

Knights, Spies, Games and Social Networks

Mark Wildon

16 February 2010



The Knights and Spies Problem

In a room there are 100 people.

- ▶ Each person is either a *knight* or a *spy*.
- ▶ Knights always tell the truth, but spies may lie or tell the truth as they see fit.
- ▶ Everyone in the room knows the identity of everyone else.
- ▶ Knights are in a strict majority.

The Knights and Spies Problem

In a room there are 100 people.

- ▶ Each person is either a *knight* or a *spy*.
- ▶ Knights always tell the truth, but spies may lie or tell the truth as they see fit.
- ▶ Everyone in the room knows the identity of everyone else.
- ▶ Knights are in a strict majority.

Asking only questions of the form

‘Person i , what is the identity of person j ?’


what is the least number of questions that will *guarantee* to find everyone’s true identity?

Outline

1. Questioning strategies
 - ▶ Can it be done at all?
 - ▶ Pairing Up: 200 questions suffice
 - ▶ Spider Interrogation Strategy: 148 questions suffice
2. A two-player game
 - ▶ 148 questions may be necessary in the worst case
3. Open problems and speculative applications
 - ▶ Public key distribution in public key cryptography
 - ▶ Building networks of reputable companies


Section 1: A Simple Questioning Strategy

If we ask the other 99 people in the room about person 1, then the majority opinion is correct.



Section 1: A Simple Questioning Strategy


If we ask the other 99 people in the room about person 1, then the majority opinion is correct.



And everyone expressing the minority opinion is a spy.

Section 1: A Simple Questioning Strategy

If we ask the other 99 people in the room about person 1, then the majority opinion is correct.









And everyone expressing the minority opinion is a spy.

But if person 1 turns out to be a spy, then we may not have made much progress. Proceeding in this way could require 2500 questions.

Pairing Up


Generalize: n people in the room

1. Put people into pairs, with one singleton if n is odd. Ask each member of the pair about the other.




Pairing Up

- Keep the pairs whose members support one another and ignore the rest. Choose a leader from each supportive pair.




Pairing Up

- Repeat the first step by pairing up leaders, until only one person is left.




Pairing Up

4. The final leader is a knight.



Pairing Up

- Use the final leader as an oracle to find all the identities that remain ambiguous.



Pairing Up Uses $O(n)$ Questions

Claim

*The final leader found by
Pairing Up is always a knight.
This strategy uses at most $2n$
questions to find all identities
in a room with n people.*

Pairing Up Uses $O(n)$ Questions

Claim


The final leader found by Pairing Up is always a knight. This strategy uses at most $2n$ questions to find all identities in a room with n people.

Proof.

By induction on n . If after the first pairing up stage there are p pairs, then we are done after

$$n + 2p + (n - 2p) = 2n$$

questions. □




The Spider Interrogation Strategy

Form *spiders*, rejecting candidates as soon as they are accused more often than they are supported.

Example room with 21 people: $S_1, S_2, K_3, K_4, \dots$

Choose the spy S_1 as first candidate.




The Spider Interrogation Strategy

Form *spiders*, rejecting candidates as soon as they are accused more often than they are supported.

Example room with 21 people: $S_1, S_2, K_3, K_4, \dots$

Choose the spy S_1 as first candidate.




The Spider Interrogation Strategy

Form *spiders*, rejecting candidates as soon as they are accused more often than they are supported.


Example room with 21 people: $S_1, S_2, K_3, K_4, \dots$

The spy S_1 is rejected after 3 questions.




The Spider Interrogation Strategy

Candidates are accepted according to a (varying) threshold for support. After S_1 , K_5 and S_{13} are rejected, K_{15} is accepted.




The Spider Interrogation Strategy

Candidates are accepted according to a (varying) threshold for support. After S_1 , K_5 and S_{13} are rejected, K_{15} is accepted.




The Spider Interrogation Strategy

Use the successful candidate to identify the rejected candidates, and those not yet involved in proceedings. Then identify people within spiders.




The Spider Interrogation Strategy

Use the successful candidate to identify the rejected candidates, and those not yet involved in proceedings. Then identify people within spiders.



The Spider Interrogation Strategy

The Spider Interrogation Strategy uses $< 3n/2$ questions in an n person room, no matter how the spies behave.



Section 2: A Two-Player Game

In the game of *Knights and Spies*, an *Interrogator* plays against a *Spy Master*.

- ▶ The Interrogator decides which questions to ask.
- ▶ The Spy Master supplies the answers, and (indirectly) determines who is a spy.

Section 2: A Two-Player Game

In the game of *Knights and Spies*, an *Interrogator* plays against a *Spy Master*.

- ▶ The Interrogator decides which questions to ask.
- ▶ The Spy Master supplies the answers, and (indirectly) determines who is a spy.

Let

$$f(n) = \begin{cases} 3m - 3 & \text{if } n = 2m - 1 \\ 3m - 2 & \text{if } n = 2m \end{cases}$$

If the Interrogator can be certain of everyone's identity after asking $< f(n)$ questions, then he wins. If exactly $f(n)$ questions are needed, the game is drawn. Otherwise the Spy Master wins.

Section 2: A Two-Player Game

In the game of *Knights and Spies*, an *Interrogator* plays against a *Spy Master*.

- ▶ The Interrogator decides which questions to ask.
- ▶ The Spy Master supplies the answers, and (indirectly) determines who is a spy.

Let


$$f(n) = \begin{cases} 3m - 3 & \text{if } n = 2m - 1 \\ 3m - 2 & \text{if } n = 2m \end{cases}$$

If the Interrogator can be certain of everyone's identity after asking $< f(n)$ questions, then he wins. If exactly $f(n)$ questions are needed, the game is drawn. Otherwise the Spy Master wins.

If all consistent interpretations of the Spy Master's answers require spies to be in the majority, the Spy Master forfeits the game. The program *GameChecker.hs* can be used to referee games.


An Optimal Strategy for the Spy Master

Answer the first $n/2 - 2$ questions with blanket accusations.




An Optimal Strategy for the Spy Master

Answer the first $n/2 - 2$ questions with blanket accusations. Then hide a single knight in each accusatory component.



An Optimal Strategy for the Spy Master

Answer the first $n/2 - 2$ questions with blanket accusations. Then hide a single knight in each accusatory component.



An Optimal Strategy for the Spy Master

This strategy ensures that the Interrogator cannot be certain of every identity until he has asked $f(n)$ questions.

The Spider Interrogation Strategy uses at most $f(n)$ questions.

Hence, with optimal play Knights and Spies is a draw. This solves the original problem: in an n person room, $f(n) \approx 3n/2$ questions are sufficient, and in the worst case, necessary.

Section 3: Open Problems and Speculative Applications

Is there a questioning strategy that never uses more than $3n/2$ questions, and on average uses αn questions for some constant $\alpha < 3/2$?

Section 3: Open Problems and Speculative Applications

Is there a questioning strategy that never uses more than $3n/2$ questions, and on average uses αn questions for some constant $\alpha < 3/2$?

Probably yes, but I can't prove it.

Section 3: Open Problems and Speculative Applications

Is there a questioning strategy that never uses more than $3n/2$ questions, and on average uses αn questions for some constant $\alpha < 3/2$?

Probably yes, but I can't prove it.

One hopeful idea: build up long chains in the search for a knight, and later identify their members by repeated bisection.




● K

Section 3: Open Problems and Speculative Applications

Is there a questioning strategy that never uses more than $3n/2$ questions, and on average uses αn questions for some constant $\alpha < 3/2$?

Probably yes, but I can't prove it.

One hopeful idea: build up long chains in the search for a knight, and later identify their members by repeated bisection.




Section 3: Open Problems and Speculative Applications

Is there a questioning strategy that never uses more than $3n/2$ questions, and on average uses αn questions for some constant $\alpha < 3/2$?

Probably yes, but I can't prove it.

One hopeful idea: build up long chains in the search for a knight, and later identify their members by repeated bisection.




Section 3: Open Problems and Speculative Applications

Is there a questioning strategy that never uses more than $3n/2$ questions, and on average uses αn questions for some constant $\alpha < 3/2$?


Probably yes, but I can't prove it.

One hopeful idea: build up long chains in the search for a knight, and later identify their members by repeated bisection.




Simulations of Chain Building

Number of questions asked in 100 person rooms with 49 randomly allocated spies. Top: spies always lie. Bottom: spies always accuse.



Simulations of Chain Building

Number of questions asked in n person rooms in which knights are just in the majority. Spies always lie. The interpolating line has gradient 1.301.



Public Key Distribution in Public Key Cryptography

Alice would like to use the RSA Protocol to email a message M to Bob. She doesn't know Bob's public key P_B , but her friend Mallory does, and he sends her it.

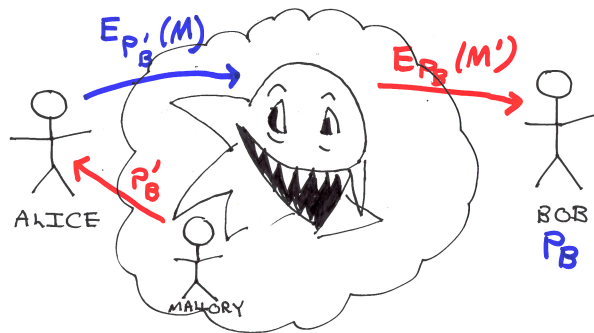
Alice sends Bob the encrypted message $E_{P_B}(M)$, and Bob acknowledges receipt (by unencrypted email say). All is well

...

Public Key Distribution in Public Key Cryptography

Alice would like to use the RSA Protocol to email a message M to Bob. She doesn't know Bob's public key P_B , but her friend Mallory does, and he sends her it.

Alice sends Bob the encrypted message $E_{P_B}(M)$, and Bob acknowledges receipt (by unencrypted email say). All is well ... unless Mallory is not what he seems.



Public Key Distribution in Public Key Cryptography

As time goes on, you will accumulate keys from other people that you may want to designate as trusted introducers. Everyone else will each choose their own trusted introducers. And everyone will gradually accumulate and distribute with their key a collection of certifying signatures from other people, with the expectation that anyone receiving it will trust at least one or two of the signatures. This will cause the emergence of a decentralized fault-tolerant web of confidence for all public keys.

Philip Zimmermann, Pretty-Good-Privacy Manual

Public Key Distribution in Public Key Cryptography

Potential difficulties:

- ▶ Doesn't work well with small ad-hoc arrangements.
- ▶ Even people acting in good faith may unwittingly behave like spies.
- ▶ Have to think hard about who one trusts, and not lazily accept every key.

Idea: use algorithms from the Knights and Spies Problem to advise the user.

Ideally most decisions would be safely handled by the computer. The user should only be bothered when a threat is strongly indicated.

Building Networks of Reputable Companies

Consider companies offering a fairly specialised professional service, say print shops. Reputable operations tend to have some idea how good their competitors are, and will preserve their reputation by answering honestly when asked about them. Fraudulent operations don't care, and anyway, are likely to lie. Spies are (hopefully) in the minority.

The Knights and Spies framework could be used to decide who is reputable. No specialised knowledge of printing is required.

Application to cloud computing: we'd like to test our agents. But maybe the work required is so great, we have to find a way to let the (untrustworthy) cloud do it for us.

Original suggestion due to Miranda Mowbray (HP Labs)

Thank you. Any questions?

Thank you. Any questions?

For my paper, and programs for playing Knights and Spies and simulating questioning strategies, go to

<http://www.maths.bris.ac.uk/~mzmjw>.

Four More Open Problems

- ▶ What is the smallest number of questions it takes to be certain of at least one person's identity?
- ▶ What is the smallest number of questions it takes to find a knight?
- ▶ Knights and Spies for logicians: 'Person i , what is your view on the formula $j \implies (k \wedge \ell)$?'
- ▶ Is the problem of deciding whether an incomplete game of Knights and Spies in a consistent state NP-complete?