On the Isolation of a Common Secret

Don Beaver, Stuart Haber and Peter Winkler*

July 22, 2002

Abstract

Two parties are said to "share a secret" if there is a question to which only they know the answer. Since possession of a shared secret allows them to communicate a bit between them over an open channel without revealing the value of the bit, shared secrets are fundamental in cryptology.

We consider below the problem of when two parties with shared knowledge can use that knowledge to establish, over an open channel, a shared secret. There are no issues of complexity or probability; the parties are not assumed to be limited in computing power, and secrecy is judged only relative to certainty, not probability. In this context the issues become purely combinatorial and in fact lead to some curious results in graph theory.

Applications are indicated in the game of bridge, and for a problem involving two sheriffs, eight suspects and a lynch mob.

1 Introduction

Suppose two parties—let us call them "Alice" and "Bob"—share a secret, that is, they have common knowledge possessed by no one else; then Alice may use her secret to transmit a bit to Bob in such a way that no eavesdropper can deduce the value of the bit. For example, if Alice and Bob are the only people in the world who know whether President Bush wears a wig, Alice may send Bob the following message (or the same message, with 0 and 1 interchanged):

"My bit is 0 if the President wears a wig, 1 otherwise."

While Eve (an eavesdropper) may believe that the President probably does not wear a wig, and therefore that Alice's bit is more likely to be 1 than 0, her inability to determine the value of the bit with certainty is all that concerns us here.

This method of encryption is called a "one-time pad"; Alice and Bob share a bit of information, and they can use it (once) to pass a bit in secret.

There are, however, situations where even though Alice and Bob appear to possess shared information not available to the public, this information does not take the form of a shared secret. Nonetheless, Alice and Bob may be able to *isolate* a shared secret by communicating with each other, even though their messages are public. Since this is precisely the situation where cryptologic methods are needed (communication lines are available, but not private), Alice and Bob are almost as well off here as if they had begun with a shared secret; they must merely spend a few preliminary rounds of communication in establishing the secret.

Let us give two examples of such situations, before proceeding further.

*This is a reprint of an article which appeared in *The Mathematics of Paul Erdős* Vol. II, R.L. Graham and J. Nešetřil, eds., Springer-Verlag, Berlin 1996.

(1) **The game of bridge:** Here two partners wish to communicate in private, but the rules of the game require that all communication be done by legal bids and plays, about which there may be no prior private understandings. Thus, there are initially no shared secrets. But there is private information: each player knows, by virtue of looking at his own hand, 13 cards that do *not* belong to his partner. Can they make use of this information to communicate in private?
(2) **The 'two sheriffs' problem:** Two sheriffs in neighboring towns are on the track of a killer, in a case involving eight suspects. By virtue of independent, reliable detective work, each has narrowed the list to only two. Now they are engaged in a telephone call; their object is to compare information, and if their pairs overlap in just one suspect, to identify him (the killer) and put out a.p.b.'s so as to catch him in either town.

The difficulty is that their telephone line has been tapped by the local lynch mob, who know the original list of suspects but not which pairs the sheriffs have arrived at. If they are able to identify the killer with certainty as a result of the phone call, he will be lynched before he can be arrested.

Can the sheriffs accomplish their objective without tipping off the mob?

2 The Mathematical Model

One natural model for common knowledge is obtained by imagining that in any situation there is an underlying finite space S of *possibilities* of which any one element may be "the truth." Alice's knowledge concerning S at any point consists of some subset X of S, meaning that X is precisely the set of truths consistent with what Alice knows. As Alice communicates with Bob she obtains more information, and her knowledge set X shrinks accordingly.

At any time the "true point" must lie both in X and in Bob's knowledge set Y, but if there are two or more points in $X \cap Y$ Alice and Bob will never be able to choose among them by communicating with each other. Hence for our purposes if X is a *possible* knowledge set for Alice and Y for Bob, our only concern is whether they intersect.

Consequently we choose to model common knowledge by using a mere vertex to represent a possible knowledge set of Alice's, and similarly for Bob; we connect vertex x of Alice's with vertex y of Bob's when the corresponding knowledge sets intersect, that is, when the two vertices are simultaneously possible. The "truth" is thus represented by some adjacent pair of vertices, i.e. an edge.

Alice and Bob's knowledge at any time thus constitutes a graph, which, in accordance with cryptographic tradition, is assumed to be known to everyone in the world. The interpretation of this graph will, we hope, become clear to the reader after some examples.

It is convenient to formalize our model as follows.

Definition 2.1. *A bigraph is a finite, non-empty collection H of ordered pairs such that if (x,y) is in H then (y,z) is not.*

Elements of the set $A(H) := \{x : (x,y) \in H \text{ for some } y\}$ will be termed "Alice's vertices" and are perforce distinct from the symmetrically defined "Bob's vertices" in $B(H)$. Thus the elements of H are edges of a bipartite graph, but note that the vertices come equipped with a labelled left-right (Alice-Bob) partition and that isolated vertices cannot arise.

Our model now consists of a bigraph H, known to all, the edges of which represent possible truths. Alice knows the endpoint in $A(H)$ of the true edge, Bob its endpoint in $B(H)$; in other words, if the true edge is (x,y) then Alice knows x and Bob knows y.

We say that Alice and Bob *share a secret* if there is a question to which they know the answer and Eve does not. In the wig example, the question "Does the President wear a wig?" can be answered only by Bob and Alice, so they indeed share a secret in this case. Here, the bigraph H consists of a pair of disjoint edges, one corresponding to "the President wears a wig" and the other to "the President does not wear a wig." The disconnectivity of H is its crucial property:

Theorem 2.2. *Two parties share a secret if and only if their bigraph is disconnected.*

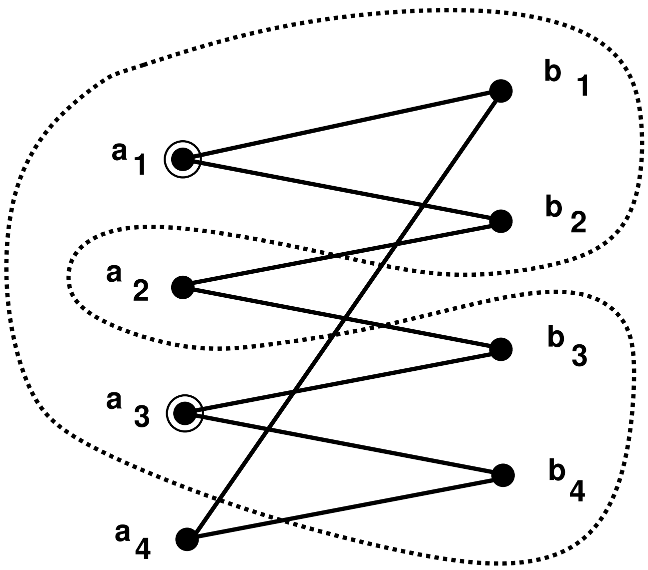

Figure 1: Alice separates an 8-cycle

Proof. It is immediate that Alice and Bob share a secret whenever their bigraph is disconnected, since if C is one of its connected components, only they can answer the question "Is the true edge in C ?". To see the converse, let Q be the given question and let $a_1, a_2, \ldots$ be its possible answers (from Eve's point of view). Write "$(u,v)\#a_i$" if it is simultaneously possible for a_i to be the answer to Q, and (u,v) to be the true edge of H.

We now note that if $(x,y)\#a_i$ and $(u,v)\#a_j$ for $i \neq j$, then (x,y) and (u,v) can be neither identical nor adjacent; if, for example, $x = u$ then when Alice's end of the true edge is x she will be unable to decide between answers a_i and a_j.

It follows that the edges consistent with the various answers a_i determine a partition of H, each part of which is a non-empty union of connected components; since the number of possible answers must be at least two, H is disconnected. □ □

If H is connected, are Alice and Bob doomed never to share a secret? Of course, if they can arrange a private (secure) conversation, they can agree on some string of bits and thus share as many secrets as they wish; this indeed is often done in traditional cryptography, in the name of agreeing on or distributing *key*. Unfortunately this phase is often dangerous and sometimes impossible; else, cryptography would be unnecessary. However, Alice and Bob may be able to use their common knowledge (reflected in the structure of H) to isolate a common secret by means of a *public* conversation; and it is just this process which we wish to investigate.

Consider, for example, the bigraph $H = \{(a_1,b_1), (a_1,b_2), (a_2,b_2), (a_2,b_3), (a_3,b_3), (a_3,b_4), (a_4,b_4), (a_4,b_1)\}$. H is an 8-cycle, thus connected, but Alice and Bob can disconnect it as follows: if Alice's end of the true edge is a_1 or a_3 she says so: "My end of the true edge is either a_1 or a_3." Bob can tell by looking at his end which of the two possibilities is the case, hence they now share a secret; this is reflected in the fact that after Alice's announcement, their bigraph is disconnected (see Fig. 1). Of course, had Alice's end been a_2 or a_4, an announcement to that effect would also have done the job. (We shall see later that two-sided conversations may be necessary to disconnect some bigraphs.)

The next example is inspired by the work of Winkler [4, 5] on cryptologic techniques for the game of bridge, and the recent works of Fischer, Paterson and Rackoff [2] and Fischer and Wright [3] using random deals of cards for cryptographic key. A card is dealt at random, face down, to each of Alice and Bob, from a three-card deck consisting (say) of an Ace, a King, and a Queen. The remaining card is discarded unseen. Here H is a 6-cycle (see Fig. 2). The six edges

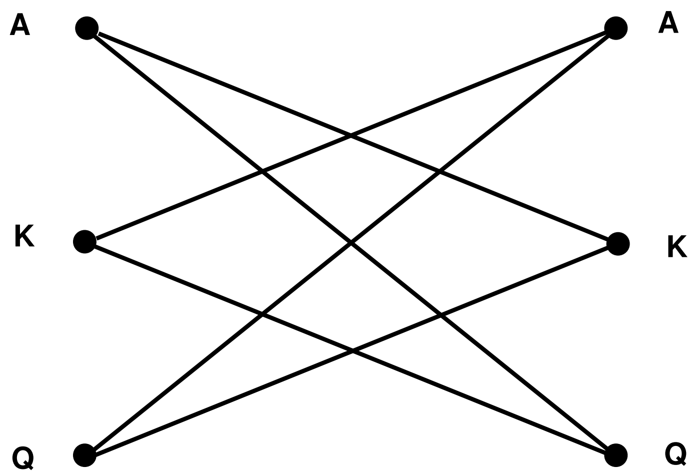

Figure 2: Bigraph for dealing 2 cards from a 3-card deck

correspond to the six possible deals; the set $A(H)$ consists of Alice's three possible holdings ("A", "K" or "Q") and similarly for $B(H)$. The fact that "K" in $A(H)$ and "K" in $B(H)$ are not adjacent corresponds to the fact that Alice and Bob cannot both hold the King. Clearly Alice and Bob share *some* knowledge in this situation, but as we shall see it is not enough for them to be assured of being able to isolate a shared secret.

For a third example, suppose H is a 10-cycle on vertices $v_0, v_1, \ldots, v_9$ with v_i adjacent to v_j just when $|i - j| = 1 \bmod 10$. (See Fig. 3.)

Let $A(H)$ consist of the vertices of even index. Then the following protocol allows Alice to disconnect the bigraph: if she holds v_i, she chooses j to be *either* $i + 4$ or $i + 6 \bmod 10$, then tells Bob:

"I hold either v_i or v_j."

This protocol is said to be *non-deterministic* because Alice has more than one choice of message for a given holding. Non-determinism as used here is thus quite different from its use in complexity theory, and in fact is more closely related in some respects to randomization (despite the absence of probability in our model). In particular non-deterministic communication protocols are quite practical.

Incidentally, we assume nothing is given away by the order in which objects are named in a semantically symmetric expression such as "v_i or v_j." In a deterministic protocol we can insure this by a naming convention, such as enforced alphabetical order, but our mathematical model for communication will obviate the problem.

In order to define "communication protocol" rigorously we need to define "conversation", even though the latter definition needs the former for interpretation. Accordingly, a *conversation* will be a finite string $m_1, \ldots, m_t$ of positive integers, communicated alternately by Alice and Bob beginning with Alice.

Although this seems perhaps a limited vocabulary for communication, it is in fact completely general because meanings can be assigned to the numbers via an agreed-upon protocol. The communication protocol specifies *under precisely what circumstances a given number may be uttered.* Thus, we operate in effect in the "political" theory of communication: when someone says something we ask not "what does that mean" but "under what circumstances would he/she have said that." This model is both stronger and simpler than one in which messages are sentences. To see its effect, consider solving the C_{10} bigraph above by using always vertices i and $i+4$. This looks reasonable at first glance but is actually a sham. For example, Alice would not then say "My end is 4 or 8" if she held 8. Hence an eavesdropper can eliminate 8 if she hears this, preventing disconnection of the bigraph and ruining the protocol.

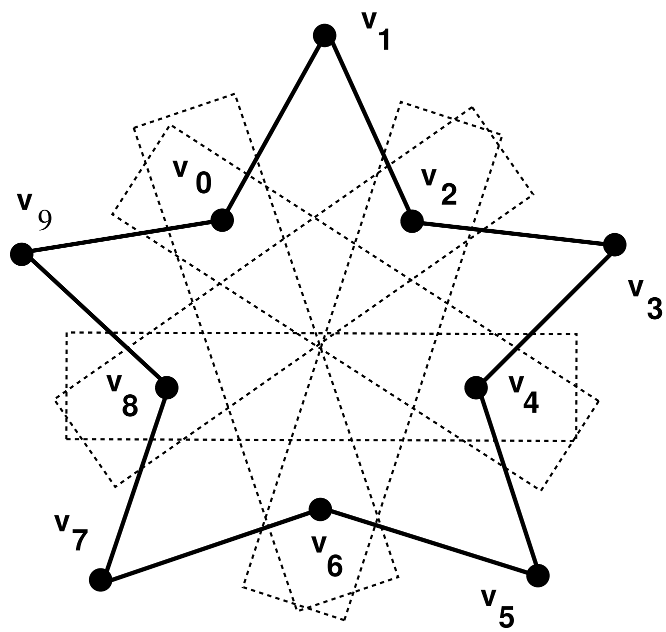

Figure 3: Non-deterministically separating a 10-cycle

3 Deterministic Separation

Definition 3.1. *A* deterministic communication protocol *for the bigraph H is a sequence of functions $f_1, \ldots$ into the positive integers, such that f_1 is a function of x (Alice's end of the true edge), and each subsequent f_i depends on the values of $f_1, \ldots, f_{i-1}$ and either on x, if i is odd, or on y (Bob's end of the true edge) if i is even.*

Thus, a deterministic communication protocol, combined with a true edge, produces a unique conversation given by $m_i = f_i(x; m_1, \ldots, m_{i-1})$ for i odd and $m_i = f_i(y; m_1, \ldots, m_{i-1})$ for i even. After step i of the conversation, the situation can be again described by a bigraph H^i consisting of those edges (u, v) such that if (u, v) had been the true edge the conversation would have been as seen. H^i is a sub-bigraph of H which contains the actual true edge; in fact H^i is obtained from H^{i-1} by deleting all edges incident to certain left-hand vertices (for i odd) or certain right-hand vertices (for i even). The vertices which survive on the left (when i is odd) are just $\{u : f_i(u; m_1, \ldots, m_{i-1}) = m_i\}$.

For our purposes a conversation $(m_1, \ldots, m_t)$ will be deemed "successful" if H^t is disconnected, and a communication protocol for H will be said to "separate H" or to be a "separation protocol for H" if it always produces a successful conversation. Finally, H itself will be termed *deterministicallyseparable* if there is a deterministic communication protocol which separates H.

We can also give a recursive characterization of the deterministically separable bigraphs. It will be useful to introduce notation for the *sub-bigraph* $H|A'$ of a bigraph H induced by a subset A' of $A(H)$; namely,

$$H|A' := H \cap (A' \times B(H)) .$$

Thus from the standpoint of ordinary graph theory, $H|A'$ is obtained by discarding isolated vertices from the subgraph of H induced by $A' \cup B(H)$. The definition of $H|B'$ for $B' \subset B(H)$ is similar.

Theorem 3.2. *Let* **DS** *be the smallest symmetric class of bigraphs which contains the disconnected bigraphs and has the following property: for any bigraph H, if there is a partition*

$A(H) = A_1 \cup A_2$ of Alice's vertices such that $H|A_1$ and $H|A_2$ are both in **DS***, then H is also in* **DS***. Then* **DS** *is the class of deterministically separable bigraphs.*

Proof. Let us check first that the class of deterministically separable bigraphs is indeed closed under the operation defined in the statement of the theorem. It is certainly symmetrical, since a protocol which separates H can be modified to one which separates the dual of H (i.e. the result of reversing the ordered pairs in H) by switching the roles of Alice and Bob, and adding a meaningless message from Alice to the beginning. If a partition is given along with separation protocols P_1 and P_2 for the two sub-bigraphs of H, we design a separation protocol P for H as follows: first Alice sends "i" iff her end of the true edge lies in A_i, then Bob sends back a meaningless "1", then protocol P_i is followed.

It remains to show that any symmetric class **C** containing the disconnected bigraphs and closed under the stated operation contains all deterministically separable bigraphs. This is done by induction on the number of edges.

Let H be connected but deterministically separable via a communication protocol P; then sooner or later P must call for a first meaningful message, say from Alice. We may assume that her options at that point are to send "1", "2", etc. up to "k" for some $k \geq 2$, according to whether her end of the true edge lies in E_1, E_2, etc; there is no dependence on previous conversation here because by assumption said conversation has up to now been predictable. Let us modify P slightly by having Alice send only a 1 or 2 at this point, the former just when her end is in E_1; if she sends a "2" Bob sends a meaningless "1" back, then Alice sends the "2", "3",... or "k" that would have been sent before and the protocol P resumes. We have thus found protocols which separate each of $H|E_1$ and $H|((A(H) - E_1)$. Each of these is thus in **C** by the induction assumption, hence $H \in$ **C** by the closure condition. □ □

4 Non-Deterministic Separation

We now introduce two ways of weakening the definition of a communication protocol. First, if the functions f_i are permitted to be *multi*-valued, so that at each point Alice or Bob has one *or more* possible messages to send, we say that the protocol is *non-deterministic*. Note that during a non-deterministic communication protocol, the bigraph shrinks as before but this time each party, at his or her turn to speak, is provided with a (message-labelled) *cover* of his or her vertices instead of a *partition*. It is still the case, however, that the knowledge of Alice and Bob is expressed at each point by the state of their bigraph.

Second, if the functions are permitted to depend for odd i on a random number m known only to Alice, and for even i on a random number n known only to Bob, then the protocol is *randomized*. Here the bigraph does *not* any longer describe the situation completely, as Bob and Alice may learn things about each other's random number.

Whether a communication protocol is non-deterministic, randomized or both, however, we continue to insist that the protocol *always* produce a successful conversation in order to qualify for separating H.

Fortunately, the three new categories of separation protocol which arise result in only one new category of bigraph.

Theorem 4.1. *The following are equivalent, for any bigraph H:*

(a) *H is separable by a non-deterministic communication protocol;*

(b) *H is separable by a randomized communication protocol;*

(c) *H is separable by a randomized, non-deterministic communication protocol.*

Proof. We need to show (c)→(b) and (c)→(a), the reverse implications being trivial. Of these the former is easy: by extending the range of the random numbers, Alice and Bob can use them to decide which message to send when there is more than one choice.

Turning random numbers into non-deterministic choices looks awkward because a random number may be used many times in the protocol, whereas there is no "consistency" built in to nondeterminism. However, this problem is illusory. Suppose, at Alice's first turn to speak, that she is supplied with a randomized separation protocol but no random number; then she chooses a random number and acts accordingly. She cannot "remember" that number at her next turn and use it again, but she can compute which random numbers are consistent with her previous action and choose one of those upon which to base her next message. Bob behaves similarly; at each turn he determines which values of his non-existent random number are consistent with his own previous actions (and his end of the true edge), then picks one such value and acts accordingly. □ □

A bigraph will be called, simply, *separable* if one (thus all) of the conditions of Theorem 4.1 obtains. The following recursive characterization is analogous to Theorem 3.2. although a small additional subtlety arises in the proof.

Theorem 4.2. *Let* $\mathbf{S}$ *be the smallest symmetric class of bigraphs which contains the disconnected bigraphs and has the following property: for any bigraph* H*, if there is a covering* $A(H) = A_1 \cup A_2$ *of Alice's vertices such that* $H|A_1$ *and* $H|A_2$ *are both in* $\mathbf{S}$*, and both* A_1 *and* A_2 *are strictly contained in* $A(H)$*, then* H *is also in* $\mathbf{S}$*. Then* $\mathbf{S}$ *is the class of separable bigraphs.*

Proof. The proof that the class of separable bigraphs is symmetrical and closed under the operation defined in the statement of the theorem is as in Theorem 3.2, except that if Alice's end of the true edge lies in $A_1 \cap A_2$ she may send *either* message "1" or message "2".

It remains to show that any symmetric class $\mathbf{C}$ containing the disconnected bigraphs and closed under the stated operation contains all separable bigraphs; this is again done by induction on the number of edges.

Let H be connected but separable via a non-deterministic communication protocol P, and suppose that H is the smallest separable bigraph not in $\mathbf{C}$. Let us call a vertex u in $A(H)$ (or, dually, in $B(H)$) *weak* if there is no proper subset A' of $A(H)$ containing u for which $H|A'$ is separable.

We claim that there is some weak vertex in $A(H)$. For, if not, define for each $x \in A(H)$ a proper subset A_x which does yield a separable sub-bigraph of H. Since the A_x's cover $A(H)$, we can find $x_1, x_2, \ldots, x_k$ such that the A_{x_i}'s cover $A(H)$ with k minimal (but necessarily greater than 1). Set $A_1 := A_{x_1}$, $A_2 := A_{x_2} \cup \cdots \cup A_{x_k}$. Then A_1 and A_2 are a proper cover of $A(H)$, and $H|A_1$ is separable by assumption. However, $H|A_2$ is also separable, since Alice can reduce it to a separable bigraph by sending some i for which her end of the true edge lies in A_{x_i}, $2 \leq i \leq k$. These bigraphs are thus both in $\mathbf{C}$ by the induction assumption, contradicting the fact that H is not in $\mathbf{C}$.

Since the class of separable bigraphs is symmetric, the dual of H is also separable but not in $\mathbf{C}$; hence the same argument produces a weak vertex in $B(H)$. It may seem to the reader that weak vertices cause trouble only if found on the true edge, and thus that Alice and Bob can't *both* be stymied as long as no two weak vertices are adjacent. However, it turns out that the mere *presence* of weak vertices on both sides is enough to render H inseparable.

To see this, let u be a weak vertex in $A(H)$; there must be some message (say m) which Alice is permitted to send when her end of the true edge is u. Let A' be the set of vertices in $A(H)$ which, like u, allow the message m; then $H|A'$ must be separable, since this is the bigraph which results when m is sent. Thus A' must not be a proper subset of $A(H)$, that is, $A' = A(H)$ and the message m is meaningless.

If m is indeed sent, the protocol turns to Bob who still has all of H before him. By the same reasoning as above, he must also have a meaningless message (i.e. a message he can send regardless of which vertex is his end of the true edge) available to him.

Now we're back to Alice with H still intact. We thus see that Alice and Bob must be allowed by the protocol to pass meaningless messages back and forth ad infinitum, irrespective of which edge of H is the true edge; but then we have a contradiction, since H is required to have a communication protocol which always separates. □ □

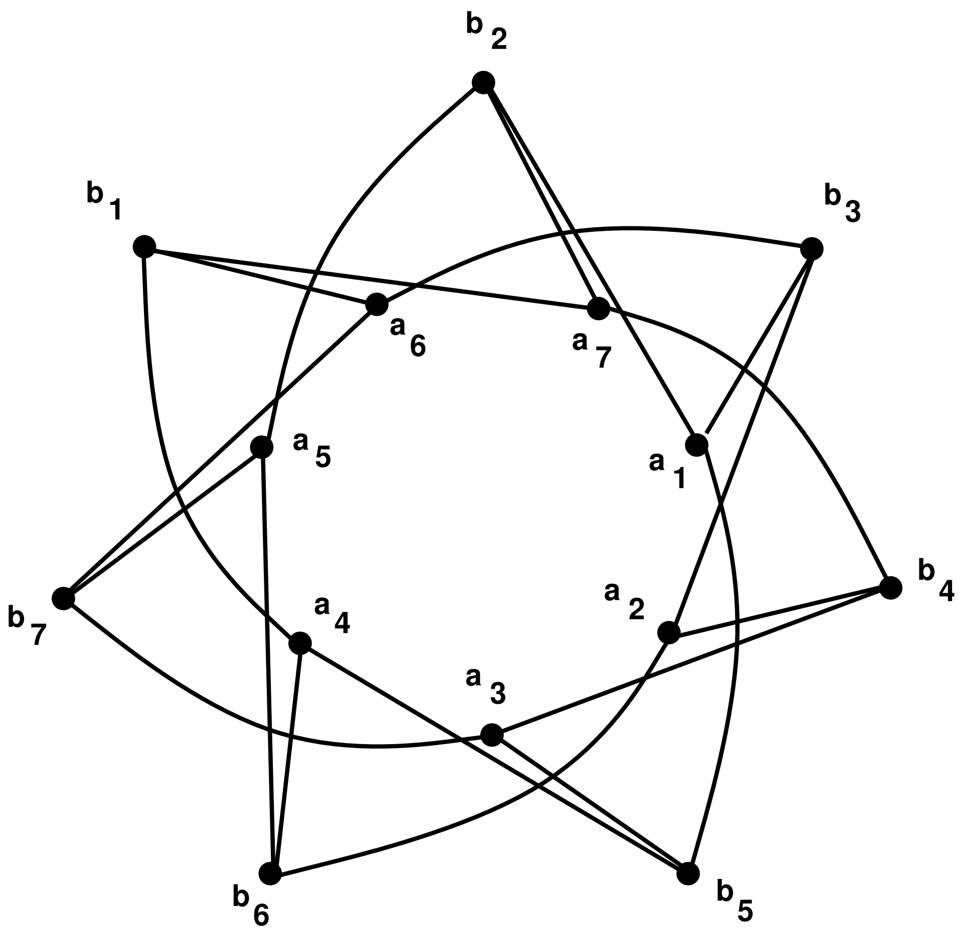

Figure 4: Incidence graph of the Fano plane

Theorem 4.2 is often useful in determining separability via case analysis. For example, it is easy to check that no path with fewer than 5 edges is separable, nor is the 6-cycle (Fig. 2) separable because for any proper subset S of Alice's or Bob's vertices, $H|S$ would be a path of length 2 or 4.

There is one class of bigraphs which is easily seen to be disjoint from the class of separable bigraphs, a fact which helps in obtaining negative results.

Theorem 4.3. *Suppose that there is an edge of H which is adjacent to all other edges of H. Then H is not separable.*

Proof. Such an edge cannot be contained in any disconnected sub-bigraph of H; thus, if it happens to be the true edge, no protocol can separate H. □ □

Note that, in particular, no separable bigraph can have a vertex which is adjacent to all the vertices on the other side.

We are now in a position to prove that that the class of separable bigraphs is strictly larger than the class of deterministically separable bigraphs.

Let $A(H) = \{a_1, \ldots, a_7\}$ and $B(H) = \{b_1, \ldots, b_7\}$, and put $(a_i, b_j) \in H$ iff $j - i = 1$, 2 or 4, where the indices are interpreted always modulo 7. Then H is the incidence graph of a Fano plane (see Fig. 4) and we have:

Theorem 4.4. *The incidence graph of the Fano plane is separable but not deterministically separable.*

Proof. We first provide a non-deterministic separation protocol. Alice begins by sending a number k such that her end of the true edge is in the set $\{a_k, a_{k+1}, a_{k+2}\}$. Bob now (deterministically) sends back "1" if his end is b_{k+2} or b_{k+6}; "2" if it is b_{k+4} or b_{k+5}; "3" if it is b_{k+3} or b_{k+1}. (It cannot be b_k). This separates H into a 2-edge component and a 1-edge component.

If, on the other hand, there were a deterministic separation protocol for H, then one of the parties would eventually have to send a meaningful message, thus effecting a partition of his or her vertices. This may as well be Alice since H is symmetrical. If one of the parts has fewer than 3 of Alice's vertices in it, or has 3 vertices whose neighborhoods intersect, then one of Bob's remaining points will be of full degree, contradicting Theorem 4.3. Otherwise the partition must be isomorphic to $\{a_1, a_2, a_3\}$ versus $\{a_4, a_5, a_6, a_7\}$. The former part induces a deterministically separable sub-bigraph as we have seen from the above protocol, but the sub-bigraph induced by the latter part contains a vertex (b_1) whose neighborhood intersects the neighborhoods of all other vertices of $B(H)$. Thus, if Bob's end of the true edge is b_1, he cannot separate the bigraph at this time. However, Alice is also stymied because she has only one vertex available (a_5) not adjacent to b_1, thus her vertices can never be partitioned so as to induce disconnected sub-bigraphs. □ □

The situation changes if we consider bigraphs which are separable *in one round*, that is, by just one message from Alice. (This is not, by the way, a symmetric class; the path on 7 vertices, for example, can be separated only by a message from the party with 4 vertices.) Before proceeding, we need a curious graph-theoretic result.

Theorem 4.5. *Let G be any graph with no vertex adjacent to all others. Then there is a partition $V_1, \dots, V_k$ of the vertices of G such that for each $i = 1, \dots, k$ the subgraph $\langle V_i \rangle$ induced by V_i is disconnected.*

Proof. If not let G be a counterexample with smallest possible number of vertices. For any subset U of the set of vertices V, let $\omega(U)$ be the number of vertices in $V - U$ which are adjacent to all other vertices in $V - U$.

Note first that if $\langle U \rangle$ is disconnected, then $\omega(U)$ must be non-zero; else we may apply the induction hypothesis to get a suitable partition of $V - U$, and appending U itself to this partition yields a partition of V suitable for G.

Choices of U for which $\langle U \rangle$ is disconnected do exist, of course, since U can be taken to be a pair of non-adjacent vertices. Hence we may choose a U for which $\langle U \rangle$ is disconnected and $\omega(U)$ is minimal.

Now let x be any full vertex in $V - U$, that is, any vertex in $V - U$ which is adjacent to all other vertices in $V - U$. By assumption x has *some* non-neighbor, say y, in V; let C be the set of vertices of the component of $\langle U \rangle$ into which y falls.

Suppose first that y is not the only vertex in C, and let $W = U \setminus \{u\}$. Then $\langle W \rangle$ is still disconnected, but x is no longer full in $V - W$ since a non-neighbor y has "moved in." Of course y is not full either, and any other vertex which is full in $V - W$ must already have been full in $V - U$. Hence $\omega(W) < \omega(U)$, a contradiction.

We are reduced to the case where y is an isolated point of $\langle U \rangle$; now we let $W = U \cup \{x\}$. Since x and y are not adjacent y is still isolated in $\langle W \rangle$. Any full vertex of $V - W$ was adjacent to x in $V - U$ and therefore already full in $V - U$; but x itself is now gone from $V - W$ and so we again have $\omega(W) < \omega(U)$, and this contradiction proves the theorem. □ □

Note that the induction hypothesis, and thus the theorem itself, can be strengthened to read "each $\langle V_i \rangle$ has at least two vertices and contains an isolated point" without changing the proof. However, we will not need the stronger statement.

Theorem 4.6. *The following are equivalent for a bigraph H.*

(a) *H is separable in one round;*

(b) *H is deterministically separable in one round;*

(c) *for every vertex u of $A(H)$ there is a vertex v of $A(H)$ such that the neighborhoods of u and v (in $B(H)$) are disjoint.*

Proof. It is enough to show (a)→(c)→(b). Suppose that H is separable in one round, let u be any vertex of $A(H)$, and let i be a message that can be sent by Alice when her end of the true edge is u. Since the bigraph that results from sending "i" is disconnected, there is a vertex v on Alice's side of it which is in a different component from u; then u and v must originally have had disjoint neighborhoods.

Now suppose that (c) is satisfied and form a graph G on the vertices in $A(H)$ by defining $\{u, v\}$ to be an edge whenever u and v have *intersecting* neighborhoods. Condition (c) says precisely that G has no vertex of full degree, hence we may apply Theorem 4.5 to obtain a partition $V_1, \ldots, V_k$ of the vertices of $A(H)$ each part of which induces a disconnected subgraph of G, hence also of H. Sending "i" when Alice's end of the true edge lies in V_i thus yields a deterministic separation protocol. □ □

We have said that a disconnected bigraph is sufficient to enable Alice and Bob to communicate a bit in secret; we are now in a position to show that separability is in fact *necessary* for such a communication, thus completing the reduction of the original cryptologic problem to a graph-theoretical one.

Let us fix a bigraph H and suppose that Alice (say) has been supplied with a bit ε which she must communicate in secret to Bob, over our usual public channel. The effect of the bit is to double the vertices on Alice's side of H; that is, each vertex $a \in A(H)$ now becomes a pair $a(0), a(1)$ each with the same neighborhood that a had in $B(H)$. The edge $(a(0), b)$ corresponds to the original (a, b) together with the statement "$\varepsilon = 0$".

At the conclusion of a successful, non-randomized communication protocol the question "What is the value of ε?" must be answerable by Bob but not by Eve, hence the bigraph must now be disconnected—and moreover (although we shall not need this fact) the vertices from $A(H)$ in each component must either all correspond to $\varepsilon = 0$ or all to $\varepsilon = 1$.

Theorem 4.7. *Alice can communicate a bit to Bob in secret, via a randomized and/or non-deterministic communication protocol, iff their bigraph is separable.*

Proof. Sufficiency has already in effect been demonstrated; Bob and Alice can cooperatively disconnect the bigraph, ignoring the bit value, then a message of the form "My bit is 0 if the true edge lies in component C, 1 otherwise" does the trick.

For the converse, we double the bigraph as above so that the protocol may be regarded as a special communication protocol, which we denote by P. If P is randomized, then we may replace the random inputs by non-determinism as in Theorem 4.1; thus we may assume P is merely non-deterministic.

Now we construct from P a randomized (!) communication protocol P' which operates on the original, undoubled bigraph plus a single random bit α for Alice. P' operates by the rule $f'_i(x; \alpha; m_1, \ldots, m_{i-1}) = f_i(x(\alpha); m_1, \ldots, m_{i-1})$ for i odd, and $f'_i = f_i$ for i even.

At each stage, the bigraph associated with P' will be precisely the image of the bigraph associated with P under the collapsing map Φ which sends $x(0)$ and $x(1)$ to x. But the image of a disconnected bigraph under this mapping is again disconnected, so P' is a separation protocol for the original bigraph, completing the proof of the theorem. □ □

Theorem 4.7 says, in effect, that if Alice and Bob know that they will need to communicate a bit in secret, then they can disconnect their bigraph *in advance*; when the bit comes in, it can then be communicated (in either direction) by a single message.

5 The 'Two Sheriffs' Problem

Let us now look now at the two sheriffs problem, but generalized as follows: one sheriff (whom we shall call "Lew") has narrowed his list of suspects to p, the other ("Ralph") to q, and the total number of suspects is n. The edges of the bigraph H here represent all possible pairs (L, R) of subsets of the set $N = \{1, 2, \ldots, n\}$ of suspects, with $|L| = p$, $|R| = q$ and $|L \cap R| \neq \varnothing$. Lew's side $A(H)$ of the bigraph will thus contain $\binom{n}{p}$ vertices and Ralph's side $\binom{n}{q}$ vertices, adjacency arising when the corresponding subsets intersect.

If Lew and Ralph succeed in determining the identity of the killer without tipping off the mob, they will share a secret and thus must have disconnected H. Conversely, suppose they manage to disconnect the bigraph; then Lew and Ralph can reduce further to two non-adjacent edges, one of which is the true edge. If that edge represents an overlap of one, the sheriffs will have found the killer.

Theorem 5.1. *If $n = 2pq$ then there is a deterministic separation protocol for solving the two sheriffs problem.*

Proof. We make use of Baranyai's Theorem [1], which says the following:

If k divides n then there is an array $\{K_{i,j}\}$, $1 \le i \le rn/k$, $1 \le j \le \binom{n}{k}/(n/k)$ of subsets of $N = \{1, 2, \ldots, n\}$ such that each $|K_{i,j}| = k$, each column $K_{1,j}, \ldots, K_{n/k,j}$ is a partition of N, and each subset of N of size k appears exactly once in the array.

Such an array (known as a 1-factorization of the complete k-uniform hypergraph on n vertices) is fixed by Lew and Ralph (publicly) for $k = p$, and Lew proceeds to tell Ralph on which column his end L of the true edge can be found.

Ralph is thus presented with a partition $L_1, \ldots, L_{n/p}$ of N, one of whose parts is Lew's narrowed-down suspect set. His job will be to split the index set $I = \{1, 2, \ldots, n/p = 2q\}$ into two parts, say I_1 and I_2, so that his suspect set R is contained in $\bigcup_{i \in I_j} L_i$. This will disconnect the bigraph.

To do this the sheriffs employ a fixed (but arbitrary) map Φ from the set of subsets of I of size at most q to the set of subsets of I of size *exactly* q, such that $\Phi(S) \cap S = \varnothing$ for every set S in the domain of Φ. Ralph forms the set $F := \{i :\ R \cap L_i \neq \varnothing\}$, then puts $F' := \Phi(F)$ and $F'' := \Phi(F')$. F' and F'' are thus complementary subsets of I of cardinality q; let I_1 be the one containing the element "1" of I, and let I_2 be the other. Ralph now identifies I_1 and I_2 and announces that his set F, defined as above, is contained either in I_1 or I_2.

The resulting bigraph will contain all vertices of $B(H)$ for which the resulting F would have been contained in I_1 or I_2 *and* would have had cardinality q, since in those cases F' and F'' are not dependent on the choice of Φ. Those vertices for which F is contained in I_e will form a connected component, for $e = 1, 2$.

Let k be the index of Lew's end of the true edge, that is, let L_k be Lew's suspect set; suppose $k \in I_e$. Let i be such that k is the ith smallest member of I_e, and let j be the ith smallest element of I_{3-e}. Lew now announces that his set of suspects is in fact either L_j or L_k.

Ralph (but not the mob) will know which of the two is Lew's suspect set; say it is L_j. If $|R \cap L_j| > 1$ then Ralph announces that the killer cannot be identified; otherwise, however, he now knows the killer (say, x). Choosing (again by order of numbers) the corresponding element y of L_k, he announces that the killer is one of x and y. This completes the protocol. □ □

Let us see how this works in the original case $p = q = 2$, $n = 8$. The following Baranyai array can be used:

$\{1,2\}$	$\{1,3\}$	$\{1,4\}$	$\{1,5\}$	$\{1,6\}$	$\{1,7\}$	$\{1,8\}$
$\{3,4\}$	$\{2,4\}$	$\{2,3\}$	$\{2,6\}$	$\{2,5\}$	$\{2,8\}$	$\{2,7\}$
$\{5,6\}$	$\{5,7\}$	$\{5,8\}$	$\{3,7\}$	$\{3,8\}$	$\{3,5\}$	$\{3,6\}$
$\{7,8\}$	$\{6,8\}$	$\{6,7\}$	$\{4,8\}$	$\{4,7\}$	$\{4,6\}$	$\{4,5\}$

Suppose that the true edge is either ({1,2},{1,3}) or ({5,6},{5,7}). Then Lew will announce that his suspect set belongs to the first column, that is, is one of {1,2}, {3,4}, {5,6} or {7,8}. If Ralph himself had one of these sets he would simply announce at this point that the killer cannot be identified; as it is, he splits the index set, telling Lew that his suspect set is either contained in $\{1,2\} \cup \{3,4\}$ or in $\{5,6\} \cup \{7,8\}$. Lew now says "My set is either {1,2} or {5,6}" and Ralph comes back with "The killer is either 1 or 5".

Non-deterministic versions of the above protocol are more easily described; Lew merely picks some partition of which his suspect set is a part, and Ralph can reduce to two possible suspect sets whose intersections with the partition indices are complementary. Here just one more message, from Lew to Ralph, completes the protocol. Moreover, this can be made to work for any $n > 2pq$ as well.

However, we can do even better when non-determinism is permitted; for example, here is a non-deterministic separation protocol for solving the case where $n = k^2$, $p = (k-1)^2 + 1$ and $q = 1$, for any $k \geq 2$.

Lew begins by choosing a $k \times k$ array $\{s_{i,j}\}$ of all the suspects, such that for some j', Lew's suspect set consists precisely of $s_{i',1}$ and all $s_{i,j}$ such that $i \neq i'$ and $j \neq 1$. Ralph, who began knowing the identity of the killer, replies as follows: if the killer is $s_{i,j}$ for $j \neq 1$ and any i, he says "The killer is either $s_{1,j}$ or $s_{i,j}$." If the killer is some $s_{i,1}$ then he picks any $j \neq 1$ and makes the same statement.

By first partitioning the suspects into possible vertices (as in the p, q, $2pq$ case) and then making a $k \times k$ array of the sort described above, but where the array elements are members of a partition instead of single suspects, we may combine the techniques for the following result:

Theorem 5.2. *The two sheriffs problem is solvable non-deterministically whenever* $n \geq q(1 + \sqrt{p-1})^2$.

It is perhaps interesting to note that we have separated a very dense bigraph here, regular on each side. In fact, related to these are the following dense bigraphs, which are *deterministically* separable: fix a large k and set H equal to

$$\{((a,b),(c,d)) : 1 \leq a,b,c,d \leq k \ \text{ and either } \ a = c \text{ and } b = d, \text{ or } a \neq c \text{ and } b \neq d\}\ .$$

To separate H, Alice's announces the first coordinate of her pair and Bob the second coordinate of his. Then each will know whether their edge is based on the equalities or the inequalities in the definition above.

6 Multi-Party Generalization

It is evident that many of our definitions and results can be extended to the case where there are more than two conversants. In this case conversation protocols, in order to remain general, allow the identity of the next speaker to depend, at each turn, on the previous conversation.

In [3] Fischer and Wright suggest using random deals to facilitate secret key exchange within a group of persons wishing to communicate privately in yet-to-be-specified subgroups. Among the negative results in [3] is a theorem (Theorem 9, p. 11) which states that no communication protocol for 3 players, each dealt one card of a 3-card deck, can enable them to isolate a secret bit. Fischer and Wright indicate that our methods can be used to generalize the result to $k > 3$ players; we show here how that can be done.

The definition of "bigraph"' extends easily as follows: a *k-graph* is a finite collection H of k-tuples $x = (x_1, \ldots, x_k)$ (which we call "blocks" to avoid confusion) such that $x, y \in H$ implies that x_i and y_j are distinct for $i \neq j$. A k-graph is thus a particular special case of k-uniform hypergraph in which the sets $\{x_i : x \in H\}$ partition the vertex set of H.

The proof of Theorem 2.2 goes through, as does an appropriate version of Theorem 4.2; thus we are once more reduced to showing that the players, say X_1 through X_k, cannot cooperatively disconnect their k-graph H_k which in this case consists of a block for each permutation of the cards.

Theorem 6.1. *The "permutation k-graph" H_k is inseparable for $k \geq 3$.*

Proof. H_2 is of course separable, indeed disconnected to begin with. Let us assume that P is a (non-deterministic) separation protocol for H_k, for some $k > 2$, and let $H^0, H^1, \ldots, H^t$ be the state of the k-graph for $X_1, \ldots, X_k$ at each stage of some (successful) conversation using P. Then $H^0 = H_k$ and $H^0, \ldots, H^{t-1}$ are connected k-graphs. We claim first that H^t must consist only of two blocks, which up to permutation of the players and cards, may as well be $1, 2, 3, \ldots, k$ and $2, 3, \ldots, k, 1$. To see this let x and y be two blocks of H^t which lie in different components; then in particular x and y are not adjacent so $x_i \neq y_i$ for $i = 1, 2, \ldots, k$. Let G be the graph on vertices 1,2 , ... , k with j adjacent to j' when $\{j, j'\} = \{x_i, y_i\}$ for some i; then since x and y are each permutations, G is regular of degree 2. If ϕ is an automorphism of G such that $\phi(j)$ is adjacent to j for all j then $(\phi(x_1), \ldots, \phi(x_k))$ is a block of H^k which is adjacent to x if ϕ has any fixed points and to y if ϕ is not the identity. Since no block of H^t can be adjacent to both x and y, every such ϕ must fix all vertices or none; hence G consists of a single cycle. By relabelling we may assume $x = (1, 2, 3, \ldots, k)$ and $y = (2, 3, \ldots, k, 1)$.

Now if H^t contains any other block some player, say X_k, must have another vertex, say $j \neq k, 1$. But then the block $(1, 2, \ldots, j-2, j-1, j+1, j+2, \ldots, k-1, k, j)$ lies in H^t; and it is adjacent to both x (at player X_1) and y (at player X_k), a contradiction. This proves the claim.

We may now assume that H^t consists exactly of the above blocks x and y, and that the last player to speak was X_1; then the vertices of H^{t-1} are exactly those appearing in x and y, plus some additional vertices held by X_1 which she eliminated in her last message. The fact that H^{t-1} contains those additional vertices means that if one of them had been X_1's "true" vertex, the conversation might have gone exactly as it did until the last message. But H^{t-1} contains no pair of non-adjacent blocks other than x and y, since in every block not equal to x or y, player X_2 holds a 2 and player X_k holds a 1. Hence, the protocol P has failed in this case and this contradiction proves the theorem. □ □

7 Final Comments

In this work we have only begun to study the combinatorial cryptology of isolating a common secret; unsettled questions abound. For example: Are there separable bigraphs requiring arbitrarily many rounds in order to separate them? Is there a good way to characterize those bigraphs that are separable in n rounds? What about the bigraphs that are deterministically separable in n rounds? Is it possible to prove that the two sheriffs problem cannot be solved deterministically when $n < 2pq$?

We hope that our "bigraphs" may prove to be a useful way of representing common knowledge, even for applications unrelated to the problem of isolating a common secret. Although they carry the same information as do other representations, they may help attract graph-theorists to knowledge problems and thus bring some powerful theorems and sharp combinatorial minds to bear.

References

[1] Zs. Baranyai, On the factorization of the complete uniform hypergraph, *Infinite and finite sets (Keszthely, 1973; dedicated to P. Erdős on his 60th birthday)*, Vol. I (1975) 91–108.

[2] M. Fischer, M. Paterson, and C. Rackoff, Secret bit transmission using a random deal of cards, *Distributed Computing and Cryptography*, American Mathematical Society (1991) 173–181.

[3] M. Fischer and R. Wright, Multiparty secret key exchange using a random deal of cards, *Proceedings of CRYPTO '91*, Springer-Verlag Lecture Notes in Computer Science, Vol 576 (1992) 141–155.

[4] P. Winkler, Cryptologic techniques in bidding and defense (Parts I, II, III, IV), *Bridge Magazine* (April-July 1981) 148–149, 186–187, 226–227, 12–13.

[5] P. Winkler, The advent of cryptology in the game of bridge, *Cryptologia*, vol. 7 (1983) 327–332.