share size	Broadcast lower bound
1	$k^*=k$	$n^*=n-1$		2H(S)	H(S)
2	$k^*=k$	$n^*<n$		2H(S)	$(n^*-k^*+1)H(S)$
3	$k^*\geq k$	$n^*<n$		2H(S)	$(n^*-k^*+1)H(S)$
4	$k^*\leq k$	$n^*\leq n$	$k-k^*<n-n^*$	2H(S)	$(n^*-k^*+1)H(S)$
5	$k^*\leq k$	$n^*\leq n$	$k-k^*\geq n-n^*$	H(S)	$\min(k-k^*,n^*-k^*+1)H(S)$
6	$k^*<k$	$n^*=n$		H(S)	$\min(k-k^*,n^*-k^*+1)H(S)$

Model 6 Open Problems

Open Problem 2

Can bounds on broadcast size of Barwick et al for making one policy update be easily extended to multiple policy updates?

Open Problem 3

There is a tradeoff between share size and broadcast size. If the share size is not minimal (but still “small”) can we find schemes with “small” broadcast sizes?

Model 6 Open Problems

Open Problem 4

Can proposed schemes in Model 6 be made robust against insiders failing to follow protocols correctly (without significant loss of efficiency)?

Open Problem 5

Is it possible to design schemes for updating access policies in Model 6 that are Type B?

(One such scheme was proposed by Tamura et al in 1999, but was later found to be flawed.)

Open Problem 5*

What about general access policies?

Case 3: Inactive dealer

Model	Initialisation Phase			Running Phase		
	D	D - P	P - P	D	D - P	P - P
7	X	S	-			S
8	X	S	-			B
9	X	S	-			-

Most appropriate for applications where shareholders are initialised via a local trusted centre but then during the running phase are remotely located in an environment lacking infrastructure (for example some ad hoc networking scenarios).

Redistributions: Model 7

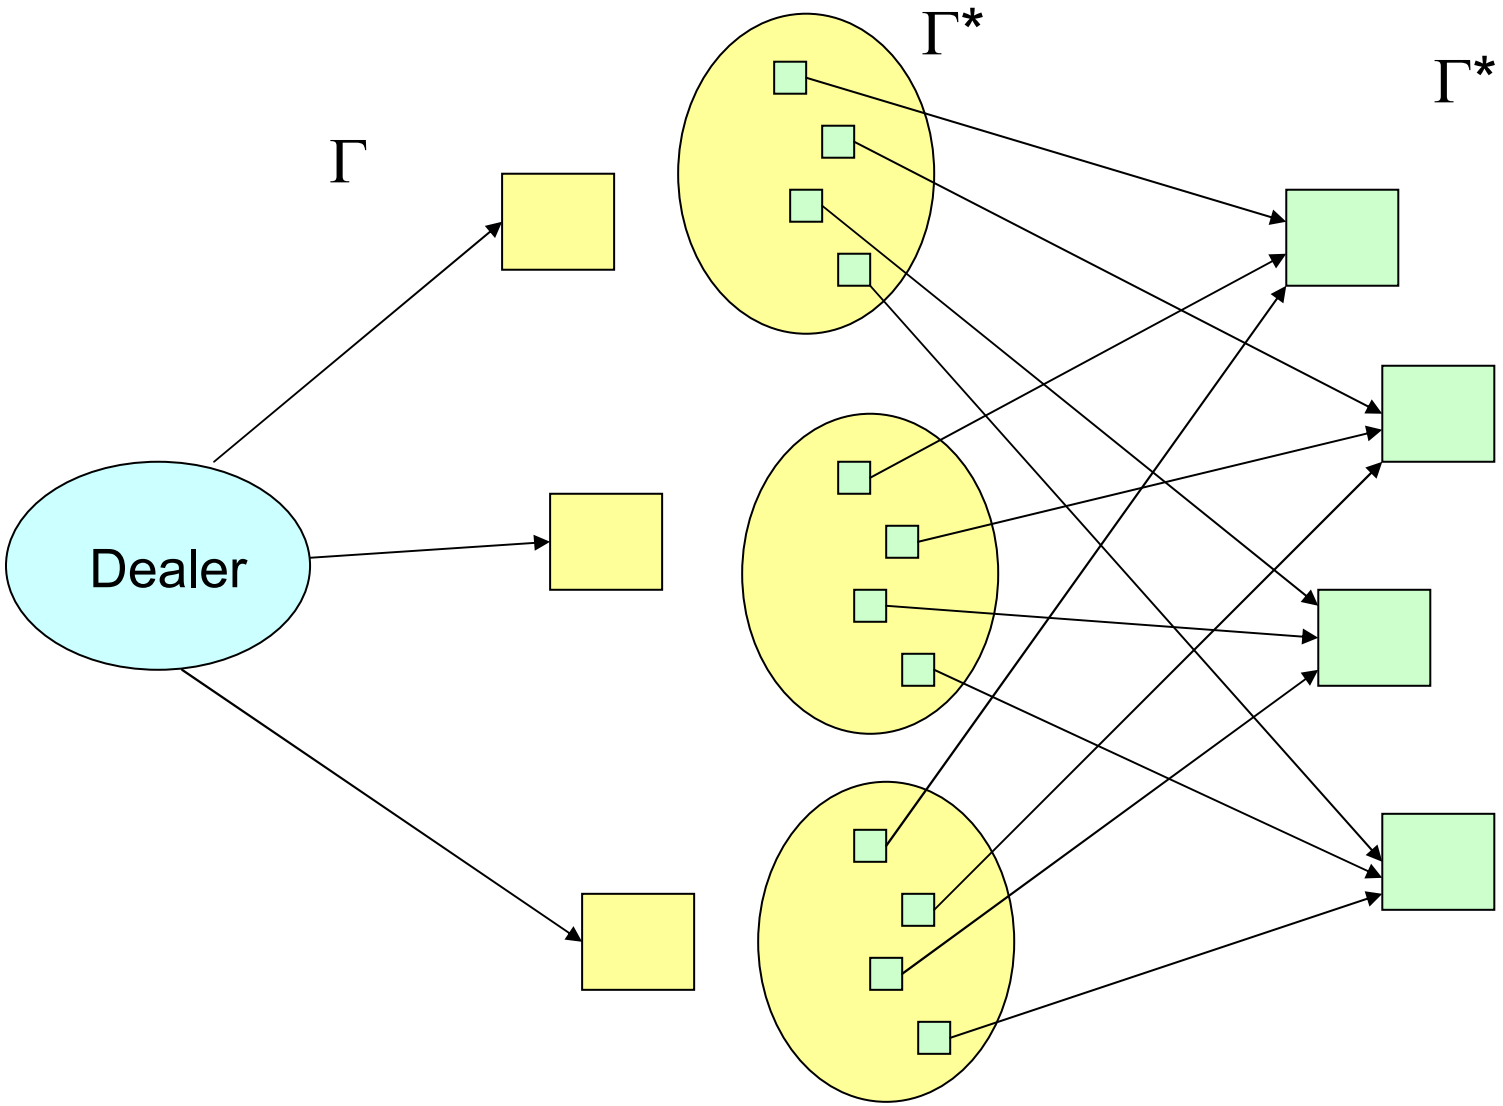

Linear redistributions: Model 7

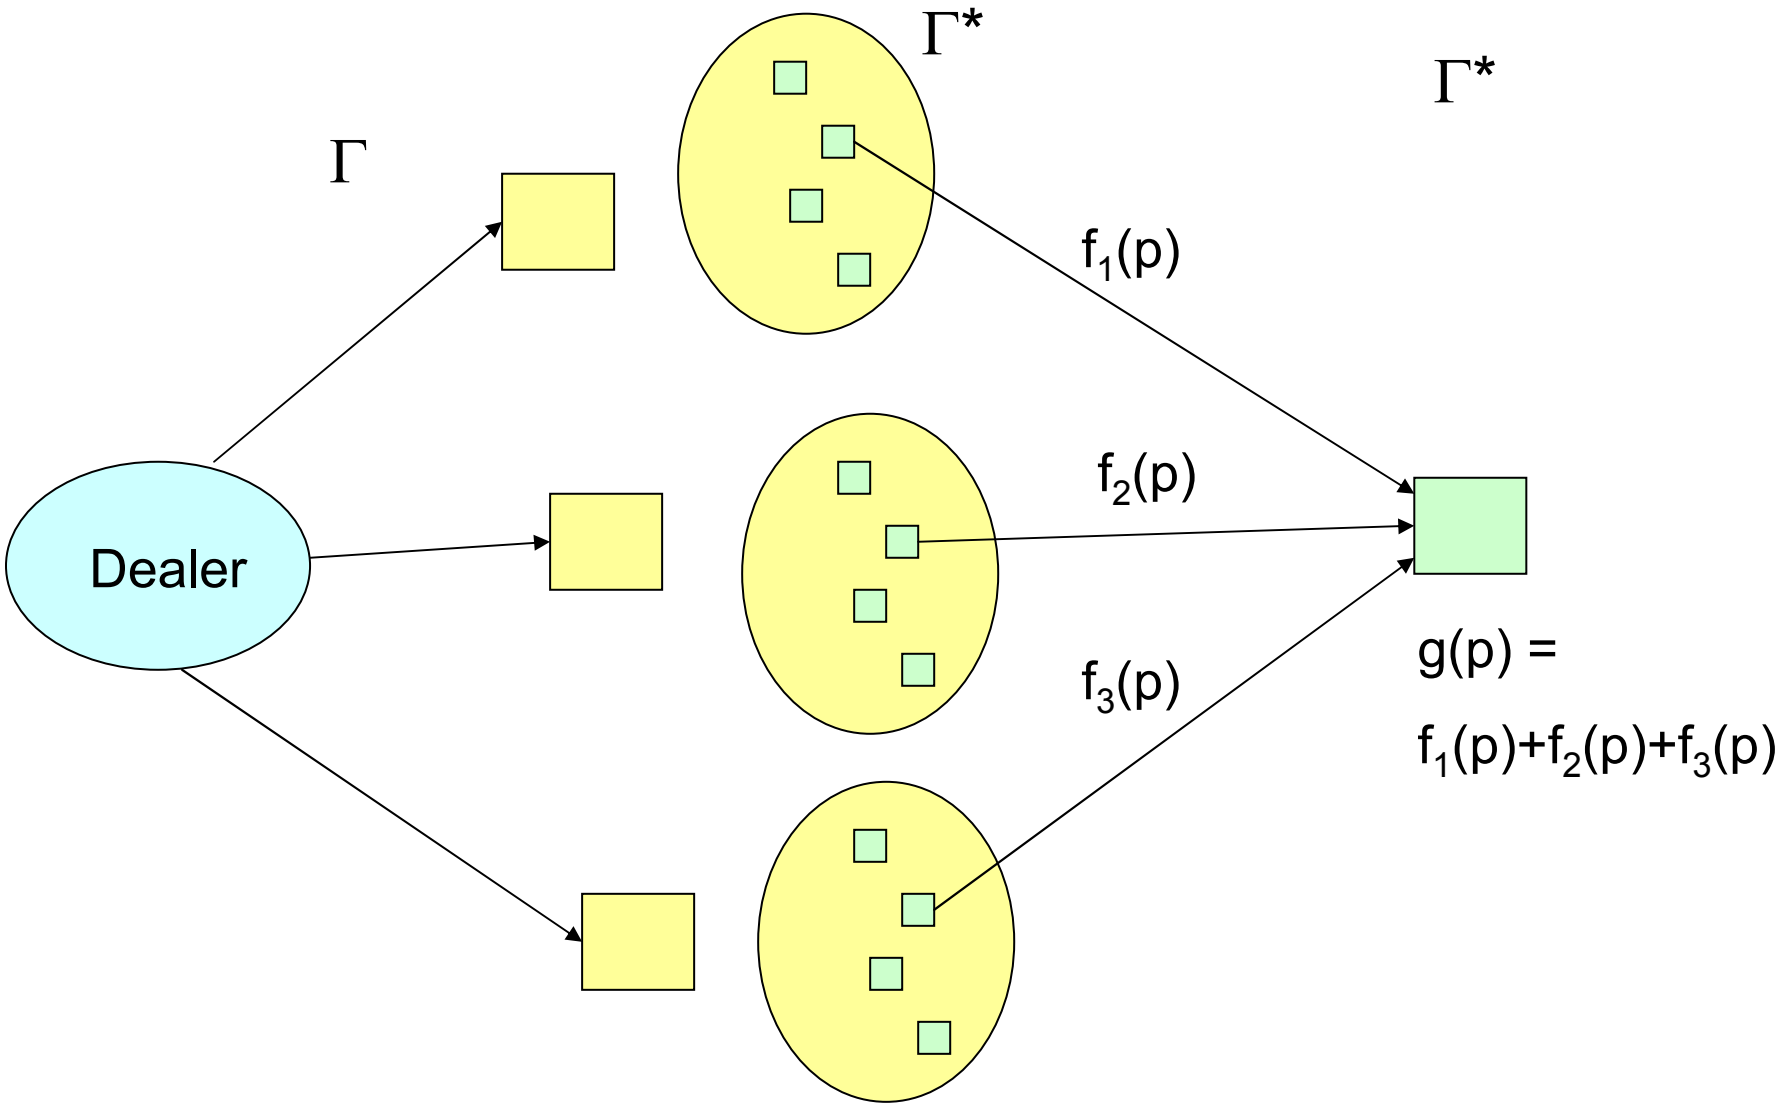

Redistribution schemes

- Basic idea employed by Herzberg et al in proactive secret sharing (Crypto 95)
- First redistribution protocols were exhibited by Desmedt and Jajodia (1997)
- Martin et al formalised the problem (Computer Journal 1999):
  - Provided a formal model
  - Proposed classes of scheme
  - Defined efficiency measures
  - Demonstrated optimal protocols

A redistribution result

Let A and B be minimal sets of the smallest size in Γ and Γ^* respectively. Then for any $\Gamma \rightarrow \Gamma^*$ redistribution we have:

Class	Bound
I/I memoryless	LT=|B|
I/I memory	LT=max{|A|,|B|}
N/I	LT=max{|A|,|B|}
I/N memoryless	LT=|P*|
I/N memory	max{|A|,|P*|} ≤ LT ≤ |A|.|P*|
N/N	max{|A|,|P*|} ≤ LT ≤ |A|.|P*|

Subshare techniques (Model 9C)

- These have been proposed for Model 9 Type C environments (1999)
- Idea is to highly structure shares so that they contain embedded subshares
- The scheme can then be operated either in:
  - Share mode (low threshold)
  - Subshare mode (high threshold)
- Such schemes are highly specialised and inflexible
- Available constructions do not generalise easily

Example subshare technique

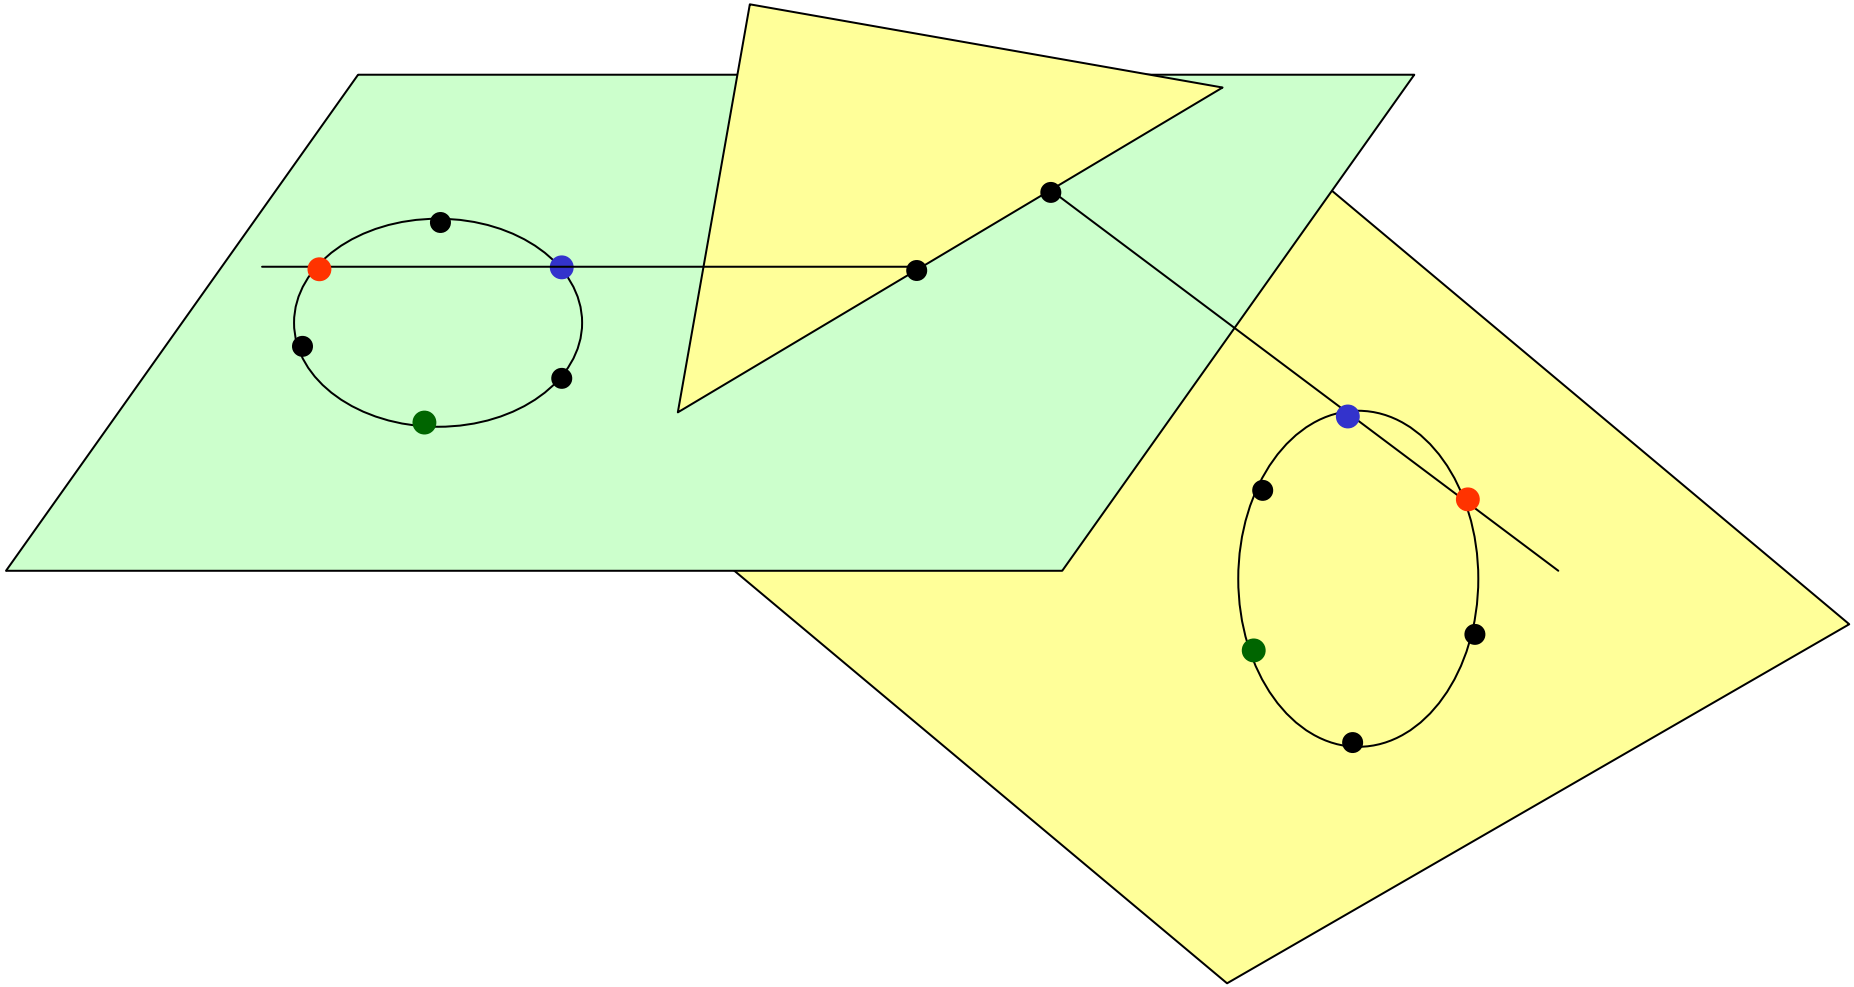

Towards a Model 9A solution?

- A challenge is to design a method of increasing the threshold of a pre-existing threshold scheme (Model 9 Type A) without any dealer assistance
- Steinfeld et al have proposed a technique based on lattice reduction (Asiacrypt 04)
- Scheme does not offer unconditional security (indeed the level of security that it offers is not entirely clear)

Case 3 Open Problems

Open Problem 6

How efficient can redistribution protocols be made if we can consider a more robust model where insiders do not necessarily follow protocols honestly and correctly?

Open Problem 7/8

Is it possible to define any practical schemes in Models 8 or 9 for any of the Types A, B or C?

Case 4: Dealer-free environments

Model	Initialisation Phase			Running Phase		
	D	D - P	P - P	D	D - P	P - P
10			S			S
11			S			B
12			S			-

Appropriate for environments lacking any formal centralised architecture – in other words fully distributed applications.

Case 4: Dealer-free

- Secret sharing in this type of environment has been studied by Ingemarsson and Simmons (Eurocrypt 90) and Jackson et al (J. Cryptology 1997)
- Some redistribution protocols use similar techniques as subroutines, so these cases may be closely related

Open Problem 9

Access policy changes have not been formally investigated in any of the Models 10, 11 or 12.

This talk...

- Conducted a brief overview of work that has been conducted on the problem of updating the access policy of a secret sharing schemes in unconditionally secure models.
- Demonstrated a variety of solutions that vary greatly, primarily depending on the environment in place during the running phase of the scheme.
- Identified a number of open problems.

For the future…

- Open problems in the unconditionally secure model worth looking at.
- The general problem of providing access policy update is of interest in more “practical” security models (supporting distributed cryptographic protocols).