Achieving Perfect Secrecy Using Correlated Random Variables

Rebecca N. Wright

YALEU/CSD/RR #1058

January 1995

ABSTRACT

Achieving Perfect Secrecy Using Correlated Random Variables

Rebecca N. Wright
Yale University
1994

We present a general model for communication among a "team" of players, P_1 through P_k, overheard by a passive eavesdropper, Eve, in which all players including Eve are given private inputs that may be correlated. We define and explore secret key exchange in this model. A *protocol* specifies how the team players behave. A *source* specifies a distribution on the team players' inputs. Together, a protocol and a source are a *system*. An adequate notion of secrecy must depend not only on the protocol, but on any relevant knowledge Eve may have about the source. We provide precise mathematical definitions of the multiparty secret key exchange problem that take Eve's knowledge into account, and allow us to determine whether a given system achieves secret key exchange with respect to any particular "type of knowledge" for Eve. Our secrecy requirements are information-theoretic and hold even if Eve is computationally unlimited.

We explore the properties of secret key exchange systems and show formally the intuitive result that the team players' inputs must be correlated in order for secret key exchange to be achieved. We show that any source has a finite *capacity* for secret key exchange even if the team players are allowed unlimited independent randomization. Specifically, the capacity of a source is the largest number of output values of any secret key exchange system using that source. We present an upper bound on the capacity of arbitrary sources.

In the second half of the dissertation, we focus our attention on sources yielding inputs consisting of hands of prespecified sizes from a randomly shuffled deck of cards. We exhibit several secret key exchange protocols designed to work even if Eve is allowed to see the remaining cards in the deck. Let each team player P_i hold s_i cards, let Eve hold s_e cards, and let k be the number of team players. We present a family of *key set* protocols, a simple family of protocols that achieve one-bit secret key exchange provided that $\max s_i + \min s_i \geq k + s_e$. We use a game-theoretic argument

to exhibit an optimal key set protocol.

Of particular interest is the symmetric case in which each player receives the same fraction β of the cards in the deck. The efficiency of a secret key exchange protocol is measured by the smallest deck size d_0 for which the protocol is guaranteed of success. The best previous bound (Fischer, Paterson, and Rackoff, 1991) was super-polynomial in $1/\beta$ and only handled the special case of $k = 2$ and $n = 1$. We present the *transformation* protocol, which achieves n-bit secret key exchange in the symmetric case for any deck size $d > d_0 = O(n(1/\beta)^{2.71})$. Using the transformation protocol, we show that if $d \geq (42.8) \cdot (1/\beta)^{2.71}(n + 0.07)$, then the capacity of a source in which each team player receives βd cards is at least 2^n.

Finally, we show that secret key exchange is not possible when $k \geq 3$ team players each hold one card from a deck of k cards by showing that the capacity of such a source is one. This is the first result showing that even if the team holds all the cards and no player's hand is empty, secret key exchange is not always possible.

Achieving Perfect Secrecy Using Correlated Random Variables

A Dissertation
Presented to the Faculty of the Graduate School
of
Yale University
in Candidacy for the Degree of
Doctor of Philosophy

by

Rebecca N. Wright

December 1994

For my mother, Barbara Wright

ACKNOWLEDGEMENTS

There are many people that have contributed directly or indirectly to the completion of this work. First and foremost is Joan Feigenbaum, who has acted as a mentor to me since I met her during my senior year of college. She encouraged me to go to graduate school, introduced me to research, and provided unlimited good advice throughout the years.

I would like to thank my advisor Michael Fischer, whose relentless search for perfection has added greatly to the quality of this dissertation. Without him, it might have been completed sooner, but would not have been as good. I thank the other members of my committee, Dana Angluin, Richard Beigel, and Peter Winkler, for their comments and encouragement on earlier versions of this manuscript. In particular, I thank Dana Angluin for consistently acting above and beyond the call of duty.

It is a pleasure to thank Nick Reingold for providing a simpler proof of an earlier version of Lemma 4.11, Richard Beigel for his suggestion leading to the development of the transformation protocol, and Michael Merritt for the idea of dividing hands used in Theorem 7.15. Much of the work described in this dissertation was completed jointly with Michael Fischer.

In addition to the technical support that has been provided by my colleagues, my friends and family have provided much needed emotional support and recreational diversion. Among these are Ramón Cáceres, Beth Multer, Tama Hasson, Jeff Westbrook, Dan Rabin, Nick Reingold, Lee Millward, Michael Black, Barbara Wright, Caryl Pofcher, Kira Enge, Christine Bruckner, Ann Cowlin, Yaledancers, the Yale Slavic Chorus, everyone at CuraGen Corporation, and, of course, my cat Picasso.

Finally, I would like to thank Koffee? for providing good coffee and a comfortable, friendly place to work whenever I needed to get out of the office for a few hours.

CONTENTS

LIST OF FIGURES

CHAPTER 1: INTRODUCTION

We present a general model for communication among a "team" of players, P_1 through P_k, overheard by a passive eavesdropper, Eve, in which all players including Eve are given private inputs that may be correlated. We define and explore the problem of secret key exchange in this model. Our secrecy requirements are information-theoretic and hold even if Eve is computationally unlimited. As a specific instance of correlated random variables, we consider the situation in which the team players are dealt hands of cards of prespecified sizes from a known deck of distinct cards.

1.1 Example 1

Consider a scenario in which Alice, Bob, and a computationally-unlimited eavesdropper, Eve, are playing a game of cards with a deck of four cards, J, Q, K, and A. Alice is given two cards, Bob is given one card, and the remaining card may or may not be given to Eve. Can Alice and Bob communicate publicly to agree on a bit that is secret from Eve? The answer to this question depends on what is meant by "secret," whether Eve is allowed to look at the remaining card, and whether Alice and Bob are allowed to use randomization.

If Eve is not allowed to look at the remaining card, the following deterministic protocol $\mathcal{P}_1$ achieves *perfect* secret bit exchange, in which Eve considers both values equally likely at the end of every run of the protocol. Alice says whichever of the following statements is true:

- "I hold $\{\mathrm{J}, \mathrm{Q}\}$ or $\{\mathrm{K}, \mathrm{A}\}$."
- "I hold $\{\mathrm{J}, \mathrm{K}\}$ or $\{\mathrm{Q}, \mathrm{A}\}$."
- "I hold $\{\mathrm{J}, \mathrm{A}\}$ or $\{\mathrm{Q}, \mathrm{K}\}$."

From Alice's message and Bob's hand, Bob can determine Alice's hand, so Alice and Bob both know the truth value of the statement "Alice holds J," and can agree on this value as their secret bit. On the other hand, Eve considers it equally likely that this statement is true or false, so does not learn anything about the secret bit.

If Eve is allowed to look at the remaining card, the following randomized protocol $\mathcal{P}_1'$ achieves perfect secret bit exchange. First, Alice randomly chooses a card x in her hand and a card y not in her hand, and asks Bob "Do you hold one of the two cards

$\{x, y\}$?" If Bob says yes, then Alice and Bob know which of them holds x and which holds y, but Eve considers both situations equally likely. Hence, Alice and Bob can agree on a secret bit, for example, by the truth value of the statement "Alice holds the smaller-valued card of x and y." On the other hand, if Bob says no, then Alice and Bob each hold one of the two remaining cards z and q and can determine the truth value of the statement "Alice holds the smaller-valued card of z and q," again obtaining the desired secret bit.

Finally, if Eve is allowed to look at the remaining card and Alice and Bob are not allowed to use randomization, then even *weak* secret bit exchange, in which we require only that Eve consider both values possible at the end of each run of the protocol, is not possible. Note that if Eve learns either player's hand, then she can learn the secret bit by simulating that player. It can be shown that if either player sends a message that depends on his or her hand, then it will sometimes be possible for Eve to learn that player's hand. Hence, in order to avoid the possibility of Eve learning the secret bit, neither player can ever send a message that depends on his or her hand. It follows that there is only one possible sequence of messages, and hence each player's output is a function only of his or her hand. Thus, for example, if Bob outputs v when he holds J, Alice must output v whenever she does not hold J in order to guarantee that their outputs always agree. In this case, if Eve holds J, then Eve knows Alice does not hold J and will output v, so she has learned the secret bit. A formal proof of this result involves a detailed case analysis.

To summarize this example, if Eve does not see the remaining card or if Alice and Bob can use randomization, then Alice and Bob can agree on a perfectly secret bit. If Eve sees the remaining card or Alice and Bob are required to behave deterministically, then Alice and Bob cannot agree even on a weakly secret bit. These results are proved formally in §5.3.

The arguments for the correctness of protocols and the nonexistence of protocols presented in the above example are informal and rely on intuition. However, intuition may be misleading when dealing with issues such as secrecy and shared knowledge, so formal definitions of secret key exchange are needed. To this end, we formalize a model for communication among a "team" of players given possibly correlated inputs that are drawn from some known joint distribution. A passive eavesdropper, Eve, hears the communication between the players and is also given an input that may be correlated to the inputs of the players.

1.2 Related Work

The scenario we consider, in which one or more players are given correlated inputs and wish to communicate to achieve some task, arises frequently in the theory of communication and secret communication. We describe some related work from these areas here.

Information Theory

Classical information theory, developed by Shannon [Sha48] in the early 1940's (see also [Gal68]), considers a noninteractive setting in which a player Alice knows a string a and wants to communicate to Bob a related value σ in such a way that Bob can recover a from σ, but the length of σ is minimized. Considering perfect secrecy in this model [Sha49], Alice and Bob share a key K and Eve is a passive computationally-unlimited eavesdropper. Alice knows a and K, Bob knows K, and both Bob and Eve hear all communication sent by Alice. Alice wishes to send a message σ (dependent on both a and K) such that Bob can recover a exactly but Eve learns nothing about a.

Interactive Communication

Since Shannon's seminal work on the theory of one-way communication, much work has been done in the area of interactive communication, where Alice and Bob communicate messages back and forth. Communication complexity ([Abe78,Yao78], see also [Lov90]) considers a scenario where Alice knows a value a and Bob knows a value b, where a and b are drawn from some known distribution. Alice and Bob wish to communicate a minimal number of bits in order to compute the value of some known function f given arguments a and b. Orlitsky [Orl90,Orl91] and Naor, Orlitsky, and Shor [NOS93] consider the same setting (Alice knows a and Bob knows b, where a and b are jointly distributed) and ask how many bits Alice and Bob must send in order for Bob to learn a if only m messages (i.e. $m/2$ rounds of interaction) are allowed.

Maurer [Mau93] and Bennett, Brassard, and Roberts [BBR88] consider the problem of almost perfect secrecy in the setting of interactive communication. In their work, Alice knows a, Bob knows b, and a passive computationally-unlimited eavesdropper Eve knows e, where a, b, and e are jointly distributed. The question asked is whether Alice and Bob can communicate interactively to determine a value y such that Alice and Bob both learn y with high probability and Eve learns at most a negligible amount of information about y. Maurer considers a general framework in which a, b, and e can be arbitrary correlated random variables, while Bennett, Brassard, and Roberts consider very specific kinds of a, b, and e motivated by quantum cryptography (see below).

Fischer, Paterson, and Rackoff [FPR91] consider the problem of completely perfect secrecy in the setting of interactive communication. In their work, as above, Alice knows a, Bob knows b, and Eve knows e, but in their case a, b and e are the hands of a deal of cards. They ask whether Alice and Bob can communicate interactively to determine a bit y such that Alice and Bob are guaranteed to learn y and Eve learns nothing about the value of y.

Beaver, Haber, and Winkler [BHW93] consider a weaker form of secrecy. In their setting, Alice knows a and Bob knows b, Eve does not get a correlated input, and the question asked is whether Alice and Bob can determine a bit y such that Alice and Bob are guaranteed to learn y and Eve does not learn y with certainty. They consider a and b being hands of a deal of cards as well as other specific kinds of correlated random variables.

The areas of quantum cryptography and public key cryptography consider the use of interactive communication to achieve secrecy against different kinds of eavesdroppers. Both quantum cryptography and public key cryptography allow two parties who share no secret information to communicate secretly, but both rely on unproven assumptions. Quantum cryptography [BB85,BBE92] considers an *active* computationally-unlimited eavesdropper and relies on assumptions about the physical world such as Heisenberg's uncertainty principle which says that measuring a quantum system both gives incomplete information about the current state of the system and changes the current state. Based on such assumptions, it is possible to use quantum techniques to construct an imperfect private channel, i.e. a channel such that Eve can learn at most a partial amount of information about the messages sent, but Eve can also tamper with these messages in certain ways such that Alice and Bob may not exactly agree on a secret. Combining the imperfect private channel with some of the ideas described in the previous paragraphs (in particular the work of Bennett, Brassard, and Roberts [BBR88]), it is possible for Alice and Bob to use a public tamperproof channel to amplify the privacy of the imperfect channel to obtain a perfectly secret or almost perfectly secret key or, if they do not succeed, to learn that they have not succeeded.

Public key cryptography [DH76,Mer78,RSA78,ElG85], on the other hand, is designed to work against a passive eavesdropper who has only limited (typically polynomial) computing power available to her. Under the unproven but widely believed assumption that computational problems such as factoring large numbers and computing discrete logarithms are computationally difficult, public key cryptography allows two players with no prior secret information to exchange a key that Eve can not learn. However, there are several problems with public key cryptography. First, even if, for example, one-way permutations are assumed to exist, this may not be useful, for Impagliazzo and Rudich [IR89] provide evidence that most of the standard techniques in cryptography cannot be used to construct a secret key exchange protocol from a one-way permutation. Second, public key cryptography is based on assumptions about the worst-case asymptotic computational complexity of certain problems. Even if public key cryptography is based on a problem that is actually asymptotically hard, it is not at all clear how to choose a key size or an actual key in order to get the desired security.

Somewhat related to public key cryptography is the concept of zero-knowledge interactive proof systems [GMR89,GMW86], in which a computationally-unlimited prover and a computationally-limited verifier communicate. The goal of a perfect zero-knowledge interactive proof system is that the prover should be able to convince the verifier that an element x belongs to a language L without the verifier learning any additional information. In some sense, the verifier plays the role of an active computationally-limited eavesdropper in this setting.

Multiparty Communication

In the setting of distributed computing, many different processes, possibly residing on many different computers, interact to carry out a given task. A problem that has been extensively studied in a distributed setting is the problem of reaching agreement in the presence of faults. Variants of this problem include Byzantine agreement and distributed consensus. In both cases, a number of processes communicate to try to agree on a value. In the Byzantine agreement problem, the processes communicate in a synchronous fashion, and some processes may be faulty i.e. they are allowed to behave arbitrarily to try to prevent the nonfaulty processes from reaching agreement. Lamport, Shostak, and Pease [LSP82] show that agreement is reachable if and only if less than one third of the processes are faulty. In the distributed consensus problem, communication is asynchronous, and a faulty process is one that stops before the protocol is complete. Fischer, Lynch, and Paterson [FLP85] show that in this setting, agreement is not reachable if even one process is faulty.

In multiprover interactive proof systems, introduced by Ben-Or, Goldwasser, Kilian, and Wigderson [BOGKW88], multiple computationally-unlimited provers interact with a computationally-limited verifier. They show without any unproven assumptions that every language in NP has a perfect zero-knowledge multiprover interactive proof system. In sharp contrast to this, it is unlikely that perfect zero-knowledge single-prover interactive proof systems (as discussed above) exist for every language in NP, since in this case the polynomial hierarchy collapses at the second level [BOGKW88].

Yao [Yao82] considers the problem of secure multiparty function computation. Players $P_1, \ldots, P_k$ are given private inputs $a_1, \ldots a_k$, respectively. As in the case of communication complexity, the players wish to evaluate some known function f on their inputs. All communication between the players is public. In addition, the players wish to preserve maximum privacy of the a_i's (some information about the a_i's will be implied by the value $f(a_1, \ldots, a_k)$. Yao shows that any function than can be computed without the privacy requirement can also be computed with the privacy requirement, but there exist functions that require exponentially many bits be sent in order to satisfy the privacy constraint. Chaum, Crépeau, and Damgård [CCD88] consider a slightly different model in which each pair of players can communicate secretly but some dishonest players may deviate from the protocol. They show that in this model secure multiparty function computation is possible if and only if at least 2/3 of the players are honest.

1.3 Our Work

Our work combines several aspects of the work described above. We consider a multiparty setting in which players $P_1, \ldots, P_k$ are given private inputs $a_1, \ldots a_k$. A passive computationally-unlimited eavesdropper Eve is given a private input a_e. The inputs $a_1, \ldots, a_k, a_e$ are drawn from a known distribution. In this setting, we define and explore the problem of *multiparty secret key exchange*. A possible application of this problem is to consider as private inputs an imperfect secret key resulting from methods of quantum cryptography. By using these inputs to achieve multiparty secret key exchange, it would be possible to amplify the privacy of the imperfectly secret key to obtain a shorter, but perfectly secret, key (as is done by Bennett, Brassard, and Roberts [BBR88] in the two-party case).

We define our model in Chapter 2. In Chapter 3, we define the secret key exchange problem. Briefly, the problem of N-valued multiparty secret key exchange is for a team of players to choose a value v from a known set of N values. After the players communicate, each team player must know v, but v must be unknown to Eve. We require that our protocols always succeed, not just succeed with high probability. An adequate notion of secrecy must depend not only on the protocol being executed, but on any relevant knowledge Eve may have that is external to the protocol. We provide precise mathematical definitions of the multiparty secret key exchange problem that take Eve's knowledge into account, and allow us to determine whether a given protocol achieves secret key exchange with respect to any particular "type of knowledge" for Eve.

We consider two kinds of secret key exchange, *perfect* and *weak*. N-valued perfect secret key exchange requires that Eve consider all N possible values for v equally likely, while N-valued weak secret key exchange requires only that Eve consider all N values for v possible. The work of [FPR91] can be formalized according to our definition of 2-valued perfect secret key exchange. Two-valued weak secret key exchange is equivalent to the concept of "sharing a secret" used in [BHW93].

In Chapter 3, we investigate the properties of secret key exchange protocols. It follows easily from our definitions that perfect secrecy implies weak secrecy and that giving Eve external knowledge can only help her. We explore how the requirements of a secret key exchange protocol restrict the inputs and behavior of the players. In particular, we show that secret key exchange is not possible if the players' inputs are not correlated.

Our model is quite general and admits on the extreme ends the case where the players' inputs are equal and the case where the players' inputs are independent. In the first case, the players can use their inputs as a secret key. We show in §3.6 that secret key exchange is not possible in the second case. Chapter 4 defines the perfect and weak secret key exchange capacity of a source of correlated random variables.

We are interested in the borderline between possibility and impossibility of secret key exchange. In order to approach this borderline, we focus our attention in Chapters 5–8 on inputs consisting of hands of prespecified sizes from a randomly shuffled deck of cards. A random deal of cards is an example of sampling without replacement. By looking at her own cards, a player gains some information about the other players' hands. Namely, she learns a set of cards that appear in no other player's hand. Peter Winkler developed bidding conventions for the game of bridge whereby one player could send her partner secret information about her hand that was totally unrelated to the actual bid and completely indecipherable to the opponents, even though the protocol was known to them [Fli81,Win81a,Win81b,Win83]. Fischer, Paterson and Rackoff [FPR91] and Beaver, Haber and Winkler [BHW93] carried this idea further, using deals of cards for secret bit transmission between two players.

We consider the use of a random deal of cards for multiparty, multivalued, secret key exchange. A *signature* $(s_1, \ldots, s_k; d)$ specifies the hand size s_i for each player and the deck size d. The perfect (respectively weak) *capacity* of a signature is the largest N such that N-valued perfect (weak) secret key exchange is possible when the deal is chosen randomly as specified by the given signature. We investigate bounds on the capacity of different signatures.

Note that if $N = 1$ or if there is only one team player, N-valued secret key exchange is not of interest. Beaver, Haber and Winkler [BHW93] and Fischer, Paterson and Rackoff [FPR91] have studied the case where $N = 2$ and there are only two team players. In Chapter 6, we present the class of key set protocols, a simple class of protocols that perform perfect secret bit exchange for any signature $(s_1, \ldots, s_k; d)$ such that $\max s_i + \min s_i \geq k + d - \sum s_i$, even if Eve is allowed to look at the remaining cards. We use a game-theoretic argument to exhibit an optimal key set protocol. In Chapter 7, we present the transformation protocol, a two-player protocol that performs n-bit perfect secret key exchange when each of the team players holds a fixed fraction of the cards, provided that the deck is large enough. We also show how to use the transformation protocol to perform multiparty secret key exchange.

In Chapter 8, we show that the capacity of $(1, \ldots, 1; k)$ is 1 if $k \geq 3$ by showing that 2-valued secret key exchange is not possible when k team players each hold one card from a deck of k cards. This is the first result showing that even if the team holds all the cards and no player's hand is empty, secret key exchange is not always possible. (By way of contrast, it follows from Lemmas 6.9, 6.8, and 6.16 that 2-valued secret key exchange is always possible if each of k players holds 2 cards from a deck of $2k$ cards.)

Much of the work described in this dissertation was completed jointly with Michael Fischer, and also appears in [FW92,FW93a,FW93b,FW].

CHAPTER 2: MULTIPARTY PROTOCOLS AND SYSTEMS

We consider a *team* of players P_1 through P_k. We denote the number of team players by k throughout the paper, and denote the set $\{1, \ldots, k\}$ by K. We will frequently use k-tuples to describe a collection of items, one for each player. Given any k-tuple x, we denote the ith component (P_i's component) of x by $\langle x \rangle_i$.

We use a synchronous distributed model of computation in which communication occurs in rounds. All protocols terminate in a fixed finite number of rounds. Each player is given a private input that will generally be chosen at random before the protocol begins. In a round, each of the players simultaneously broadcasts a message to all of the other players. The message sent by P_i at a given round depends on P_i's input and the messages sent by the team in previous rounds. On termination, each player P_i produces a private output that depends on her own input and the messages sent by the team in all rounds. We define protocols formally in §2.1.

We describe how the players' random inputs are generated in §2.3. The information given to Eve is formalized in §2.4. Briefly, a global value is chosen randomly according to a known, fixed distribution. Each player has a view of the global value that constitutes her private input. Eve also gets a view of the global value. In general, the global value and the players' views will be chosen so that the players' inputs consist of correlated values (such as the hands of a deal of cards) as well as independent random values that play the role of private coin flips or dice rolls for the individual players. A simple protocol $\mathcal{P}_2$ is given in §2.2. An example of random inputs to $\mathcal{P}_2$ is given in §2.6.

2.1 Protocols

Formally, a *protocol* is a 7-tuple $\mathcal{P} = (k, t, U, V, M, \mu, \nu)$.

- k is the number of players.
- t is the number of rounds.
- $U = U_1 \times \cdots \times U_k$ is the input set.
- V is the output set.
- M is the set of messages.

- μ is a k-tuple $(\mu_1, \ldots, \mu_k)$ of *message functions.*
- ν is a k-tuple $(\nu_1, \ldots, \nu_k)$ of *output functions.*

An element $u \in U$ is called an *input vector* or simply an *input.* The component $\langle u \rangle_i$ is *the input for* P_i. U_i is the input set for P_i. An element $v \in V$ is called an *output.* We say $\mathcal{P}$ is $|V|$-valued.

A k-tuple $m = (c_1, \ldots, c_k) \in M^k$ is called a *message vector.* It represents the messages sent by all team players in a given round. The component $\langle m \rangle_i$ is the message sent by P_i. A sequence of at most t message vectors is called a *conversation.* A conversation τ is *complete* if $|\tau| = t$ and *partial* if $|\tau| < t$. We let $\mathsf{c} = \bigcup_{\ell=0}^{t}(M^k)^\ell$ denote the set of conversations, $\mathsf{cc} = (M^k)^t$ denote the set of complete conversations, and $\mathsf{pc} = \mathsf{c} - \mathsf{cc}$ denote the set of partial conversations. Given a conversation τ, we let τ_j denote the jth message vector of τ, and for $\ell \leq |\tau|$, we write $\tau[\ell] = (\tau_1, \ldots, \tau_\ell)$ to denote the conversation consisting of the length ℓ prefix of τ. We denote the concatenation of a partial conversation τ and a message vector m by $\tau \cdot m$ or τm, so $\tau \cdot m = \tau m = (\tau_1, \ldots, \tau_{|\tau|}, m)$.

The message function $\mu_i : U_i \times \mathsf{pc} \to M$ specifies the messages for P_i to send. Let $a_i \in U_i$, $\ell < t$ and $\tau \in \mathsf{pc}$. The message $\mu_i(a_i, \tau)$ is the message that P_i sends at round $\ell + 1$ when P_i has input a_i and the conversation through round ℓ is τ. The output function $\nu_i : U_i \times \mathsf{cc} \to V$ specifies the output for P_i. Let $a_i \in U_i$ and $\sigma \in \mathsf{cc}$. The value $\nu_i(a_i, \sigma)$ is the output for P_i when P_i has input a_i and the complete conversation is σ.

When the protocol $\mathcal{P}$ is clear from context, we will use the notation k, t, U, U_i, V, M, μ, μ_i, ν, ν_i, c, cc and pc as above. Otherwise, we attach $\mathcal{P}$ as a superscript.

Given an input $u \in U$, we say σ is a *conversation of* u if σ is complete and $\langle \sigma_\ell \rangle_i = \mu_i(\langle u \rangle_i, \sigma[\ell - 1])$ for all $i \in K$ and $1 \leq \ell \leq t$. It is easily seen that for each input u, σ is unique. This is because the messages of the first round of σ are completely determined by u, and successive rounds of σ are determined by u and the previous rounds of σ. We denote this unique conversation σ by $\mathrm{conv}(u)$. We denote P_i's output $\nu_i(\langle u \rangle_i, \mathrm{conv}(u))$ by $\mathrm{out}_i(u)$, and we denote the vector $(\mathrm{out}_1(u), \ldots, \mathrm{out}_k(u))$ of all players' outputs by $\mathrm{out}(u)$.

It is possible to "interpolate" between a number of input vectors to construct a new input vector (the "interpolant") by giving each team player P_i his input from one of the original vectors u_i. Let ℓ be the last round where the conversations of all the original vectors are the same. Because the behavior of the team at each round is completely determined by the input vector and the messages received up to that round, the conversation of the interpolant agrees with the original conversations up to round ℓ. Furthermore, each player P_i sends the same message at round $\ell + 1$ with the interpolant as input as he did with the original vector u_i as input.

We say that two input vectors u and u' *touch at coordinate* i, or simply *touch*, if $\langle u \rangle_i = \langle u' \rangle_i$. If u and u' touch at coordinate i, we write $u =_i u'$. Obviously, $=_i$ is an equivalence relation.

Let $U' \subseteq U$. We say that u is *an interpolant of* U' (also u *interpolates* U') if for every $i \in K$, there is some $u' \in U'$ such that $u =_i u'$. We call any k-tuple $(u_1, \ldots, u_k)$ a U'*-derivation* of u if $u_i \in U'$ and $u =_i u_i$ for all $i \in K$. Note that u interpolates U' if and only if there exists a U'-derivation of u.

Lemma 2.1 (First Interpolation Lemma) *Let $\mathcal{P}$ be a protocol, let $\tau \in \mathsf{c}$, let $\ell = |\tau|$, and let $U' \subseteq \{U : \mathrm{conv}(u)[\ell] = \tau\}$. Let $\widehat{u}$ interpolate U' and let $(u_1, \ldots, u_k)$ be a U'-derivation of $\widehat{u}$. Then*

1. $\mathrm{conv}(\widehat{u})[\ell] = \tau$.
2. *If $\ell < t$, then* $\mathrm{conv}(\widehat{u})_{\ell+1} = (\langle \mathrm{conv}(u_1)_{\ell+1} \rangle_1, \ldots, \langle \mathrm{conv}(u_k)_{\ell+1} \rangle_k)$.
3. *If $\ell = t$, then* $\mathrm{out}(\widehat{u}) = (\mathrm{out}_1(u_1), \ldots, \mathrm{out}_k(u_k))$.

Proof: Suppose the conditions of the lemma.

1. Since $\mathrm{conv}(u)[\ell] = \tau$ for all $u \in U'$, we have $\mu_i(\langle u \rangle_i, \tau[j-1]) = \langle \tau_j \rangle_i$ for all $u \in U'$, $i \in K$, and $1 \leq j \leq \ell$. It follows that $\mathrm{conv}(\widehat{u})[\ell] = \tau$.
2. Suppose $\ell < t$. Then by (1), $\mathrm{conv}(\widehat{u})[\ell] = \tau$. It follows that $\mathrm{conv}(\widehat{u})_{\ell+1} = (\mu_1(\langle u_1 \rangle_1, \tau), \ldots, \mu_k(\langle u_k \rangle_k, \tau)) = (\langle \mathrm{conv}(u_1)_{\ell+1} \rangle_1, \ldots, \langle \mathrm{conv}(u_k)_{\ell+1} \rangle_k)$.
3. Suppose $\ell = t$. Then by (1), $\mathrm{conv}(\widehat{u})[\ell] = \mathrm{conv}(\widehat{u}) = \tau$. Thus $\mathrm{out}(\widehat{u}) = (\nu_1(\langle u_1 \rangle_1, \tau), \ldots, \nu_k(\langle u_k \rangle_k, \tau)) = (\mathrm{out}_1(u_1), \ldots, \mathrm{out}_k(u_k))$. ■

2.2 Example 2

A simple example of a protocol is one in which each of two players is given an input bit, there is no communication, and each player outputs her given input bit. Formally, this protocol is defined as $\mathcal{P}_2 = (2, 0, \{0,1\} \times \{0,1\}, \{0,1\}, \emptyset, \mu, \nu)$, where each μ_i is any function defined over an empty domain and each ν_i is defined by $\nu_i(a_i, \epsilon) = a_i$ for $i \in \{1, 2\}$ and $a_i \in \{0, 1\}$. We will investigate the properties of $\mathcal{P}_2$ more fully in §2.6, §3.5, and §4.2.

2.3 Sources

A protocol specifies the possible inputs to the players and how the players behave given these inputs. We are generally interested in the case that inputs are generated by a random source. Before proceeding, we introduce some definitions and notations from basic probability theory.

Given an arbitrary distribution $\mathcal{G}$ over an arbitrary finite set Γ, we write $\Pr_{\mathcal{G}}(g)$ to denote the probability assigned by $\mathcal{G}$ to $g \in \Gamma$. A subset $X \subseteq \Gamma$ is called an *event*. We write $\Pr_{\mathcal{G}}[X]$ to denote the probability $\sum_{g \in X} \Pr_{\mathcal{G}}(g)$ assigned to X by $\mathcal{G}$. When the distribution $\mathcal{G}$ is clear from context, we omit the subscript and write simply $\Pr(g)$ and $\Pr[X]$. We say an element g is *feasible* if $\Pr(g) > 0$. We say an event X is *feasible* if $\Pr[X] > 0$ or, equivalently, if X contains a feasible element. We write $\text{feas}(X)$ to denote the set $\{g \in X : g \text{ is feasible}\}$.

A function f over domain Γ is called a *random variable*. The event $\{g \in \Gamma : f(g) = x\}$ is denoted by $f^{-1}(x)$. The random variables $f_1, \ldots, f_k$ are *independent* if $\Pr\left[\bigcap_{i=1}^{k} f_i^{-1}(x_i)\right] = \prod_{i=1}^{k} \Pr\left[f_i^{-1}(x_i)\right]$ for all $x_1, \ldots, x_k$. The events $X_1, \ldots, X_k$ are *independent* if the random variables $f_1, \ldots, f_k$ defined by $f_i(g) = 1$ if $g \in X_i$ and $f_i(g) = 0$ if $g \notin X_i$ are independent. We say an event X *respects* a random variable f if $f(g_1) = f(g_2)$ implies that g_1 and g_2 are either both elements of X or both not elements of X. Note that if $f_1, \ldots, f_k$ are independent random variables and $X_1, \ldots, X_k$ are events such that each X_i respects f_i, then $X_1, \ldots, X_k$ are independent.

Given a set of distributions $\mathcal{G}_i$ on Γ_i, we define the distribution $\mathcal{G} = \mathcal{G}_1 \times \cdots \times \mathcal{G}_k$ on $\Gamma_1 \times \cdots \times \Gamma_k$ by $\Pr_{\mathcal{G}}((g_1, \ldots, g_k)) = \prod_{i=1}^{k} \Pr_{\mathcal{G}_i}(g_i)$. Note that if $\mathcal{G} = \mathcal{G}_1 \times \cdots \times \mathcal{G}_k$ and $f_1, \ldots, f_k$ are random variables over $\Gamma_1 \times \cdots \times \Gamma_k$ such that each f_i depends only on the i^{th} component of its argument, then $f_1, \ldots, f_k$ are independent.

A source specifies how the inputs to a protocol are generated. Specifically, the players' inputs are chosen randomly according to some fixed distribution described by a source. Formally, a *source* is a quadruple $\mathcal{T} = (\Omega, \mathcal{F}, U, \theta)$, where

- Ω is a finite set.
- $\mathcal{F}$ specifies a probability distribution on Ω.

- $U = U_1 \times \cdots \times U_k$ is a k-tuple of sets.
- $\theta = (\theta_1, \ldots, \theta_k)$, where each $\theta_i : \Omega \to U_i$ is a *view function.*

We sometimes say $\mathcal{T}$ is a *source for U*. When a source is clear from context, we will use the notation Ω, $\mathcal{F}$, U, U_i, θ, and θ_i as specified above. When we wish to make a source $\mathcal{T}$ explicit, we use $\mathcal{T}$ as a superscript.

An element $\omega \in \Omega$ is called a *point* or a *global value*. Each view function θ_i is a random variable. Given a point $\omega \in \Omega$, we define the input $\theta(\omega) = (\theta_1(\omega), \ldots, \theta_k(\omega))$, and given an input u, we write $\theta^{-1}(u)$ to denote the event $\{\omega : \theta(\omega) = u\}$. A source $\mathcal{T}$ for U can itself be regarded as a distribution on U, where $\Pr_{\mathcal{T}}(u) = \Pr_{\mathcal{F}}[\theta^{-1}(u)]$. $\mathcal{T}$ also induces a distribution on each player P_i's input set $U_i^{\mathcal{T}}$.

A source $\mathcal{S}$ for U is *canonical* if $\Omega = U$ and $\theta_i(u) = \langle u \rangle_i$ for all $u \in U$. We allow noncanonical sources in order to allow greater flexibility in the information given to Eve. (See §2.4.) The following proposition states that for any distribution on U, it is always possible to construct a canonical source for U that realizes this distribution.

Proposition 2.2 *Let $U = U_1 \times \cdots \times U_k$ be a set, and let $\mathcal{F}$ be a distribution over U. Then there exists a canonical source $\mathcal{S}$ for U such that $\Pr_{\mathcal{S}}(u) = \Pr_{\mathcal{F}}(u)$ for all $u \in U$.*

We say that u *spans* U' if u touches every $u' \in U'$. We say a set $U' \subseteq U$ is *coverable* if there exists a feasible interpolant u of U' that spans U'. We also say u *covers* U'.

The special case of $U' = \{u_0, u_1\}$ arises frequently in the sequel. The following propositions are immediate from the definitions.

Proposition 2.3 *Let $\mathcal{T}$ be a source. If $\{u_0, u_1\} \subseteq \mathrm{feas}(U)$ and u_0 touches u_1, then $\{u_0, u_1\}$ is coverable.*

Proposition 2.4 *Let $\mathcal{T}$ be a source. If $u \in \mathrm{feas}(U) - \{u_0, u_1\}$ and u interpolates $\{u_0, u_1\}$, then $\{u_0, u_1\}$ is coverable.*

The distribution that $\mathcal{T}$ defines on $U^{\mathcal{T}}$ can incorporate both correlated and independent initial information for the team players. In the study of randomized algorithms and protocols, it is often desirable to consider private independent random information (such as coin flips) as being separate from any correlated initial information. To this end, we allow any source $\mathcal{T}$ to be augmented by additional independent randomization. Formally, a source $\mathcal{S}$ is a *randomized extension of* $\mathcal{T}$ if there exist finite sets $R_1, \ldots, R_k$ and distributions $\mathcal{R}_1, \ldots, \mathcal{R}_k$ such that $\mathcal{R}_i$ is a distribution on R_i for each i and

- $U^{\mathcal{S}} = (R_1 \times U_1^{\mathcal{T}}) \times \cdots \times (R_k \times U_k^{\mathcal{T}})$.
- $\Pr_{\mathcal{S}}((r_1, q_1), \ldots, (r_k, q_k)) = \Pr_{\mathcal{T}}(q_1, \ldots, q_k) \cdot \prod_{i=1}^{k} \Pr_{\mathcal{R}_j}(r_i)$.

We also say that $\mathcal{S}$ is $\mathcal{T}$ *extended by* $\mathcal{R}_1, \ldots, \mathcal{R}_k$. Here, r_i models the independent private random information for player P_i. Let $u = ((r_1, q_1), \ldots, (r_k, q_k)) \in U^{\mathcal{S}}$. We call $\operatorname{corr}(u) = (q_1, \ldots, q_k) \in U^{\mathcal{T}}$ the *correlated part of the input*; $\operatorname{corr}(\langle u \rangle_i) = \langle \operatorname{corr}(u) \rangle_i = q_i$ is P_i*'s share* of $\operatorname{corr}(u)$. We call $(r_1, \ldots, r_k)$ the *independent part of the input*, and we denote r_i by $\langle \operatorname{ind}(u) \rangle_i$. Note that q is feasible as an input of $\mathcal{T}$ if and only if q is the correlated part of a feasible input in $\mathcal{S}$.

2.4 Views for Eve

Given a source $\mathcal{S}$ with global value set Ω, a *view function for Eve* is a function $\theta_e : \Omega \to U_e$, where U_e is an arbitrary set. This formalizes the information that Eve is given about the global value. She also hears the conversation of the team players. Thus, if ω is the chosen global value, Eve is given the view $\theta_e(\omega)$ and hears the conversation $\operatorname{conv}(\theta(\omega))$.

A view function θ_e for Eve is *empty* if $\theta_e(\omega) = \theta_e(\omega')$ for all $\omega, \omega' \in \Omega$. Hence, an empty view gives Eve no additional information beyond the conversation. Note that if Eve were only given empty views, it would be sufficient to consider only canonical sources. However, the more general definition of sources allows consideration of information for Eve such as "Eve sees a random card from Alice's hand."

2.5 Systems

Together, a protocol $\mathcal{P}$ and a source $\mathcal{T}$ for $U^{\mathcal{P}}$ are a called a *system*, denoted $\mathcal{P}_{\mathcal{T}}$. We say $\mathcal{P}_{\mathcal{T}}$ is N*-valued* if $|V^{\mathcal{P}}| = N$. Since a source for $U^{\mathcal{P}}$ defines a distribution on $U^{\mathcal{P}} = U^{\mathcal{T}}$, $\mathcal{P}_{\mathcal{T}}$ induces a distribution on conversations and on the players' outputs. We extend the term 'feasible' to conversations: we say τ is *feasible* if there is a feasible input $u \in U^{\mathcal{T}}$ such that $\operatorname{conv}(u)[\ell] = \tau$. A system $\mathcal{P}_{\mathcal{S}}$ is a *randomized* $\mathcal{T}$*-system* if $\mathcal{S}$ is a randomized extension of $\mathcal{T}$.

Lemma 2.5 (Second Interpolation Lemma) *Let $\mathcal{P}_{\mathcal{S}}$ be a randomized $\mathcal{T}$-system, let $U' \subseteq \text{feas}\left(U^{\mathcal{S}}\right)$, and let $\widehat{u}$ interpolate U'. If $\text{corr}(\widehat{u})$ is feasible, then $\widehat{u}$ is feasible.*

Proof: Suppose the conditions of the lemma and let $(u_1, \ldots, u_k)$ be a U'-derivation of $\widehat{u}$. Then $\Pr_{\mathcal{S}}(\widehat{u}) = \Pr_{\mathcal{T}}(\text{corr}(\widehat{u})) \cdot \prod_{j=1}^{k} \Pr_{\mathcal{R}_j}(\langle \text{ind}(\widehat{u}) \rangle_j)$. Since every $u \in U'$ is feasible, then for $i \in K$,

$$\Pr_{\mathcal{S}}(u_i) = \Pr_{\mathcal{T}}(\text{corr}(u_i)) \cdot \prod_{j=1}^{k} \Pr_{\mathcal{R}_j}(\langle \text{ind}(u_i) \rangle_j) > 0.$$

Thus, in particular, $\Pr_{\mathcal{R}_i}(\langle \text{ind}(u_i) \rangle_i) > 0$ for $i \in K$. If, in addition, $\text{corr}(\widehat{u})$ is feasible, then $\Pr_{\mathcal{T}}(\text{corr}(\widehat{u})) > 0$. Hence, $\Pr_{\mathcal{S}}(\widehat{u}) = \Pr_{\mathcal{T}}(\text{corr}(\widehat{u})) \prod_{i=1}^{k} \Pr_{\mathcal{R}_i}(\langle \text{ind}(u_i) \rangle_i) > 0$, so $\widehat{u}$ is feasible. ■

We say two sources $\mathcal{S}$ and $\mathcal{S}'$ are *team-equivalent* if $U^{\mathcal{S}} = U^{\mathcal{S}'}$ and $\Pr_{\mathcal{S}}(u) = \Pr_{\mathcal{S}'}(u)$ for all $u \in U^{\mathcal{S}}$. Team-equivalence defines an equivalence relation on sources. Proposition 2.6 states that team-equivalence and randomized extensions behave nicely together.

Proposition 2.6 *Let $\mathcal{S}$ be a randomized extension of $\mathcal{T}$, let $\mathcal{S}'$ be team-equivalent to $\mathcal{S}$, and let $\mathcal{T}'$ be team-equivalent to $\mathcal{T}$. Then $\mathcal{S}'$ is a randomized extension of $\mathcal{T}'$.*

2.6 Example 2, Revisited

Recall the simple protocol $\mathcal{P}_2 = (2, 0, \{0,1\} \times \{0,1\}, \{0,1\}, \emptyset, \mu, \nu)$ introduced in §2.2, where $\nu_i(a_i, \epsilon) = a_i$, in which each player receives an input and outputs that input. An example of a source $\mathcal{S}_2$ for $U^{\mathcal{P}_2}$ is one that models the player's inputs as being generated by a single flip of a shared biased coin. Formally, $\mathcal{S}_2 = (\{0,1\}, \mathcal{F}, \{0,1\} \times \{0,1\}, \theta)$, where $\Pr_{\mathcal{F}}(0) = 1/3$, $\Pr_{\mathcal{F}}(1) = 2/3$, and $\theta_i(\omega) = \omega$ for $i \in \{1,2\}$ and $\omega \in \{0,1\}$. We will see in §3.5 that $(\mathcal{P}_2)_{(\mathcal{S}_2)}$ performs weak but not perfect secret key exchange.

CHAPTER 3: SECRET KEY EXCHANGE

"Secret key exchange" is used informally to mean the following. A "key" is a value that is chosen randomly from some fixed set of values. A key is "exchanged" if all the players learn the key. A key is "secret" if a passive computationally-unlimited eavesdropper, Eve, who may have some information about the players' inputs, cannot learn the key. We formalize each of these notions independently, as *uniformity*, *agreement*, and *secrecy*.

Uniformity restricts the *a priori* probability of an output. Secrecy restricts the relation between an output's *a priori* probability and its *a posteriori* probability. Hence, together uniformity and secrecy restrict the *a posteriori* probability. Uniformity and secrecy are both defined in two strengths: perfect and weak. The perfect uniformity and perfect secrecy conditions together imply that Eve has *no information* about the players' outputs, while the weak uniformity and weak secrecy conditions imply that Eve does not learn the players' outputs *with certainty.*

These conditions may also be useful for defining other problems in our model. For example, a system could be considered to perform secret message transmission if secrecy and agreement are satisfied along with a third condition that the output values must be equal to the specified message given to the designated sender as part of his view. In this dissertation, we limit our attention to the secret key exchange problem.

The remainder of this chapter is organized as follows. §3.1 defines the agreement, uniformity and secrecy conditions. §3.2 defines secret key exchange. §3.3 explores the effect of the information given to Eve. §3.4 shows two results relating the behavior of individual players and the behavior of the team as a whole. §3.5 shows that the system given in §2.6 performs weak but not perfect secret key exchange. Finally, §3.6 explores how the agreement, uniformity, and secrecy conditions restrict the behavior of the players and shows that secret key exchange is not possible if the players' inputs are not correlated. In contrast, §3.7 shows that secret key exchange can be possible when the players' inputs are pairwise independent.

3.1 Conditions for Secret Key Exchange

Fix a system $\mathcal{P}_{\mathcal{S}}$ and a view function θ_e over Ω for Eve, and let $i, j \in K$, $v \in V$, $\sigma \in \mathsf{c}$, $\ell = |\sigma|$, and $a_e \in U_e$. We define several events over Ω.

$$
\begin{aligned}
\mathrm{O}_i(v) &= \{\omega : \mathrm{out}_i(\theta(\omega)) = v\} \\
\mathrm{C}(\sigma) &= \{\omega : \mathrm{conv}(\theta(\omega))[\ell] = \sigma\} \\
\mathrm{E}(a_e) &= \{\omega : \theta_e(\omega) = a_e\}
\end{aligned}
$$

Thus $\mathrm{O}_i(v)$ is the event that P_i outputs v, $\mathrm{C}(\sigma)$ is the event that the conversation at round ℓ is σ, and $\mathrm{E}(a_e)$ is the event that Eve has view a_e. The events $\mathrm{O}_i(v)$ and $\mathrm{C}(\sigma)$ depend only on $\mathcal{P}_{\mathcal{S}}$, while the event $\mathrm{E}(a_e)$ depends on $\mathcal{P}_{\mathcal{S}}$ and θ_e. When we wish to explicitly make clear the system $\Psi = \mathcal{P}_{\mathcal{S}}$, we write $\mathrm{C}^{\Psi}(\sigma)$, $\mathrm{O}_i^{\Psi}(\sigma)$, and $\mathrm{E}^{\Psi}(a_e)$. When we wish to explicitly make clear Eve's view function, we write $\mathrm{E}_{\theta_e}(a_e)$.

We say $\mathcal{P}_{\mathcal{S}}$ satisfies agreement if the following condition holds.

- **Agreement**: $\Pr[\mathrm{O}_i(v_1) \cap \mathrm{O}_j(v_2)] = 0$ for all $i, j \in K$ and all pairs $v_1, v_2 \in V$ such that $v_1 \neq v_2$.

Thus, $\mathcal{P}_{\mathcal{S}}$ satisfies agreement if and only if the outputs of all team players agree at every feasible point.

We define two uniformity conditions to capture two types of distributions on the players' outputs.

- **Perfect uniformity**: $\Pr[\mathrm{O}_i(v_1)] = \Pr[\mathrm{O}_i(v_2)]$ for all $i \in K$ and all $v_1, v_2 \in V$.

- **Weak uniformity**: $\Pr[\mathrm{O}_i(v)] > 0$ for all $i \in K$ and all $v \in V$.

Thus, $\mathcal{P}_{\mathcal{S}}$ satisfies perfect uniformity if each team player's output is uniformly distributed over the output set, while $\mathcal{P}_{\mathcal{S}}$ satisfies weak uniformity if each team player outputs each value with positive probability.

Analogously, we define two levels of secrecy that limit the amount of information Eve is given by her view and the conversation. While the agreement and uniformity conditions apply to a system, the secrecy conditions apply to a system together with a view function for Eve. If the perfect (weak) secrecy condition holds, we say $\mathcal{P}_{\mathcal{S}}$ satisfies perfect (weak) secrecy *against* θ_e.

- **Perfect secrecy**: $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)] = \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] \cdot \Pr[\mathrm{O}_i(v)]$ for all $a_e \in U_e$, $\sigma \in \mathsf{cc}$, $i \in K$, and $v \in V$.

- **Weak secrecy**: $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)] > 0$ if and only if $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] > 0$ and $\Pr[\mathrm{O}_i(v)] > 0$ for all $a_e \in U_e$, $\sigma \in \mathsf{cc}$, $i \in K$, and $v \in V$.

The perfect secrecy condition, formulated as an independence condition, is essentially Shannon's formulation of perfect secrecy [Sha49]. $\mathcal{P}_{\mathcal{S}}$ satisfies perfect secrecy against θ_e if each team player's output is independent of the information available to Eve, i.e., her view and the conversation. Equivalently, perfect secrecy requires that Eve's probability of guessing a player's output correctly be the same whether or not she takes into account her view and the conversation.

Weak secrecy, on the other hand, requires only that an eavesdropper not be able to rule out any output for a given player. Specifically, weak secrecy requires that Eve consider each initially possible output for each team player to still be possible after hearing the conversation. Note that the weak secrecy condition could be equivalently formulated based on conditional probability to say that if $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] > 0$ and $\Pr[\mathrm{O}_i(v)] > 0$, then $\Pr[\mathrm{O}_i(v) \mid \mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] > 0$.

Proposition 3.1 *Let $\mathcal{P}_{\mathcal{S}}$ be a system and let θ_e be a view function over Ω for Eve. If $\mathcal{P}_{\mathcal{S}}$ satisfies perfect uniformity, then $\mathcal{P}_{\mathcal{S}}$ satisfies weak uniformity. If $\mathcal{P}_{\mathcal{S}}$ satisfies perfect secrecy against θ_e, then $\mathcal{P}_{\mathcal{S}}$ satisfies weak secrecy against θ_e.*

Proof: Let $\mathcal{P}_{\mathcal{S}}$ be a system and let θ_e be a view function over Ω for Eve. Suppose $\mathcal{P}_{\mathcal{S}}$ satisfies perfect uniformity. Then $\Pr[\mathrm{O}_i(v_1)] = \Pr[\mathrm{O}_i(v_2)]$ for all $v_1, v_2 \in V$. It follows that $\Pr[\mathrm{O}_i(v)] = 1/|V^{\mathcal{P}}| > 0$ for all $v \in V$, so $\mathcal{P}_{\mathcal{S}}$ satisfies weak uniformity.

Now suppose $\mathcal{P}_{\mathcal{S}}$ satisfies perfect secrecy against θ_e. If $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)] > 0$ then obviously $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] > 0$ and $\Pr[\mathrm{O}_i(v)] > 0$. Conversely, if $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] > 0$ and $\Pr[\mathrm{O}_i(v)] > 0$, then by perfect secrecy, $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)] = \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] \cdot \Pr[\mathrm{O}_i(v)] > 0$. Hence $\mathcal{P}_{\mathcal{S}}$ satisfies weak secrecy against θ_e. ■

We say a property of systems *respects* team-equivalence if, for every protocol $\mathcal{P}$ and every pair of team-equivalent sources $\mathcal{S}$ and $\mathcal{S}'$ for $U^{\mathcal{P}}$, $\mathcal{P}_{\mathcal{S}}$ satisfies the property if and only if $\mathcal{P}_{\mathcal{S}'}$ satisfies the property. Since the team players behave the same given inputs generated by team-equivalent sources, the following proposition holds.

Proposition 3.2 *The following properties respect team-equivalence: agreement, weak uniformity, perfect uniformity, weak secrecy against the empty view, perfect secrecy against the empty view.*

3.2 Secret Key Exchange Systems

Let $\mathcal{P}_{\mathcal{S}}$ be a system and let θ_e be a view function for Eve. $\mathcal{P}_{\mathcal{S}}$ performs *perfect (respectively weak) secret key exchange against* θ_e if $\mathcal{P}_{\mathcal{S}}$ satisfies agreement, perfect (weak) uniformity, and perfect (weak) secrecy against θ_e. We also say $\mathcal{P}_{\mathcal{S}}$ is a system *for perfect (weak) secret key exchange against* θ_e to mean the same thing. We say that $\mathcal{P}_{\mathcal{S}}$ performs, or is for, *perfect (weak) secret key exchange* if $\mathcal{P}_{\mathcal{S}}$ performs perfect (weak) secret key exchange against the empty view over Ω. A source $\mathcal{T}$ is *adequate* for N-valued perfect (weak) secret key exchange if there exists an N-valued randomized $\mathcal{T}$-system for perfect (weak) secret key exchange.

We combine our terminology for systems in the obvious way; for example, we say $\mathcal{P}_{\mathcal{S}}$ performs N-valued weak secret key exchange against θ_e if $\mathcal{P}_{\mathcal{S}}$ is N-valued and $\mathcal{P}_{\mathcal{S}}$ performs weak secret key exchange against θ_e. If $N = 2^n$, we also say $\mathcal{P}_{\mathcal{S}}$ performs n-bit secret key exchange. We abbreviate '2-valued secret key exchange' by 'secret bit exchange'. A secret bit is equivalent to a shared secret coin flip for the team.

3.3 The Effect of Eve's View

In this section, we explore the effect of Eve's view. Proposition 3.3 states that secrecy against an empty view can be rewritten without reference to Eve's view. Its proof is omitted.

Proposition 3.3 *Let $\mathcal{P}_{\mathcal{S}}$ be a system and let θ_e be an empty view for Eve.*

1. *$\mathcal{P}_{\mathcal{S}}$ satisfies perfect secrecy against θ_e if and only if*

$$\Pr\left[\mathrm{C}(\sigma) \cap \mathrm{O}_i(v)\right] = \Pr\left[\mathrm{C}(\sigma)\right] \cdot \Pr\left[\mathrm{O}_i(v)\right]$$

for all $\sigma \in \mathsf{cc}$, $i \in K$, and $v \in V$.

2. *$\mathcal{P}_{\mathcal{S}}$ satisfies weak secrecy against θ_e if and only if*

$$\Pr\left[\mathrm{C}(\sigma) \cap \mathrm{O}_i(v)\right] > 0 \Longleftrightarrow \Pr\left[\mathrm{C}(\sigma)\right] > 0 \text{ and } \Pr\left[\mathrm{O}_i(v)\right] > 0$$

for all $\sigma \in \mathsf{cc}$, $i \in K$, and $v \in V$.

Let θ_e and θ'_e be view functions over Ω for Eve. We say θ_e is a *refinement* of θ'_e if for all $\omega_0, \omega_1 \in \Omega$, $\theta_e(\omega_0) = \theta_e(\omega_1) \Rightarrow \theta'_e(\omega_0) = \theta'_e(\omega_1)$. Lemma 3.4 states that giving Eve more information in the form of a refined view can only help her.

Lemma 3.4 *Let $\theta_e : \Omega \to U_e$ be a refinement of $\theta'_e : \Omega \to U'_e$. If $\mathcal{P}_{\mathcal{S}}$ satisfies perfect (weak) secrecy against θ_e, then $\mathcal{P}_{\mathcal{S}}$ satisfies perfect (weak) secrecy against θ'_e.*

Proof: Let $\theta_e : \Omega \to U_e$ be a refinement of $\theta'_e : \Omega \to U'_e$. Let $a' \in U'_e$. Since θ_e is a refinement of θ'_e, $\mathrm{E}_{\theta'_e}(a')$ is the disjoint union of all distinct sets $\mathrm{E}_{\theta_e}(a)$ such that $\mathrm{E}_{\theta_e}(a) \subseteq \mathrm{E}_{\theta'_e}(a')$. Let $\widehat{U}_e = \{a \in U_e : \mathrm{E}_{\theta_e}(a) \subseteq \mathrm{E}_{\theta'_e}(a')\}$. Then

$$\Pr\left[\mathrm{E}_{\theta'_e}(a')\right] = \sum_{a \in \widehat{U}_e} \Pr\left[\mathrm{E}_{\theta_e}(a)\right]$$

Let $\sigma \in \mathsf{cc}$, $i \in K$, and $v \in V$. Suppose $\mathcal{P}_{\mathcal{S}}$ satisfies perfect secrecy against θ_e. Then $\Pr\left[\mathrm{E}_{\theta_e}(a) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)\right] = \Pr\left[\mathrm{E}_{\theta_e}(a) \cap \mathrm{C}(\sigma)\right] \cdot \Pr\left[\mathrm{O}_i(v)\right]$ for all $a \in U_e$. Summing over all $a \in \widehat{U}_e$ yields $\Pr\left[\mathrm{E}_{\theta'_e}(a') \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)\right] = \Pr\left[\mathrm{E}_{\theta'_e}(a') \cap \mathrm{C}(\sigma)\right] \cdot \Pr\left[\mathrm{O}_i(v)\right]$, so $\mathcal{P}_{\mathcal{S}}$ satisfies perfect secrecy against θ'_e.

Now suppose $\mathcal{P}_{\mathcal{S}}$ satisfies weak secrecy against θ_e. If $\Pr\left[\mathrm{E}(a') \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)\right] > 0$, then trivially $\Pr\left[\mathrm{E}(a') \cap \mathrm{C}(\sigma)\right] > 0$ and $\Pr\left[\mathrm{O}_i(v)\right] > 0$. Conversely, suppose $\Pr\left[\mathrm{E}(a') \cap \mathrm{C}(\sigma)\right] > 0$ and $\Pr\left[\mathrm{O}_i(v)\right] > 0$. Since $\Pr\left[\mathrm{E}_{\theta'_e}(a') \cap \mathrm{C}(\sigma)\right] > 0$, there is some $\mathrm{E}_{\theta_e}(a) \subseteq \mathrm{E}_{\theta'_e}(a')$ such that $\Pr\left[\mathrm{E}_{\theta'_e}(a) \cap \mathrm{C}(\sigma)\right] > 0$. Hence by weak secrecy against θ_e, $\Pr\left[\mathrm{E}_{\theta_e}(a) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)\right] > 0$, so $\Pr\left[\mathrm{E}_{\theta'_e}(a') \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)\right] \geq \Pr\left[\mathrm{E}_{\theta'_e}(a) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)\right] > 0$. Hence $\mathcal{P}_{\mathcal{S}}$ satisfies weak secrecy against θ'_e. ■

Corollary 3.5 *Let $\theta_e : \Omega \to U_e$ be a refinement of $\theta'_e : \Omega \to U'_e$. If $\mathcal{P}_{\mathcal{S}}$ performs N-valued perfect (weak) secret key exchange against θ_e, then $\mathcal{P}_{\mathcal{S}}$ performs N-valued perfect (weak) secret key exchange against θ'_e.*

Corollary 3.5 implies that if a system performs secret key exchange against any view, then it performs secret key exchange against any empty view.

Corollary 3.6 *Let $\mathcal{P}_{\mathcal{S}}$ be a system and let θ_e be a view function over Ω for Eve. If $\mathcal{P}_{\mathcal{S}}$ performs perfect (weak) secret key exchange against θ_e, then $\mathcal{P}_{\mathcal{S}}$ performs perfect (weak) secret key exchange against any empty view over Ω.*

Proof: Let $\mathcal{P}_{\mathcal{S}}$ be a system, let θ_e be a view function over Ω for Eve, and suppose that $\mathcal{P}_{\mathcal{S}}$ performs perfect (weak) secret key exchange against θ_e, let θ'_e be an empty view over Ω. Then any view function over Ω, in particular θ_e, is a refinement of θ'_e. By Corollary 3.5, $\mathcal{P}_{\mathcal{S}}$ satisfies perfect (weak) secrecy against θ'_e. ■

Let $\theta_e^*(x) = \emptyset$ for all x. Then θ_e^* is an empty view over Ω. Corollary 3.6 implies that if a system $\mathcal{P}_\mathcal{S}$ performs secret key exchange against θ_e^*, then $\mathcal{P}_\mathcal{S}$ satisfies performs secret key exchange against any empty view over Ω. We call θ_e^* the *canonical empty view*, or simply *the empty view.*

Theorem 3.7 shows that giving Eve independent random values can not help her thwart perfect secret key exchange and giving Eve "weakly independent" random values (in which every value for Eve is possible for every element of Ω) can not help her thwart weak secret key exchange.

Theorem 3.7 *Let $\mathcal{P}_\mathcal{S}$ be a system, let θ_e be a view function for Eve, and let $\mathcal{R}_e$ is a distribution over a set R_e. Let $\mathcal{S}'$ be such that $\Omega^{\mathcal{S}'} = \Omega \times R_e$, $\theta_i^{\mathcal{S}'}(\omega, r_e) = \theta_i^{\mathcal{S}}(\omega)$ for all $i \in K$, and $\theta_e^{\mathcal{S}'}(\omega, r_e) = (\theta_e(\omega), r_e)$.*

1. *If $\mathcal{F}^{\mathcal{S}'} = \mathcal{F} \times \mathcal{R}_e$ and $\mathcal{P}_\mathcal{S}$ performs perfect secret key exchange against θ_e, then $\mathcal{P}_{\mathcal{S}'}$ performs perfect secret key exchange against $\theta_e^{\mathcal{S}'}$.*

2. *If $\mathcal{F}^{\mathcal{S}'}$ is such that $\Pr_{\mathcal{S}'}(\omega, r_e) > 0$ whenever $\Pr_\mathcal{S}(\omega) > 0$, and if $\mathcal{P}_\mathcal{S}$ performs weak secret key exchange against θ_e, then $\mathcal{P}_{\mathcal{S}'}$ performs weak secret key exchange against $\theta_e^{\mathcal{S}'}$.*

Proof: Suppose the conditions of the theorem and suppose that $\mathcal{P}_\mathcal{S}$ performs perfect (weak) secret key exchange against θ_e. We denote the system $\mathcal{P}_\mathcal{S}$ by Ψ and the system $\mathcal{P}_{\mathcal{S}'}$ by Ψ'; we let Pr denote probability according to $\mathcal{F}^\mathcal{S}$ and $\Pr'$ denote probability according to $\mathcal{F}^{\mathcal{S}'}$. $\mathcal{S}$ and $\mathcal{S}'$ are team-equivalent, so by Proposition 3.2, Ψ' satisfies agreement and perfect (weak) uniformity.

Since $\Omega^{\mathcal{S}'} = \Omega \times R_e$,

$$\Pr'((\omega, r_e)) = \Pr(\omega) \cdot \Pr_{\mathcal{R}_e}(r_e)$$

for $\omega \in \Omega$ and $r_e \in R_e$. Hence, for any event $X \subseteq \Omega$ and $r_e \in R_e$,

$$\Pr'[X \times \{r_e\}] = \Pr[X] \cdot \Pr_{\mathcal{R}_e}(r_e)$$

Furthermore, let $a_e \in U_e$, $r_e \in R_e$, $\sigma \in \mathrm{cc}$, $v \in V$, and $i \in K$. Then

$$\begin{aligned} \mathrm{E}^{\Psi'}((a_e, r_e)) &= \mathrm{E}^{\Psi}(a_e) \times \{r_e\} \\ \mathrm{C}^{\Psi'}(\sigma) &= \mathrm{C}^{\Psi}(\sigma) \times R_e \\ \mathrm{O}_i^{\Psi'}(v) &= \mathrm{O}_i^{\Psi}(v) \times R_e \end{aligned}$$

Suppose that $\mathcal{F}^{\mathcal{S}'} = \mathcal{F}^{\mathcal{S}} \times \mathcal{R}_e$ and Ψ satisfies perfect secrecy against θ_e. Then it follows from the above that

$$\begin{aligned} &\Pr'\left[\mathrm{E}^{\Psi'}((a_e, r_e)) \cap \mathrm{C}^{\Psi'}(\sigma) \cap \mathrm{O}_i^{\Psi'}(v)\right] \\ &= \Pr'\left[\left(\mathrm{E}^{\Psi}(a_e) \times \{r_e\}\right) \cap \left(\mathrm{C}^{\Psi}(\sigma) \times R_e\right) \cap \left(\mathrm{O}_i^{\Psi}(v) \times R_e\right)\right] \\ &= \Pr'\left[\left(\mathrm{E}^{\Psi}(a_e) \cap \mathrm{C}^{\Psi}(\sigma) \cap \mathrm{O}_i^{\Psi}(v)\right) \times \{r_e\}\right] \\ &= \Pr\left[\mathrm{E}^{\Psi}(a_e) \cap \mathrm{C}^{\Psi}(\sigma) \cap \mathrm{O}_i^{\Psi}(v)\right] \cdot \Pr_{\mathcal{R}_e}(r_e) \end{aligned}$$

Similarly,

$$\begin{aligned} &\Pr'\left[\mathrm{E}^{\Psi'}((a_e, r_e)) \cap \mathrm{C}^{\Psi'}(\sigma)\right] \\ &= \Pr\left[\mathrm{E}^{\Psi}(a_e) \cap \mathrm{C}^{\Psi}(\sigma)\right] \cdot \Pr_{\mathcal{R}_e}(r_e) \end{aligned}$$

Therefore,

$$\begin{aligned} &\Pr'\left[\mathrm{E}^{\Psi'}((a_e, r_e)) \cap \mathrm{C}^{\Psi'}(\sigma) \cap \mathrm{O}_i^{\Psi'}(v)\right] \\ &= \Pr\left[\mathrm{E}^{\Psi}(a_e) \cap \mathrm{C}^{\Psi}(\sigma) \cap \mathrm{O}_i^{\Psi}(v)\right] \cdot \Pr_{\mathcal{R}_e}(r_e) \\ &= \Pr\left[\mathrm{E}^{\Psi}(a_e) \cap \mathrm{C}^{\Psi}(\sigma)\right] \cdot \Pr\left[\mathrm{O}_i^{\Psi}(v)\right] \cdot \Pr_{\mathcal{R}_e}(r_e) \\ &= \Pr'\left[\mathrm{E}^{\Psi'}((a_e, r_e)) \cap \mathrm{C}^{\Psi'}(\sigma)\right] \cdot \Pr'\left[\mathrm{O}_i^{\Psi'}(v)\right] \end{aligned}$$

so Ψ' satisfies perfect secrecy against $\theta_e^{\mathcal{S}'}$.

Now suppose that $\mathcal{F}^{\mathcal{S}'}$ is such that $\Pr_{\mathcal{S}'}(\omega, r_e) > 0$ whenever $\Pr_{\mathcal{S}}(\omega) > 0$ and Ψ satisfies weak secrecy against θ_e. If $\Pr'\left[\mathrm{E}^{\Psi'}(a_e, r_e) \cap \mathrm{C}^{\Psi'}(\sigma) \cap \mathrm{O}_i^{\Psi'}(v)\right] > 0$, then clearly $\Pr'\left[\mathrm{E}^{\Psi'}(a_e, r_e) \cap \mathrm{C}^{\Psi'}(\sigma)\right] > 0$ and $\Pr'\left[\mathrm{O}_i^{\Psi'}(v)\right] > 0$. Conversely, suppose that $\Pr'\left[\mathrm{E}^{\Psi'}(a_e, r_e) \cap \mathrm{C}^{\Psi'}(\sigma)\right] > 0$ and $\Pr'\left[\mathrm{O}_i^{\Psi'}(v)\right] > 0$. Then

$$\begin{aligned} \Pr\left[\mathrm{E}^{\Psi}(a_e) \cap \mathrm{C}^{\Psi}(\sigma)\right] &= \Pr'\left[\left(\bigcup_{r \in R_e} \mathrm{E}^{\Psi'}((a_e, r))\right) \cap \left(\mathrm{C}^{\Psi'}(\sigma)\right)\right] \\ &\geq \Pr'\left[\mathrm{E}^{\Psi'}(a_e, r_e) \cap \mathrm{C}^{\Psi'}(\sigma)\right] \\ &> 0 \end{aligned}$$

and $\Pr\left[\mathrm{O}_i^{\Psi}(v)\right] = \Pr'\left[\mathrm{O}_i^{\Psi'}(v)\right] > 0$. Hence, since Ψ satisfies weak secrecy against θ_e, we have that $\Pr\left[\mathrm{E}^{\Psi}(a_e) \cap \mathrm{C}^{\Psi}(\sigma) \cap \mathrm{O}_i^{\Psi}(v)\right] > 0$. By the assumption that $\Pr_{\mathcal{S}'}(\omega, r_e) > 0$ whenever $\Pr_{\mathcal{S}}(\omega) > 0$, it follows that $\Pr'\left[\mathrm{E}^{\Psi'}(a_e, r_e) \cap \mathrm{C}^{\Psi'}(\sigma) \cap \mathrm{O}_i^{\Psi'}(v)\right] > 0$. Hence Ψ' satisfies weak secrecy against θ_e. ∎

3.4 Behavior of the Players

In this section, we examine the relationship between the behavior of individual players and the behavior of the team as a whole. To this end, we define some additional events. For $i \in K$, $v \in V$, and $\sigma \in \mathsf{cc}$, let

$$\begin{array}{lcl} \mathrm{O}(v) & = & \bigcap_{i=1}^{k} \mathrm{O}_i(v) \\ \mathrm{CO}(\sigma, v) & = & \mathrm{C}(\sigma) \cap \mathrm{O}(v) \end{array}$$

Thus $\mathrm{O}(v)$ is the event that v is output by all players and $\mathrm{CO}(\sigma, v)$ is the event that the conversation is σ and all players output v. As before, when we wish to explicitly make clear the system $\Psi = \mathcal{P}_{\mathcal{S}}$, we write $\mathrm{O}^{\Psi}(v)$ and $\mathrm{CO}^{\Psi}(\sigma, v)$.

We are also interested in the behavior of individual players with regard to an arbitrary complete conversation σ and output v. We define several events that express how player P_i would behave in certain hypothetical circumstances.

$$\begin{array}{lcl} \overline{\mathrm{O}}_i(\sigma, v) & = & \{\omega : \nu_i(\theta_i(\omega), \sigma) = v\} \\ \overline{\mathrm{C}}_i(\sigma) & = & \{\omega : \mu_i(\theta_i(\omega), \sigma[\ell - 1]) = \langle \sigma_\ell \rangle_i \text{ for } 1 \leq \ell \leq t\} \\ \overline{\mathrm{CO}}_i(\sigma, v) & = & \overline{\mathrm{C}}_i(\sigma) \cap \overline{\mathrm{O}}_i(\sigma, v) \end{array}$$

Thus $\overline{\mathrm{O}}_i(\sigma, v)$ is the event that P_i would output v if presented with the fixed conversation σ. $\overline{\mathrm{C}}_i(\sigma)$ is the event that P_i would play according to σ at every round ℓ if presented with the length ℓ prefix of σ. $\overline{\mathrm{CO}}_i(\sigma, v)$ is the event that P_i would play according to σ at every round if presented with a prefix of σ and would output v if presented with σ.

Lemma 3.8 states the intuitive fact that for each point ω, if each player plays according to σ when presented with any prefix of σ, then σ really is the conversation of $\theta(\omega)$. Lemma 3.9 states that if each player plays according to σ when presented with any prefix of σ and outputs v when presented with σ, then all players really do play according to σ and output v.

Lemma 3.8 *Let $\sigma \in \mathsf{cc}$. Then $\bigcap_{i=1}^{k} \overline{\mathrm{C}}_i(\sigma) = \mathrm{C}(\sigma)$.*

Proof: Let $\sigma \in \mathsf{cc}$. Then

$$\begin{array}{lcl} \displaystyle\bigcap_{i=1}^{k} \overline{\mathrm{C}}_i(\sigma) & = & \{\omega : \mu_i(\theta_i(\omega), \sigma[\ell - 1]) = \langle \sigma_\ell \rangle_i \text{ for } 1 \leq \ell \leq t \text{ and } i \in K\} \\ & = & \{\omega : \mathrm{conv}(\theta(\omega)) = \sigma\} \\ & = & \mathrm{C}(\sigma) \end{array}$$

as desired. ■

Lemma 3.9 *Let $\sigma \in \mathrm{cc}$ and $v \in V$. Then $\bigcap_{i=1}^{k} \overline{\mathrm{CO}}_i(\sigma, v) = \mathrm{CO}(\sigma, v)$.*

Proof: Let $\sigma \in \mathrm{cc}$ and $v \in V$. Then

$$\begin{aligned}
&\mathrm{C}(\sigma) \cap \bigcap_{i=1}^{k} \overline{\mathrm{O}}_i(\sigma, v) \\
&= \{\omega : \mathrm{conv}(\theta(\omega)) = \sigma \text{ and } \nu_i(\theta_i(\omega), \sigma) = v \text{ for } i \in K\} \\
&= \{\omega : \mathrm{conv}(\theta(\omega)) = \sigma \text{ and } \nu_i(\theta_i(\omega), \mathrm{conv}(\theta(\omega))) = v \text{ for } i \in K\} \\
&= \{\omega : \mathrm{conv}(\theta(\omega)) = \sigma \text{ and } \mathrm{out}_i(\theta(\omega)) = v \text{ for } i \in K\} \\
&= \mathrm{C}(\sigma) \cap \mathrm{O}(v) \\
&= \mathrm{CO}(\sigma, v)
\end{aligned}$$

Since it follows from Lemma 3.8 that $\bigcap_{i=1}^{k} \overline{\mathrm{CO}}_i(\sigma, v) = \mathrm{C}(\sigma) \cap \bigcap_{i=1}^{k} \overline{\mathrm{O}}_i(\sigma, v)$, we have $\bigcap_{i=1}^{k} \overline{\mathrm{CO}}_i(\sigma, v) = \mathrm{CO}(\sigma, v)$, as desired. ■

3.5 Example 2, Revisited Again

We show here that the system $(\mathcal{P}_2)_{(\mathcal{S}_2)}$ discussed in §2.2 and §2.6 performs weak but not perfect secret key exchange. Recall that $\mathcal{P}_2 = (2, 0, \{0,1\} \times \{0,1\}, \{0,1\}, \emptyset, \mu, \nu)$, where $\nu_i(a_i, \epsilon) = a_i$ and $\mathcal{S}_2 = (\{0,1\}, \mathcal{F}, \{0,1\} \times \{0,1\}, \theta)$, where $\Pr_{\mathcal{F}}(0) = 1/3$, $\Pr_{\mathcal{F}}(1) = 2/3$, and $\theta_i(\omega) = \omega$ for $i \in \{1,2\}$ and $\omega \in \{0,1\}$.

Theorem 3.10 *$(\mathcal{P}_2)_{(\mathcal{S}_2)}$ performs weak secret bit exchange but does not perform perfect secret bit exchange.*

Proof: There are two global values for $(\mathcal{P}_2)_{(\mathcal{S}_2)}$: 0 (with probability 1/3) and 1 (with probability 2/3). When the global value ω is chosen, both players output ω. The conversation is always the empty conversation ϵ. Hence, $\mathrm{O}_1(0) = \mathrm{O}_2(0) = \{0\}$ and $\mathrm{O}_1(1) = \mathrm{O}_2(1) = \{1\}$, so $\Pr[\mathrm{O}_1(0)] = \Pr[\mathrm{O}_2(0)] = 1/3$ and $\Pr[\mathrm{O}_1(1)] = \Pr[\mathrm{O}_2(1)] = 2/3$. Since $\mathrm{C}(\epsilon) = \Omega$, it follows that $\Pr[\mathrm{C}(\epsilon) \cap \mathrm{O}_1(0)] = \Pr[\mathrm{C}(\epsilon) \cap \mathrm{O}_2(0)] = 1/3$ and $\Pr[\mathrm{C}(\epsilon) \cap \mathrm{O}_1(1)] = \Pr[\mathrm{C}(\epsilon) \cap \mathrm{O}_2(1)] = 2/3$.

Then $\mathrm{O}_1(0) \cap \mathrm{O}_2(1) = \mathrm{O}_1(1) \cap \mathrm{O}_2(0) = \emptyset$, so agreement is satisfied. $\Pr[\mathrm{O}_i(v)] > 0$ for all $i \in \{1,2\}$ and $v \in \{0,1\}$, so weak uniformity is satisfied, but $\Pr[\mathrm{O}_1(0)] \neq \Pr[\mathrm{O}_1(1)]$, so perfect uniformity is not satisfied. As computed above, $\Pr[\mathrm{C}(\epsilon) \cap \mathrm{O}_i(v)] > 0$ for all $i \in \{1,2\}$ and $v \in \{0,1\}$. By Proposition 3.3, weak secrecy against the empty view is satisfied. Since $(\mathcal{P}_2)_{(\mathcal{S}_2)}$ is 2-valued and satisfies agreement, weak uniformity, and weak secrecy against the empty view, $(\mathcal{P}_2)_{(\mathcal{S}_2)}$ performs weak secret bit exchange. Since $(\mathcal{P}_2)_{(\mathcal{S}_2)}$ does not satisfy perfect uniformity, $(\mathcal{P}_2)_{(\mathcal{S}_2)}$ does not perform perfect secret bit exchange. ■

We will see in §4.2 that $\mathcal{S}_2$ is not adequate for perfect secret bit exchange.

3.6 Properties of Secret Key Exchange Systems

In this section, we examine how the secret key exchange conditions restrict the inputs, conversation, and outputs of the players. Lemma 3.11 provides several alternative characterizations of the agreement condition.

Lemma 3.11 *Let $\mathcal{P}_\mathcal{S}$ be a system. The following are equivalent.*

1. *$\mathcal{P}_\mathcal{S}$ satisfies agreement.*
2. *$\mathrm{out}_i(u) = \mathrm{out}_j(u)$ for all feasible $u \in U$ and $i, j \in K$.*
3. *$\omega \in \mathrm{O}_i(v) \Rightarrow \omega \in \mathrm{O}(v)$ for all feasible $\omega \in \Omega$, $v \in V$, and $i \in K$.*
4. *$\Pr[X \cap \mathrm{O}_i(v)] = \Pr[X \cap \mathrm{O}(v)]$ for all $X \subseteq \Omega$, $i \in K$ and $v \in V$.*
5. *$\Pr[X] = \sum_{v \in V} \Pr[X \cap \mathrm{O}(v)]$ for all $X \subseteq \Omega$.*

Proof: (1) $\Rightarrow$ (2): Suppose that (1) holds. Let $u \in U$ be feasible, and let $i, j \in K$. By definition of $\mathrm{out}_i(u)$, $\theta^{-1}(u) \subseteq \mathrm{O}_i(\mathrm{out}_i(u)) \cap \mathrm{O}_j(\mathrm{out}_j(u))$. Since u is feasible, $\theta^{-1}(u)$ contains a feasible point. Since agreement is satisfied, $\mathrm{out}_i(u) = \mathrm{out}_j(u)$.

(2) $\Rightarrow$ (3): Suppose that (2) holds. Let $v \in V$, let $\omega \in \mathrm{O}_i(v)$ be feasible, and let $i \in K$. Then $v = \mathrm{out}_i(\theta(\omega))$. $\theta(\omega)$ is feasible, so by (2), $\mathrm{out}_j(\theta(\omega)) = v$ for all $j \in K$. It follows that $\omega \in \mathrm{O}(v)$.

(3) $\Rightarrow$ (4): Suppose that (3) holds. Let $X \subseteq \Omega$, $i \in K$ and $v \in V$. By (3), if $\omega \in \mathrm{O}_i(v) - \mathrm{O}(v)$, then ω is not feasible. It follows that $\Pr[\mathrm{O}_i(v) - \mathrm{O}(v)] = 0$. Since in addition $\mathrm{O}(v) \subseteq \mathrm{O}_i(v)$, we have

$$\begin{aligned} \Pr[X \cap \mathrm{O}_i(v)] &= \Pr[X \cap \mathrm{O}(v)] + \Pr[X \cap (\mathrm{O}_i(v) - \mathrm{O}(v))] \\ &= \Pr[X \cap \mathrm{O}(v)] \end{aligned}$$

(4) $\Rightarrow$ (5): Suppose that (4) holds. Let $X \subseteq \Omega$ and let $i \in K$. Then X is the disjoint union over $v \in V$ of $X \cap \mathrm{O}_i(v)$ since $\{\mathrm{O}_i(v) : v \in V\}$ is a partition of Ω. Hence $\Pr[X] = \sum_{v \in V} \Pr[X \cap \mathrm{O}_i(v)]$. It follows by (4) that $\Pr[X] = \sum_{v \in V} \Pr[X \cap \mathrm{O}(v)]$.

(5) $\Rightarrow$ (1): Suppose that (5) holds. Let $i, j \in K$ and $v_1, v_2 \in V$ such that $v_1 \neq v_2$, and let $X = \mathrm{O}_i(v_1) \cap \mathrm{O}_j(v_2)$. Then by (5), $\Pr[X] = \sum_{v \in V} \Pr[X \cap \mathrm{O}(v)]$. However, since $v_1 \neq v_2$, it follows that $X \cap \mathrm{O}(v) = \mathrm{O}_i(v_1) \cap \mathrm{O}_j(v_2) \cap \mathrm{O}(v) = \emptyset$. Hence $\Pr[X] = \sum_{v \in V} \Pr[X \cap \mathrm{O}(v)] = 0$, so agreement is satisfied. ■

Lemma 3.12 shows that a system performs N-valued perfect secret key exchange if and only if, for each conversation and input for Eve, each output value occurs with conditional probability $1/N$. Similarly, Lemma 3.14 shows that if a system performs N-valued weak secret key exchange, then for each feasible conversation and input for Eve, each output value must occur with nonzero conditional probability.

Lemma 3.12 *Let $\mathcal{P}_{\mathcal{S}}$ be N-valued and let θ_e be a view function for Eve. The following are equivalent.*

1. *$\mathcal{P}_{\mathcal{S}}$ performs N-valued perfect secret key exchange against θ_e.*

2. *For all $a_e \in U_e$ and $\sigma \in \mathrm{cc}$, $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] = \sum_{v \in V} \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)]$ and all the terms of the sum are equal.*

3. *For all $a_e \in U_e$, $\sigma \in \mathrm{cc}$, and $v \in V$, $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] = N \cdot \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)]$.*

Proof: Let $\mathcal{P}_{\mathcal{S}}$ be an N-valued system and let θ_e be a view function for Eve.

(1) $\Rightarrow$ (2): Suppose that (1) holds. Then $\mathcal{P}_{\mathcal{S}}$ satisfies agreement, uniformity, and perfect secrecy against θ_e. Let $a_e \in U_e$, let $\sigma \in \mathrm{cc}$, and let $i \in K$. By Lemma 3.11 (part 5),

$$\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] = \sum_{v \in V} \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)] \tag{3.1}$$

For $v \in V$,

$$\begin{aligned} \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)] &= \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)] && (3.2) \\ &= \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] \cdot \Pr[\mathrm{O}_i(v)] && (3.3) \end{aligned}$$

where (3.2) is by Lemma 3.11 (part 4) and (3.3) is by perfect secrecy against θ_e. Let $v_1, v_2 \in V$. By perfect uniformity, $\Pr[\mathrm{O}_i(v_1)] = \Pr[\mathrm{O}_i(v_2)]$. Hence, by (3.3),

$$\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v_1)] = \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v_2)]$$

It follows that all terms in the sum (3.1) are equal.

(2) $\Rightarrow$ (3): Suppose that (2) holds. Let $a_e \in U_e$, and let $\sigma \in \mathsf{cc}$. By (2), $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] = \sum_{v \in V} \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)]$, and $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v_1)] = \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v_2)]$ for all $v_1, v_2 \in V$. Since $|V| = N$, it follows that $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] = N \cdot \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)]$ for all $v \in V$.

(3) $\Rightarrow$ (1): Suppose that (3) holds. Let $v \in V$. By (3),

$$\begin{aligned} 1 &= \sum_{a_e \in U_e} \sum_{\sigma \in \mathsf{cc}} \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] \\ &= \sum_{a_e \in U_e} \sum_{\sigma \in \mathsf{cc}} N \cdot \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)] \\ &= N \cdot \Pr[\mathrm{O}(v)] \end{aligned}$$

Hence,

$$\Pr[\mathrm{O}(v)] = \frac{1}{N} \tag{3.4}$$

Let $i \in K$. Since Ω is the disjoint union $\bigcup_{v \in V} \mathrm{O}_i(v)$, we have $\sum_{v \in V} \Pr[\mathrm{O}_i(v)] = 1$. By (3.4), $\sum_{v \in V} \Pr[\mathrm{O}(v)] = 1$. Hence

$$\sum_{v \in V} \Pr[\mathrm{O}_i(v)] = \sum_{v \in V} \Pr[\mathrm{O}(v)] \tag{3.5}$$

Let $v \in V$ and let $\omega \in \Omega$ be feasible. Since $\mathrm{O}_i(v) \supseteq \mathrm{O}(v)$, it follows from (3.5) that $\mathrm{feas}(\mathrm{O}_i(v)) = \mathrm{feas}(\mathrm{O}(v))$ and hence that

$$\Pr[\mathrm{O}_i(v)] = \Pr[\mathrm{O}(v)] = \frac{1}{N} \tag{3.6}$$

Thus $\mathcal{P}_{\mathcal{S}}$ satisfies perfect uniformity.

Finally, let $v \in V$, let $\sigma \in \mathsf{cc}$, let $a_e \in U_e$, and let $i \in K$. Then

$$\begin{aligned} \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)] &= \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)] && (3.7) \\ &= \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] \cdot \frac{1}{N} && (3.8) \\ &= \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] \cdot \Pr[\mathrm{O}_i(v)] && (3.9) \end{aligned}$$

where (3.7) follows from Lemma 3.11 (part 4), (3.8) follows from (3), and (3.9) follows from (3.6). Hence $\mathcal{P}_{\mathcal{S}}$ satisfies perfect secrecy against θ_e.

We conclude that $\mathcal{P}_{\mathcal{S}}$ performs N-valued perfect secret key exchange, as desired.

■

Corollary 3.13 *Let $\mathcal{P}_{\mathcal{S}}$ be N-valued. Then $\mathcal{P}_{\mathcal{S}}$ performs perfect secret key exchange (against the empty view) if and only if* $\Pr[\mathrm{C}(\sigma)] = N \cdot \Pr[\mathrm{CO}(\sigma, v)]$ *for all* $\sigma \in \mathsf{cc}$ *and* $v \in V$.

Lemma 3.14 *Let $\mathcal{P}_{\mathcal{S}}$ be N-valued and let θ_e be a view function for Eve. $\mathcal{P}_{\mathcal{S}}$ performs weak secret key exchange against θ_e if and only if for all $a_e \in U_e$ and $\sigma \in \mathsf{cc}$,*

1. $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] = \sum_{v \in V} \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)]$, *and*

2. *if* $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] > 0$ *then* $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)] > 0$ *for all* $v \in V$.

Proof: Suppose the conditions of the lemma and suppose that $\mathcal{P}_{\mathcal{S}}$ performs N-valued weak secret key exchange against θ_e. Then $\mathcal{P}_{\mathcal{S}}$ satisfies agreement, weak uniformity, and weak secrecy against θ_e. By Lemma 3.11 (part 5), $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] = \sum_{v \in V} \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)]$. Let $v \in V$ and suppose that $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] > 0$. By weak uniformity, $\Pr[\mathrm{O}_i(v)] > 0$. Hence, by Lemma 3.11 (part 4) and weak secrecy against θ_e, $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)] = \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)] > 0$.

For the converse, suppose that $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] = \sum_{v \in V} \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)]$, where if $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] > 0$ then $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)] > 0$ for all $v \in V$. Let $v \in V$, let $i \in K$, and let $\omega \in \Omega$ be feasible. Since Ω is the disjoint union $\bigcup_{v \in V} \mathrm{O}_i(v)$, we have $\sum_{v \in V} \Pr[\mathrm{O}_i(v)] = 1$. By assumption, $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] = \sum_{v \in V} \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)]$. Summing over all $a_e \in U_e$ and $\sigma \in \mathsf{cc}$ yields $\sum_{v \in V} \Pr[\mathrm{O}(v)] = 1$. Hence $\sum_{v \in V} \Pr[\mathrm{O}_i(v)] = \sum_{v \in V} \Pr[\mathrm{O}(v)]$. Since $\mathrm{O}_i(v) \supseteq \mathrm{O}(v)$, it follows that $\Pr[\mathrm{O}_i(v)] = \Pr[\mathrm{O}(v)]$, and therefore that $\omega \in \mathrm{O}_i(v) \Rightarrow \omega \in \mathrm{O}(v)$. By Lemma 3.11 (part 1), $\mathcal{P}_{\mathcal{S}}$ satisfies agreement. Let $a'_e = \theta_e(\omega)$ and $\sigma' = \mathrm{conv}(\omega)$. Then $\omega \in \mathrm{E}(a'_e) \cap \mathrm{C}(\sigma')$, so $\Pr[\mathrm{E}(a'_e) \cap \mathrm{C}(\sigma')] > 0$. Hence, by assumption, $\Pr[\mathrm{E}(a'_e) \cap \mathrm{C}(\sigma') \cap \mathrm{O}(v)] > 0$. It follows that $\Pr[\mathrm{O}_i(v)] = \Pr[\mathrm{O}(v)] > 0$, so weak uniformity is satisfied. Let $a_e \in U_e$ and $\sigma \in \mathsf{cc}$ and suppose $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] > 0$. Then $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)] = \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)] > 0$, so weak secrecy against θ_e is satisfied. It follows that $\mathcal{P}_{\mathcal{S}}$ performs weak secret key exchange against θ_e. ∎

Corollary 3.15 *Let $\mathcal{P}_{\mathcal{S}}$ be N-valued. $\mathcal{P}_{\mathcal{S}}$ performs weak secret key exchange (against the empty view) if and only if for $\sigma \in \mathsf{cc}$,*

1. $\Pr[\mathrm{C}(\sigma)] = \sum_{v \in V} \Pr[\mathrm{CO}(\sigma, v)]$, *and*

2. *if* $\Pr[\mathrm{C}(\sigma)] > 0$ *then* $\Pr[\mathrm{CO}(\sigma, \mathrm{O}(v))] > 0$ *for all* $v \in V$.

If two inputs that give rise to the same conversation have a nontrivial feasible interpolant (i.e. the set of two inputs is coverable), then some team players can not distinguish the first input vector from the interpolant and some team players can not distinguish the second input vector from the interpolant. It follows that all team players must output the same value on all three inputs. The following lemma formalizes this argument to show that in order for a system to perform secret key exchange, there must be inputs that are not coverable.

Lemma 3.16 *Let $\mathcal{P}_{\mathcal{S}}$ be an N-valued system such that $N \geq 2$, and suppose that every set of two feasible inputs is coverable. Then $\mathcal{P}_{\mathcal{S}}$ does not perform weak (perfect) secret key exchange.*

Proof: By Proposition 3.1, it suffices to show the lemma holds for weak secret key exchange. Suppose the conditions of the lemma and suppose by way of contradiction that $\mathcal{P}_{\mathcal{S}}$ performs weak secret key exchange. Let σ be a feasible complete conversation and let $\omega, \omega' \in \mathrm{feas}\,(\mathrm{C}(\sigma))$. Let $u = \theta(\omega)$ and let $u' = \theta(\omega')$. By assumption, there exists $\widehat{u}$ that covers $\{u, u'\}$. Since $\widehat{u}$ spans $\{u, u'\}$, there exist $i, i' \in K$ such that $u =_i \widehat{u}$ and $u' =_{i'} \widehat{u}$. Let $v = \mathrm{out}_i(u)$ and $v' = \mathrm{out}_{i'}(u')$. Then $\omega \in \mathrm{O}_i(v)$ and $\omega' \in \mathrm{O}_{i'}(v')$, so by Lemma 3.11 (part 3), $\omega \in \mathrm{O}(v)$ and $\omega' \in \mathrm{O}(v')$. Hence $\omega \in \mathrm{CO}(\sigma, v)$ and $\omega' \in \mathrm{CO}(\sigma, v')$. By Lemma 2.1 (parts 1 and 3), $\mathrm{conv}(\widehat{u}) = \sigma$, $\mathrm{out}_i(\widehat{u}) = v$, and $\mathrm{out}_{i'}(\widehat{u}) = v'$. Since $\widehat{u}$ is feasible, it follows from Lemma 3.11 (part 2) that $v = v'$, so $\omega' \in \mathrm{CO}(\sigma, v)$.

Since ω and ω' were chosen arbitrarily from $\mathrm{feas}\,(\mathrm{C}(\sigma))$, it follows that $\omega \in \mathrm{CO}(\sigma, v)$ for all $\omega \in \mathrm{feas}\,(\mathrm{C}(\sigma))$. Let $y \in V - \{v\}$. Then $\Pr[\mathrm{CO}(\sigma, y)] = 0$, a contradiction to Corollary 3.15. We conclude $\mathcal{P}_{\mathcal{S}}$ does not perform weak secret key exchange. ∎

A direct consequence of Lemma 3.16 is the intuitive result that the players' inputs must be correlated in order for secret key exchange to be possible. Specifically, if the team players can be divided into two sets such that the inputs of one set are independent of the other, regardless of the correlations within the sets, then secret key exchange is not possible. Given a set $K' \subseteq K$ and an input $u \in U$, we define the event $\mathrm{I}_{K'}(u) = \{u' : u' =_i u \text{ for all } i \in K'\}$ over U. Thus $\mathrm{I}_{K'}(u)$ is the event that the team players in K' have inputs as specified by u. In particular, if there are sets K_0 and K_1 such that $\mathrm{I}_{K_0}(u)$ and $\mathrm{I}_{K_1}(u)$ are independent events for every input u, then there is no prior shared secret information between the team players in K_0 and the team players in K_1.

Theorem 3.17 *Let $\mathcal{P}_{\mathcal{S}}$ be an N-valued system such that $N \geq 2$, and suppose there is a partition of K into nonempty sets K_0 and K_1 such that $\mathrm{I}_{K_0}(u)$ and $\mathrm{I}_{K_1}(u)$ are* *independent events for all $u \in U$. Then $\mathcal{P}_S$ does not perform weak (perfect) secret key exchange.*

Proof: Suppose the conditions of the theorem, and let u and u' be arbitrary distinct feasible inputs. Let $u_i = u$ for $i \in K_0$, let $u_i = u'$ for $i \in K_1$, and let $\widehat{u} = (\langle u_1 \rangle_1, \ldots, \langle u_k \rangle_k)$. Then $\widehat{u}$ interpolates $\{u, u'\}$. Since K_0 and K_1 are nonempty, $\widehat{u}$ spans $\{u, u'\}$.

By the independence condition,

$$\begin{aligned}
\Pr[\mathrm{I}_K(\widehat{u})] &= \Pr[\mathrm{I}_{K_0}(\widehat{u}) \cap \mathrm{I}_{K_1}(\widehat{u})] \\
&= \Pr[\mathrm{I}_{K_0}(\widehat{u})] \cdot \Pr[\mathrm{I}_{K_1}(\widehat{u})]
\end{aligned}$$

Since $u \in \mathrm{I}_{K_0}(\widehat{u})$ and u is feasible, $\Pr[\mathrm{I}_{K_0}(\widehat{u})] > 0$. Similarly, $\Pr[\mathrm{I}_{K_1}(\widehat{u})] > 0$. It follows that $\Pr[\mathrm{I}_K(\widehat{u})] > 0$. Since by definition $\mathrm{I}_K(\widehat{u}) = \{\widehat{u}\}$, it follows that $\widehat{u}$ is feasible. Hence $\widehat{u}$ covers $\{u, u'\}$. By Lemma 3.16, $\mathcal{P}_S$ does not perform weak (perfect) secret key exchange. ■

Theorem 3.17 implies the folklore result that public key cryptography is not possible in the presence of a computationally unlimited eavesdropper, since in public key cryptography the participants are assumed to have no prior shared secret information. The first written reference to such results we are aware of is in Rudich's dissertation [Rud88], where he shows that public key cryptography is not possible against a suitably powerful adversary.

3.7 Example 3

Theorem 3.17 shows that if the players' inputs are jointly independent, secret key exchange can not take place. However, it is not the case that pairwise independence precludes the possibility of secret key exchange, as is illustrated by the following simple system.

Consider three team players, Alice, Bob and Carol. Alice is given two random bits a_1 and a_2, Bob is given two random bits b_1 and b_2, and Carol is given $c_1 = a_1 \oplus b_1$ and $c_2 = a_2 \oplus b_2$ and a random bit d. Then the inputs of Alice, Bob, and Carol are pairwise independent (but not jointly independent), and the following one-round protocol achieves perfect secret bit exchange: Alice announces a_1, Bob announces b_2, Carol announces $e_1 = c_1 \oplus d$ and $e_2 = c_2 \oplus d$. From Bob's message, Alice can compute $c_2 = a_2 \oplus b_2$. From Alice's message, Bob can compute $c_1 = a_1 \oplus b_1$. Thus, from Carol's message, Alice can compute $d = c_2 \oplus e_2$ and Bob can both compute $d = c_2 \oplus e_2$. Hence, after the first round all players know d, which they output. Eve does not learn d because given any conversation, both values for d are equally likely.

Formally, let $\mathcal{P}_3 = (3, 1, \{0,1\}^2 \times \{0,1\}^2 \times \{0,1\}^3, \{0,1\} \cup \{0,1\}^2, \{0,1\}, \mu, \nu)$, where

$$\begin{aligned}
\mu_1((a_1, a_2), \epsilon) &= a_1 \\
\mu_2((b_1, b_2), \epsilon) &= b_2 \\
\mu_3((c_1, c_2, d), \epsilon) &= (c_1 \oplus d, c_2 \oplus d)
\end{aligned}$$

$$\begin{aligned}
\nu_1((a_1, a_2), (a, b, (e_1, e_2))) &= e_2 \oplus a_2 \oplus b \\
\nu_2((b_1, b_2), (a, b, (e_1, e_2))) &= e_1 \oplus a \oplus b_1 \\
\nu_3((c_1, c_2, d), (a, b, (e_1, e_2))) &= d
\end{aligned}$$

Let $\mathcal{S}_3$ be the canonical source for $U^{\mathcal{P}_3} = (\{0,1\}^2 \times \{0,1\}^2 \times \{0,1\}^3)$ such that

$$\Pr((a_1, a_2), (b_1, b_2), (c_1, c_2, d)) = \begin{cases} \frac{1}{32} & \text{if } c_1 = a_1 \oplus b_1 \text{ and } c_2 = a_2 \oplus b_2 \\ 0 & \text{otherwise} \end{cases}$$

Theorem 3.18 *$(\mathcal{P}_3)_{(\mathcal{S}_3)}$ performs perfect secret bit exchange against the empty view.*

Proof: By Corollary 3.13, it suffices to show that $\Pr[\mathrm{C}(\sigma)] = 2 \cdot \Pr[\mathrm{CO}(\sigma, v)]$ for all $\sigma \in \mathrm{cc}$ and $v \in \{0,1\}$. Let $\sigma \in \mathrm{cc}$ and $v \in \{0,1\}$. If $\Pr[\mathrm{C}(\sigma)] = 0$, then $\Pr[\mathrm{CO}(\sigma, v)] = 0$, so $\Pr[\mathrm{C}(\sigma)] = 2 \cdot \Pr[\mathrm{CO}(\sigma, v)]$, as desired. Otherwise, $\sigma = (a, b, (e_1, e_2))$ and for every $u = ((a_1, a_2), (b_1, b_2), (c_1, c_2, d)) \in \mathrm{feas}\,(\mathrm{C}(\sigma))$,

$$\begin{aligned}
a &= a_1 & \quad b &= b_2 \\
e_1 &= c_1 \oplus d & \quad e_2 &= c_2 \oplus d \\
c_1 &= a_1 \oplus b_1 & \quad c_2 &= a_2 \oplus b_2
\end{aligned}$$

Then for every $((a_1, a_2), (b_1, b_2), (c_1, c_2, d)) \in \mathrm{feas}\,(\mathrm{C}(\sigma))$,

$$\begin{aligned}
\nu_1((a_1, a_2), (a, b, (e_1, e_2))) &= e_2 \oplus a_2 \oplus b_2 = c_2 \oplus d \oplus c_2 = d \\
\nu_2((b_1, b_2), (a, b, (e_1, e_2))) &= e_1 \oplus a_1 \oplus b_1 = c_1 \oplus d \oplus c_1 = d \\
\nu_3((c_1, c_2, d), (a, b, (e_1, e_2))) &= d
\end{aligned}$$

Hence if $((a_1, a_2), (b_1, b_2), (c_1, c_2, d)) \in \mathrm{feas}\,(\mathrm{C}(\sigma))$, then $((a_1, a_2), (b_1, b_2), (c_1, c_2, d)) \in \mathrm{feas}\,(\mathrm{CO}(\sigma, d))$. Thus $\Pr[\mathrm{C}(\sigma)] = \Pr[\mathrm{CO}(\sigma, 0)] + \Pr[\mathrm{CO}(\sigma, 1)]$.

Furthermore, a_1, a_2, b_1, b_2, and d are independently and uniformly distributed over $\{0,1\}$. It is easily confirmed by brute force that a_1, b_2, $a_1 \oplus b_1 \oplus d$, $a_2 \oplus b_2 \oplus d$, and d are independently and uniformly distributed over $\{0,1\}$. In particular, this implies that the conversation $\sigma = (a_1, b_2, (a_1 \oplus b_1 \oplus d, a_2 \oplus b_2 \oplus d))$ is independent of the output d. In addition, $\Pr[\mathrm{O}(d)] = \frac{1}{2}$ for $d \in \{0,1\}$, so $\Pr[\mathrm{CO}(\sigma, d)] = \Pr[\mathrm{C}(\sigma) \cap \mathrm{O}(d)] = \frac{1}{2} \cdot \Pr[\mathrm{C}(\sigma)]$ for $d \in \{0,1\}$, as desired. ■

CHAPTER 4: THE CAPACITY OF A SOURCE

Theorem 3.17 shows that independent random inputs alone are not sufficient for secret key exchange. However, as seen by Example 1, there are cases where secret key exchange is not possible given correlated inputs alone, but is possible given independent random inputs in addition to the correlated inputs. Thus, independent random inputs can make the difference between possibility and impossibility of secret key exchange. In this chapter, we will investigate the ability of a source $\mathcal{T}$ plus additional independent random inputs to perform secret key exchange.

We imagine a scenario in which a protocol designer is given a source $\mathcal{T}$ of correlated inputs. The designer, whose goal is to obtain the largest secret key possible, is allowed to specify independent randomness (as modeled by a randomized extension of $\mathcal{T}$) and a protocol. For a given source $\mathcal{T}$, we define the *perfect capacity of* $\mathcal{T}$, denoted $\mathrm{pcap}(\mathcal{T})$ to be the maximum N such $\mathcal{T}$ is adequate for N-valued perfect secret key exchange. Since every source is adequate for 1-valued perfect secret key exchange, this maximum is well defined whenever there is an upper bound on N. If there is no upper bound, we take the perfect capacity of $\mathcal{T}$ to be infinite. We similarly define the weak capacity of $\mathcal{T}$, denoted $\mathrm{wcap}(\mathcal{T})$, with respect to weak secret key exchange. The following proposition is a consequence of Proposition 3.1.

Proposition 4.1 *Let $\mathcal{T}$ be a source. Then* $1 \leq \mathrm{pcap}(\mathcal{T}) \leq \mathrm{wcap}(\mathcal{T})$.

We will see shortly that the capacity of $\mathcal{T}$ is always finite. §4.2 shows that there exist sources $\mathcal{T}$ with $\mathrm{pcap}(\mathcal{T}) < \mathrm{wcap}(\mathcal{T})$ by showing that the source defined in §2.6 has a higher weak capacity than perfect capacity.

The following proposition, which follows from Propositions 2.6 and 3.2, states that team-equivalent sources have the same capacity.

Proposition 4.2 *Let $\mathcal{T}$ and $\mathcal{T}'$ be two team-equivalent sources. Then* $\mathrm{pcap}(\mathcal{T}) = \mathrm{pcap}(\mathcal{T}')$ *and* $\mathrm{wcap}(\mathcal{T}) = \mathrm{wcap}(\mathcal{T}')$.

It will be shown in §4.3 that for every $N \leq \mathrm{pcap}(\mathcal{T})$ (respectively, $N \leq \mathrm{wcap}(\mathcal{T})$), $\mathcal{T}$ is adequate for N-valued perfect (respectively weak) secret key exchange. Hence, the weak and perfect capacity of $\mathcal{T}$ can be interpreted as measures of the information provided to the players by $\mathcal{T}$. Theorem 3.17 implies that if $\mathcal{T}$ is a source in which the views of one set of team players are independent from the views of another set of team players, then $\mathrm{wcap}(\mathcal{T}) = \mathrm{pcap}(\mathcal{T}) = 1$.

4.1 Capacity Is Finite

We begin by showing some general properties of randomized $\mathcal{T}$-systems. Fix a source $\mathcal{T}$ with at least two players. Let $\mathcal{P}_{\mathcal{S}}$ be a randomized $\mathcal{T}$-system, let $q \in U^{\mathcal{T}}$, let $\tau \in \mathsf{c}$, and let $\ell = |\tau|$. We define

$$\mathrm{F}(q,\tau) = \{u \in \mathrm{feas}\left(U^{\mathcal{S}}\right) : \mathrm{corr}(u) = q \text{ and } \mathrm{conv}(u)[\ell] = \tau\}$$

Thus $\mathrm{F}(q,\tau)$ consists of the feasible inputs that have correlated part q and give rise to the conversation τ. We say q is *compatible* with τ (and τ is compatible with q) if $\mathrm{F}(q,\tau)$ is nonempty. We define $\mathrm{compat}(\tau) = \{q : \mathrm{F}(q,\tau) \neq \emptyset\}$, the set of q compatible with τ. Note that if q is compatible with τ, then q and τ are both feasible.

Let $q \in U^{\mathcal{T}}$ and $\sigma \in \mathsf{cc}$. If there exists a unique $v \in V$ such that $\mathrm{out}_i(u) = v$ for all $i \in K$ and $u \in \mathrm{F}(q,\sigma)$, then we define $\mathrm{out}(q,\sigma) = v$. Hence, if defined, $\mathrm{out}(q,\sigma)$ is the value that all team players output given any feasible input in which the correlated part is q and the conversation is σ.

Lemma 4.3 *Let $\mathcal{P}_{\mathcal{S}}$ be a randomized $\mathcal{T}$-system satisfying agreement such that $k \geq 2$, let $q \in U^{\mathcal{T}}$, and let $\sigma \in \mathsf{cc}$. If q is compatible with σ, then $\mathrm{out}(q,\sigma)$ is defined.*

Proof: Suppose the conditions of the lemma and suppose that q is compatible with σ. Then $\mathrm{F}(q,\sigma)$ is nonempty. Let $u, u' \in \mathrm{F}(q,\sigma)$ and let $i, i' \in K$. Let $v = \mathrm{out}_i(u)$ and $v' = \mathrm{out}_{i'}(u')$. It suffices to show that $v' = v$.

Since $u, u' \in \mathrm{F}(q,\sigma)$, it follows that u and u' are feasible and the correlated part $\mathrm{corr}(u) = \mathrm{corr}(u') = q$ is feasible. Let $\widehat{u}$ interpolate $\{u, u'\}$ such that $\widehat{u} =_i u$ and $\widehat{u} =_{i'} u'$. Then $\mathrm{corr}(\widehat{u}) = q$, so $\mathrm{corr}(\widehat{u})$ is feasible. Hence, by Lemma 2.5, $\widehat{u}$ is feasible. By Lemma 2.1 (part 3), $\mathrm{out}_i(\widehat{u}) = v$ and $\mathrm{out}_{i'}(\widehat{u}) = v'$. By Lemma 3.11 (part 2), $v' = v$. ■

Fix a complete conversation σ and two correlated parts $q, q' \in \mathrm{compat}(\sigma)$. If $q =_x q'$, then P_x can not determine whether the correlated part is q or q'. It follows that all the team players must output the same value in both situations.

Lemma 4.4 *Let $\mathcal{P}_{\mathcal{S}}$ be a randomized $\mathcal{T}$-system satisfying agreement such that $k \geq 2$, let $\sigma \in \mathsf{cc}$, and let $q, q' \in \mathrm{compat}(U^{\mathcal{T}})$. If q touches q', then $\mathrm{out}(q,\sigma) = \mathrm{out}(q',\sigma)$.*

Proof: Suppose the conditions of the lemma and suppose that $q =_x q'$ for some x. By Lemma 4.3, $\mathrm{out}(q,\sigma)$ and $\mathrm{out}(q',\sigma)$ are defined. Let $u \in \mathrm{F}(q,\sigma)$ and $u' \in \mathrm{F}(q',\sigma)$.

Let $\widehat{u}$ interpolate $\{u, u'\}$ such that $\widehat{u} =_x u$ and $\widehat{u} =_i u'$ for all $i \neq x$. Then $\mathrm{corr}(\widehat{u}) = q'$ since $\mathrm{corr}(\widehat{u}) =_x \mathrm{corr}(u) =_x q =_x q'$ and for $i \neq x$, $\mathrm{corr}(\widehat{u}) =_i \mathrm{corr}(u') =_i q'$; hence $\mathrm{corr}(\widehat{u})$ is feasible. By Lemma 2.5, $\widehat{u}$ is feasible. By Lemma 2.1 (part 1), $\mathrm{conv}(\widehat{u}) = \sigma$. It follows that $\widehat{u} \in \mathrm{F}(q', \sigma)$. By Lemma 2.1 (part 3), $\mathrm{out}_x(\widehat{u}) = \mathrm{out}_x(u)$. Furthermore, since $u \in \mathrm{F}(q, \sigma)$, $\mathrm{out}(q, \sigma) = \mathrm{out}_x(u)$. Similarly, since $\widehat{u} \in \mathrm{F}(q', \sigma)$, $\mathrm{out}(q', \sigma) = \mathrm{out}_x(\widehat{u})$. It follows that $\mathrm{out}(q, \sigma) = \mathrm{out}(q', \sigma)$. ■

Fix a complete conversation σ. It follows from the results of Section 3.2 that all output values must be possible when the conversation is σ. By the above, all players output the same value whenever P_i holds a given share. Therefore, the number of output values of a secret key exchange system is at most the number of feasible shares for any player P_i. Formally, we have the following.

Theorem 4.5 *Let $\mathcal{T}$ be a source. Then* $\mathrm{wcap}(\mathcal{T}) \leq \min_{i \in K} \left|\mathrm{feas}\left(U_i^{\mathcal{T}}\right)\right|$.

Proof: Suppose $\mathcal{P}_{\mathcal{S}}$ is a randomized $\mathcal{T}$-system for N-valued weak secret key exchange for some N and let $i \in K$. We complete the proof by showing $N \leq \left|\mathrm{feas}\left(U_i^{\mathcal{T}}\right)\right|$. Fix a feasible complete conversation σ. If $q, q' \in \mathrm{compat}(\sigma)$ and $\langle q \rangle_i = \langle q' \rangle_i$, then by Lemma 4.4, $\mathrm{out}(q, \sigma) = \mathrm{out}(q', \sigma)$. Thus,

$$
\begin{aligned}
& |\{v \in V : \exists q \in \mathrm{compat}(\sigma) \text{ s.t. } \mathrm{out}(q, \sigma) = v\}| \\
\leq \quad & |\{q_i \in U_i^{\mathcal{T}} : \exists q \in \mathrm{compat}(\sigma) \text{ s.t. } \langle q \rangle_i = q_i\}|
\end{aligned}
$$

Since by Corollary 3.15, $\Pr[\mathrm{CO}(\sigma, v)] > 0$ for all $v \in V$, we have

$$|\{v \in V : \exists q \in \mathrm{compat}(\sigma) \text{ s.t. } \mathrm{out}(q, \sigma) = v\}| = |V| = N$$

Since, in addition, $|\{q_i \in U_i^{\mathcal{T}} : \exists q \in \mathrm{compat}(\sigma) \text{ s.t. } \langle q \rangle_i = q_i\}| \leq \mathrm{feas}\left(U_i^{\mathcal{T}}\right)$, it follows that $N \leq \left|\mathrm{feas}\left(U_i^{\mathcal{T}}\right)\right|$. ■

Lemma 4.6 generalizes Lemma 2.1 (First Interpolation Lemma) to randomized $\mathcal{T}$-systems.

Lemma 4.6 (Third Interpolation Lemma) *Let $\mathcal{P}_{\mathcal{S}}$ be a randomized $\mathcal{T}$-system, let $\tau \in$* c*, let $U' \subseteq$* $\mathrm{compat}(\tau)$*, and suppose that $\widehat{q}$ is a feasible interpolant of U'. Let $(q_1, \ldots, q_k)$ be a U'-derivation of $\widehat{q}$. Then*

1. *$\widehat{q}$ is compatible with τ.*

2. *If $|\tau| < t$ and $m_1, \ldots, m_k$ are message vectors such that q_i is compatible with τm_i for $i \in K$, then $\widehat{q}$ is compatible with $\tau \cdot (\langle m_1 \rangle_1, \langle m_2 \rangle_2, \ldots, \langle m_k \rangle_k)$.*

3. *If $|\tau| = t$, $\mathcal{P}_{\mathcal{S}}$ satisfies agreement, and $k \geq 2$, then $\mathrm{out}(\widehat{q}, \tau) = \mathrm{out}(q_i, \tau)$ for $i \in K$.*

Proof: Suppose the conditions of the lemma and let $\ell = |\tau|$.

1. For $i \in K$, $q_i \in \mathrm{compat}(\tau)$, so there is some $u_i \in \mathrm{F}(q_i, \tau)$. Then $\widehat{u} = (\langle u_1 \rangle_1, \ldots, \langle u_k \rangle_k)$ interpolates $\{u_1, \ldots, u_k\}$ and $\mathrm{corr}(\widehat{u}) = \widehat{q}$. By Lemma 2.1 (part 1), $\mathrm{conv}(\widehat{u})[\ell] = \tau$ and by Lemma 2.5, $\widehat{u}$ is feasible. Thus, $\widehat{u} \in \mathrm{F}(\widehat{q}, \tau)$, so $\widehat{q}$ is compatible with τ.

2. Suppose that $\ell < t$ and q_i is compatible with τm_i for $i \in K$. Let $u_i \in \mathrm{F}(q_i, \tau m_i)$ for $i \in K$. Then $\widehat{u} = (\langle u_1 \rangle_1, \ldots, \langle u_k \rangle_k)$ interpolates $\{u_1, \ldots, u_k\}$ and $\mathrm{corr}(\widehat{u}) = \widehat{q}$. By Lemma 2.1 (part 2), $\mathrm{conv}(\widehat{u})[\ell + 1] = \tau \cdot (\langle m_1 \rangle_1, \langle m_2 \rangle_2, \ldots, \langle m_k \rangle_k)$ and by Lemma 2.5, $\widehat{u}$ is feasible. Thus, $\widehat{u} \in \mathrm{F}(\widehat{q}, \tau \cdot (\langle m_1 \rangle_1, \langle m_2 \rangle_2, \ldots, \langle m_k \rangle_k))$, so $\widehat{q}$ is compatible with $\tau \cdot (\langle m_1 \rangle_1, \langle m_2 \rangle_2, \ldots, \langle m_k \rangle_k)$.

3. Suppose $\ell = t$, $\mathcal{P}_{\mathcal{S}}$ satisfies agreement, $k \geq 2$ and $i \in K$. Since $(q_1, \ldots, q_k)$ is a U'-derivation of $\widehat{q}$, $\widehat{q} =_i q_i$. By part 1, $\widehat{q}$ is compatible with τ, so by Lemma 4.4, $\mathrm{out}(\widehat{q}, \tau) = \mathrm{out}(q_i, \tau)$. ∎

Lemma 4.7 *Let $\mathcal{P}_{\mathcal{S}}$ be a randomized $\mathcal{T}$-system satisfying agreement, let $\sigma \in \mathrm{cc}$, and let $U' \subseteq \mathrm{compat}(\sigma)$ be coverable. Then $\mathrm{out}(q, \sigma) = \mathrm{out}(q', \sigma)$ for all $q, q' \in U'$.*

Proof: Suppose the conditions of the lemma, let $q, q' \in U'$, and let $\widehat{q}$ be a covering of U'. Then $\widehat{q}$ is a feasible spanning interpolant of U'. It follows that there is a U'-derivation $(q_1, \ldots, q_k)$ of $\widehat{q}$ such that $q = q_i$ and $q' = q_{i'}$ for for some $i, i' \in K$. By Lemma 4.6 (part 3), $\mathrm{out}(q, \sigma) = \mathrm{out}(q_i, \sigma) = \mathrm{out}(\widehat{q}, \sigma) = \mathrm{out}(q_{i'}, \sigma) = \mathrm{out}(q', \sigma)$. ∎

4.2 Example 2, Final Visit

Recall the source $\mathcal{S}_2 = (\{0,1\}, \mathcal{F}, \{0,1\} \times \{0,1\}, \theta)$ introduced in §2.6, where $\Pr_{\mathcal{F}}(0) = 1/3$, $\Pr_{\mathcal{F}}(1) = 2/3$, and $\theta_i(\omega) = \omega$ for $i \in \{1,2\}$ and $\omega \in \{0,1\}$. $\mathcal{S}_2$ provides a shared biased coin flip to the team players. It follows from Theorem 3.10 that $\mathrm{wcap}(\mathcal{S}_2) \geq 2$. The following theorem shows that there is no way the team players can use independent randomization to obtain a perfectly secret bit from their shared biased coin flip. Thus, $\mathcal{S}_2$ is an example of a source that has a higher weak capacity than perfect capacity.

Theorem 4.8 $\text{pcap}(\mathcal{S}_2) = 1$.

Proof: Suppose by way of contradiction that there exists a protocol $\mathcal{P}$ and a randomized extension $\mathcal{S}$ of $\mathcal{S}_2$ such that $\mathcal{P}_{\mathcal{S}}$ performs 2-valued perfect secret key exchange. Let $\sigma \in \text{cc}$, let $i \in \{1, 2\}$, and suppose without loss of generality that $V = \{0, 1\}$. It follows from Lemma 4.4 that if two feasible inputs have the same correlated part and the same conversation, then they have the same output. Therefore, there are output values v and v' such that

$$\begin{aligned} \text{feas}\left(\mathrm{C}(\sigma) \cap \theta_i^{-1}(0)\right) &= \text{feas}\left(\mathrm{CO}(\sigma, v)\right) \\ \text{feas}\left(\mathrm{C}(\sigma) \cap \theta_i^{-1}(1)\right) &= \text{feas}\left(\mathrm{CO}(\sigma, v')\right) \end{aligned}$$

If $v = v'$, then $\Pr\left[\mathrm{C}(\sigma)\right] = \Pr\left[\left(\mathrm{C}(\sigma) \cap \theta_i^{-1}(0)\right) \cup \left(\mathrm{C}(\sigma) \cap \theta_i^{-1}(1)\right)\right] = \Pr\left[\mathrm{CO}(\sigma, v)\right]$, a contradiction to Corollary 3.13. Otherwise, $\{v, v'\} = \{0, 1\}$. By Corollary 3.13, it follows that $\Pr\left[\mathrm{C}(\sigma) \cap \theta_i^{-1}(0)\right] = \Pr\left[\mathrm{C}(\sigma) \cap \theta_i^{-1}(1)\right] = 1/2$. Summing over all $\sigma \in \text{cc}$ yields $\Pr\left[\theta_i^{-1}(0)\right] = \Pr\left[\theta_i^{-1}(1)\right]$, a contradiction. Hence $\text{pcap}(\mathcal{S}_2) < 2$, so $\text{pcap}(\mathcal{S}_2) = 1$. ■

4.3 Capacity as a Measure of Information

Let $\mathcal{T}$ be a source. We will show that for every $N \leq \text{pcap}(\mathcal{T})$ (respectively, $N \leq \text{wcap}(\mathcal{T})$), there is a randomized $\mathcal{T}$-system for N-valued perfect (respectively weak) secret key exchange. Hence in some sense, the capacity of a source is a measure of the shared secret information provided to the team by the source.

Let $N' \leq N$, let $V = \{0, \ldots, N-1\}$, and let $V' = \{0, \ldots, N'-1\}$. We first consider a simple N'-valued protocol $\mathcal{P}^{(N:N')}$, in which all the team players are given as input an N-valued secret key v and, in addition, P_1 is given a random value $v' \in \{1, \ldots, N'-1\}$. In this protocol, P_1 sends $x = (v' + v) \bmod N$. Each team player recovers v' by computing $(x - v) \bmod N$.

The value v is used as a one-time pad for v', although not in the usual sense where v and v' are bit strings of equal length. (Fischer, Paterson and Rackoff [FPR91] use such generalized one-time pads in the case where $v' \in \{0, 1\}$.) Unlike typical bit-string one-time pads, each fixed one-time pad $v \in V$ does not necessarily produce a uniformly distributed encrypted result $x \in V$ even if the cleartext message $v' \in V'$

is uniformly distributed. (For example, if $N = 3$, $N' = 2$, and $v = 0$, then $x \neq 2$.) However, like bit-string one-time pads, generalized one-time pads always produce a uniformly distributed encrypted result x for each particular cleartext message v', provided the one-time pad v is uniformly distributed. Hence, although sending x may reveal information about the pad v, it reveals nothing about the cleartext message v',

Note that if $N' > N$, this technique can not yield a secret key exchange system because the constructed protocol $\mathcal{P}'$ will never output values between N and N', so the constructed system can not satisfy perfect or weak uniformity.

Formally, $\mathcal{P}^{(N:N')} = (k, 1, (V \times V') \times V^{(k-1)}, V', V \cup \emptyset, \mu, \nu)$, where

$$
\begin{array}{lcll}
\mu_1((v,v'),\epsilon) & = & (v+v') \bmod N & \\
\mu_i(v,\epsilon) & = & \emptyset & \text{if } 2 \le i \le k
\end{array}
$$

$$
\begin{array}{lcll}
\nu_1((v,v'),\sigma) & = & v & \\
\nu_i(v,\sigma) & = & (\langle\sigma_1\rangle_1 - v) \bmod N & \text{if } 2 \le i \le k
\end{array}
$$

Theorem 4.9 *Let $\mathcal{S}^{(N:N')}$ be a source for $U^{\mathcal{P}^{(N:N')}}$ and suppose that* $\text{feas}\left(U^{\mathcal{P}^{(N:N')}}\right) = \{((v,v'),v,\ldots,v)\ : v \in V \text{ and } v' \in V'\}$.

1. *If* $\Pr(u) = \Pr(u')$ *for all* $u, u' \in \text{feas}\left(U^{\mathcal{P}^{(N:N')}}\right)$, *then* $(\mathcal{P}^{(N:N')})_{(\mathcal{S}^{(N:N')})}$ *performs N'-valued perfect secret key exchange.*
2. *If* $\Pr(u) > 0$ *for all* $u \in \text{feas}\left(U^{\mathcal{P}^{(N:N')}}\right)$, *then* $(\mathcal{P}^{(N:N')})_{(\mathcal{S}^{(N:N')})}$ *performs N'-valued weak secret key exchange.*

Proof: Suppose the conditions of the theorem. Then $\mathcal{P}^{(N:N')}$ is N'-valued. Let $u \in \text{feas}\left(U^{\mathcal{P}^{(N:N')}}\right)$. Then $u = ((v,v'),v,\ldots,v)$ for some $v \in V$ and $v' \in V'$. It follows that $\text{conv}(u) = ((v+v') \bmod N, \emptyset, \ldots, \emptyset)$ and $\text{out}_i(u) = v'$ for all $i \in K$. Then by Lemma 3.11 (parts 2 and 5),

$$\Pr[\mathrm{C}(\sigma)] = \sum_{v' \in V'} \Pr[\mathrm{CO}(\sigma, v')] \tag{4.1}$$

for all $\sigma \in \text{cc}$. Furthermore, for all $\sigma \in \text{cc}$ and $v' \in V'$, we have $\sigma = (x, \emptyset, \ldots, \emptyset)$ for some $x \in V$ and

$$\Pr[\mathrm{CO}(\sigma, v')] = \Pr((v,v'),v,\ldots,v) \tag{4.2}$$

where $v \in V$ is the unique value such that $(v+v') \bmod N = x$.

Suppose $\Pr(u) = \Pr(u')$ for all $u, u' \in \text{feas}\left(U^{\mathcal{P}^{(N:N')}}\right)$. Then $\Pr(u) = \frac{1}{N \cdot N'}$ for all $u \in \text{feas}\left(U^{\mathcal{P}^{(N:N')}}\right)$. Let $\sigma \in \mathsf{cc}$ and $v' \in V'$. If σ is not feasible, then $\Pr[\mathrm{C}(\sigma)] = N' \cdot \Pr[\mathrm{CO}(\sigma, v')] = 0$. Otherwise σ is feasible, so $\sigma = (x, \emptyset, \ldots, \emptyset)$ and $\Pr[\mathrm{C}(\sigma)] = \frac{1}{N}$. Let $v \in V$ be the unique value such that $(v + v') \bmod N = x$. Then by (4.2), $\Pr[\mathrm{C}(\sigma)] = \frac{1}{N} = N' \cdot \Pr((v, v'), v, \ldots, v) = N' \cdot \Pr[\mathrm{CO}(\sigma, v')]$. By Corollary 3.13, $(\mathcal{P}^{(N:N')})_{(\mathcal{S}^{(N:N')})}$ performs N'-valued perfect secret key exchange.

Now suppose $\Pr(u) > 0$ for all $u \in \text{feas}\left(U^{\mathcal{P}^{(N:N')}}\right)$. Let $\sigma \in \mathsf{cc}$ be feasible and let $v \in V$ be the unique value such that $(v + v') \bmod N = x$. Then by (4.2),$\Pr[\mathrm{CO}(\sigma, v')] = \Pr((v, v'), v, \ldots, v) > 0$. Since (4.1) holds, it follows from Corollary 3.15 that $(\mathcal{P}^{(N:N')})_{(\mathcal{S}^{(N:N')})}$ performs N'-valued weak secret key exchange. ■

Let $\mathcal{P}_{\mathcal{S}}$ be a randomized $\mathcal{T}$-system, and assume without loss of generality that $V = \{0, \ldots, N-1\}$. Let $N' \leq N$ and let $V' = \{0, \ldots, N'-1\}$. We create an N'-valued secret key exchange system $\mathcal{P}'_{\mathcal{S}'}$ by "composing" $\mathcal{P}_{\mathcal{S}}$ with $\mathcal{P}^{N:N'}$. That is, $\mathcal{P}'_{\mathcal{S}'}$ operates as follows. First it runs $\mathcal{P}_{\mathcal{S}}$ to exchange a key $v \in V$. Then P_1 chooses a random $v' \in V'$, and sends $x = (v' + v) \bmod N$. Each team player knows v, and recovers v' by computing $(x - v) \bmod N$.

Formally, $\mathcal{P}' = (k, t+1, U', M', V', \mu', \nu')$, where $U' = (U_1 \times V') \times U_2 \times \cdots \times U_k$ and $M' = M \cup V \cup \{\emptyset\}$. Let c', cc', and pc' denote the sets of conversations, complete conversations, and partial conversations of $\mathcal{P}'$, respectively. Let $a_i \in U_i$ for $1 \leq i \leq k$, let $v' \in V'$ and let $\tau \in \mathsf{c}$. The relevant portions of the message functions μ'_i are defined as follows.

$$\begin{array}{lcll}
\mu'_1((a_1, v'), \tau) & = & \mu_1(a_1, \tau) & \text{if } |\tau| < t \\
\mu'_i(a_i, \tau) & = & \mu_i(a_i, \tau) & \text{if } 2 \leq i \leq k \text{ and } |\tau| < t \\
\mu'_1((a_1, v'), \tau) & = & (v' + \nu_1(a_1, \tau)) \bmod N & \text{if } |\tau| = t \\
\mu'_i(a_i, \tau) & = & \emptyset & \text{if } 2 \leq i \leq k \text{ and } |\tau| = t
\end{array}$$

Let $\sigma \in \mathsf{cc}'$ such that $\sigma[t] \in \mathsf{c}$, $\langle \sigma_{(t+1)} \rangle_1 \in V'$ and $\langle \sigma_{t+1} \rangle_i = \emptyset$ for $2 \leq i \leq k$. The relevant portions of the output functions ν'_i are defined as follows.

$$\begin{array}{lcll}
\nu'_1((a_i, v'), \sigma) & = & v' & \text{if } |\sigma| = t+1 \\
\nu'_i(a_i, \sigma) & = & (\langle \sigma_{t+1} \rangle_1 - \nu_i(a_i, \sigma[t])) \bmod N & \text{if } 2 \leq i \leq k \text{ and } |\sigma| = t+1
\end{array}$$

Let $\mathcal{F}_{N'}$ be the uniform distribution on the set V'. To construct a source for $U^{\mathcal{P}'}$, we consider the global value set $\Omega' = \Omega \times V'$ and the distribution $\mathcal{F}' = \mathcal{F} \times \mathcal{F}_{N'}$ on

Ω'. Let $\omega' \in \Omega'$. Then $\omega' = (\omega, v')$ for some $\omega \in \Omega$ and $v' \in V'$. We define θ' by $\theta'_1(\omega') = (\theta_1(\omega), v')$ and $\theta'_i(\omega') = \theta_i(\omega)$ for $2 \leq i \leq k$. We let $\mathcal{S}' = (\Omega', \mathcal{F}', \theta')$.

Let $\theta_e : \Omega \to U_e$ be a view function over Ω for Eve and let $\omega' = (\omega, v) \in \Omega'$. We consider the view function $\theta'_e : \Omega' \to U_e$ over Ω' for Eve defined by $\theta'_e(\omega') = \theta_e(\omega)$.

Theorem 4.10 *Let $\mathcal{P}_\mathcal{S}$ be a randomized $\mathcal{T}$-system for N-valued perfect (weak) secret key exchange against θ_e and let $N' \leq N$. Let $\mathcal{P}'_{\mathcal{S}'}$ be the system constructed as above from $\mathcal{P}$, $\mathcal{S}$, and N', and let θ'_e be the view function for Eve as constructed above from θ_e. Then $\mathcal{P}'_{\mathcal{S}'}$ is randomized $\mathcal{T}$-system for N'-valued perfect (weak) secret key exchange against θ'_e.*

Proof (sketch): $\mathcal{P}'_{\mathcal{S}'}$ runs $\mathcal{P}_\mathcal{S}$, then uses the outputs of $\mathcal{P}_\mathcal{S}$ as the inputs to run $\mathcal{P}^{(N:N')}$. By assumption, $\mathcal{P}_\mathcal{S}$ performs perfect (weak) secret key exchange, and therefore provides the right kind of inputs to $\mathcal{P}^{(N:N')}$ for secret key exchange (i.e. the kind of inputs as specified in Theorem 4.9). In the composite protocol, Eve hears the conversation τ of $\mathcal{P}_\mathcal{S}$ before $\mathcal{P}^{(N:N')}$ is run. However, since $\mathcal{P}_\mathcal{S}$ performs perfect (weak) secret key exchange, the conversation τ of $\mathcal{P}_\mathcal{S}$ is (weakly) independent of the output value v of $\mathcal{P}_\mathcal{S}$, so by Theorem 3.7, ${\mathcal{P}^{(N:N')}}_{\mathcal{S}(N:N')}$ performs perfect (weak) secret key exchange even if Eve hears the conversation τ. It follows that the composite system $\mathcal{P}'_{\mathcal{S}'}$ performs N'-valued perfect (weak) secret key exchange. ■

4.4 An Upper Bound on Perfect Capacity

Fischer, Paterson, and Rackoff [FPR91] show that for two players dealt hands of cards, perfect secret bit exchange is not possible if the probability that randomly chosen hands from the deck overlap is too large. Fischer and Wright [FW92,FW] generalize this result to arbitrarily large teams dealt hands of cards to obtain an upper bound on the perfect capacity of deals of cards. Here, we further generalize these results to obtain an upper bound on the perfect capacity of any source. This bound is an improvement over the bound implied by Theorem 4.5.

Fix a source $\mathcal{T}$. We define

$$\psi_\mathcal{T} = \left(\prod_{i=1}^{k} |\mathrm{feas}\,(U_i)\,| \right) \cdot \max_{q \in \mathrm{feas}(U)} \Pr(q)$$

$\psi_\mathcal{T}$ measures the distance between the distribution induced on the players' inputs by $\mathcal{T}$ and the uniform distribution on the players' inputs. The uniform distribution on the players' inputs results in the players receiving statistically independent inputs, so the players receive no shared information in this case. In some sense, the smaller $\psi_{\mathcal{T}}$ is, the less shared information a random legal deal contains for the team players. This is made precise in Theorem 4.12 below.

We will need the following lemma about real numbers. It is proved using the arithmetic and geometric means inequality (AGM), which says that if a_1 through a_m are nonnegative, then $\sqrt[m]{\prod_{i=1}^{m} a_i} \leq \left(\sum_{i=1}^{m} a_i\right)/m$.

Lemma 4.11 *Let x_i^j be nonnegative for $1 \leq i \leq p$ and $1 \leq j \leq q$. Then*

$$\min_{j\in\{1,\ldots,q\}} \left(\prod_{i=1}^{p} x_i^j\right) \leq \frac{1}{q^p} \prod_{i=1}^{p} \sum_{j=1}^{q} x_i^j$$

Proof: Let x_i^j be nonnegative for $1 \leq i \leq p$ and $1 \leq j \leq q$. Then

$$\min_{j\in\{1,\ldots,q\}} \left(\prod_{i=1}^{p} x_i^j\right) \leq \sqrt[q]{\prod_{j=1}^{q} \prod_{i=1}^{p} x_i^j} \tag{4.3}$$

$$= \prod_{i=1}^{p} \sqrt[q]{\prod_{j=1}^{q} x_i^j} \tag{4.4}$$

$$\leq \prod_{i=1}^{p} \left(\frac{\sum_{j=1}^{q} x_i^j}{q}\right) \tag{4.5}$$

$$= \frac{1}{q^p} \prod_{i=1}^{p} \sum_{j=1}^{q} x_i^j \tag{4.6}$$

Here, (4.3) holds because the qth root of the product of q positive numbers is always at least as big as the smallest of the numbers. (4.5) is by the AGM. (4.4) and (4.6) are direct algebraic manipulation. ∎

Theorem 4.12 *Let $\mathcal{T}$ be a source. Then* $\mathrm{pcap}(\mathcal{T}) \leq \left\lfloor \sqrt[k-1]{\psi_{\mathcal{T}}} \right\rfloor$.

Proof: Let $\mathcal{T}$ be a source and let $N = \mathrm{pcap}(\mathcal{T})$. We show that $N \leq \left\lfloor \sqrt[k-1]{\psi_{\mathcal{T}}} \right\rfloor$.

By the definition of perfect capacity, there exists a randomized $\mathcal{T}$-system $\mathcal{P}_{\mathcal{S}}$ for N-valued perfect secret key exchange. Let $\mathcal{S}_1$ be a canonical source that is team-equivalent to $\mathcal{S}$. (Such a source exists by Proposition 2.2.) It follows from Proposition 2.6 that $\mathcal{P}_{\mathcal{S}_1}$ is a randomized $\mathcal{T}$-system and it follows from Proposition 3.2 that $\mathcal{P}_{\mathcal{S}_1}$ performs N-valued perfect secret key exchange. Hence $U^{\mathcal{S}_1} = U^{\mathcal{P}} = (R_1 \times U_1^{\mathcal{T}}) \times \ldots \times (R_k \times U_k^{\mathcal{T}})$ for some $R_1, \ldots, R_k$, and for $u \in U^{\mathcal{P}}$,

$$\Pr_{\mathcal{S}_1}(u) = \Pr_{\mathcal{T}}(\mathrm{corr}(u)) \cdot \prod_{i=1}^{k} \Pr_{\mathcal{R}_i}(\langle \mathrm{ind}(u) \rangle_i)$$

where $\mathcal{R}_i$ is the distribution $\mathcal{S}_1$ induces on R_i.

We construct another canonical source $\mathcal{S}_2$ for $U^{\mathcal{P}}$ in which the distribution of the independent parts is the same as in $\mathcal{S}_1$, but the distribution of the correlated parts is different. Specifically, we let $\mathcal{S}_2$ be a canonical source for $U^{\mathcal{P}}$ such that for every $u \in U^{\mathcal{P}}$,

$$\Pr_{\mathcal{S}_2}(u) = \begin{cases} \prod_{i=1}^{k} \dfrac{\Pr_{\mathcal{R}_i}(\langle \text{ind}(u)\rangle_i)}{|\text{feas}\,(U_i^{\mathcal{T}})\,|} & \text{if } \text{corr}(u) \in \text{feas}\left(U_1^{\mathcal{T}}\right) \times \cdots \times \text{feas}\left(U_k^{\mathcal{T}}\right) \\ 0 & \text{otherwise} \end{cases}$$

(Such a source exists by Proposition 2.2.) Then $\Omega^{\mathcal{S}_1} = \Omega^{\mathcal{S}_2} = U^{\mathcal{P}}$.

Let $u \in U^{\mathcal{P}}$. Recall that

$$\psi_{\mathcal{T}} = \left(\prod_{i=1}^{k} |\text{feas}\,(U_i)\,|\right) \cdot \max_{q \in \text{feas}(U)} \Pr(q)$$

If $\text{corr}(u) \notin \text{feas}\left(U^{\mathcal{T}}\right)$, then $\Pr_{\mathcal{S}_1} = 0$, so trivially $\Pr_{\mathcal{S}_1}(u) \leq \psi_{\mathcal{T}} \Pr_{\mathcal{S}_2}(u)$. Otherwise, $\text{corr}(u) \in \text{feas}\left(U^{\mathcal{T}}\right) \subseteq \text{feas}\left(U_1^{\mathcal{T}}\right) \times \cdots \times \text{feas}\left(U_k^{\mathcal{T}}\right)$, so $\Pr_{\mathcal{S}_2}(u) = \prod_{i=1}^{k} \frac{\Pr_{\mathcal{R}_i}(\langle \text{ind}(u)\rangle_i)}{|\text{feas}\left(U_i^{\mathcal{T}}\right)|}$. Then

$$\begin{aligned} \Pr_{\mathcal{S}_1}(u) &= \Pr_{\mathcal{T}}(\text{corr}(u)) \cdot \prod_{i=1}^{k} \Pr_{\mathcal{R}_i}(\langle \text{ind}(u)\rangle_i) \\ &\leq \max_{q \in \text{feas}(U^{\mathcal{T}})} \Pr(q) \cdot \prod_{i=1}^{k} \Pr_{\mathcal{R}_i}(\langle \text{ind}(u)\rangle_i) \\ &= \psi_{\mathcal{T}} \Pr_{\mathcal{S}_2}(u) \end{aligned}$$

Hence $\Pr_{\mathcal{S}_1}(u) \leq \psi_{\mathcal{T}} \Pr_{\mathcal{S}_2}(u)$ for every $u \in U^{\mathcal{P}}$. Thus, for any event $X \in U^{\mathcal{P}}$,

$$\Pr_{\mathcal{S}_1}[X] \leq \psi_{\mathcal{T}} \Pr_{\mathcal{S}_2}[X] \tag{4.7}$$

Let $\mathcal{Q}_i$ be the distribution induced on $U_i^{\mathcal{T}}$ by $\mathcal{S}_2$. It follows from the construction of $\mathcal{S}_2$ that $\Pr_{\mathcal{Q}_i}(q_i) = 1/|\text{feas}\left(U_i^{\mathcal{T}}\right)|$ if $q_i \in \text{feas}\left(U_i^{\mathcal{T}}\right)$ and $\Pr_{\mathcal{Q}_i}(q_i) = 0$ if $q_i \notin \text{feas}\left(U_i^{\mathcal{T}}\right)$. Hence

$$\begin{aligned} \Pr_{\mathcal{S}_2}(u) &= \prod_{i=1}^{k} \left(\Pr_{\mathcal{R}_i}(\langle \text{ind}(u)\rangle_i) \cdot \Pr_{\mathcal{Q}_i}(\langle \text{corr}(u)\rangle_i)\right) \\ &= \prod_{i=1}^{k} \Pr_{(\mathcal{R}_i \times \mathcal{Q}_i)}(\langle u\rangle_i) \end{aligned}$$

That is, $\mathcal{F}^{\mathcal{S}_2} = (\mathcal{R}_1 \times \mathcal{Q}_1) \times \cdots \times (\mathcal{R}_k \times \mathcal{Q}_k)$. Since $\theta_i^{\mathcal{S}_2}(u) = \langle u\rangle_i$, each $\theta_i^{\mathcal{S}_2}$ depends only on the i^{th} component of its argument. It follows that the random variables $\theta_1^{\mathcal{S}_2}, \ldots, \theta_k^{\mathcal{S}_2}$ are independent.

Let $\sigma \in \mathrm{cc}$. The event $\overline{\mathrm{C}}_i(\sigma)$ respects the random variable $\theta_i^{\mathcal{S}_2}$ for $i \in K$. It follows that the events $\overline{\mathrm{C}}_1(\sigma), \ldots, \overline{\mathrm{C}}_k(\sigma)$ are independent. Hence, by Lemma 3.8,

$$\Pr_{\mathcal{S}_2}[\mathrm{C}(\sigma)] = \prod_{i=1}^{k} \Pr_{\mathcal{S}_2}\left[\overline{\mathrm{C}}_i(\sigma)\right] \tag{4.8}$$

Similarly, if $v \in V$, then the event $\overline{\mathrm{CO}}_i(\sigma, v)$ respects $\theta_i^{\mathcal{S}_2}$. Hence, for any $v \in V$, the events $\overline{\mathrm{CO}}_1(\sigma, v), \ldots, \overline{\mathrm{CO}}_k(\sigma, v)$ are independent. By Lemma 3.9,

$$\Pr_{\mathcal{S}_2}[\mathrm{CO}(\sigma, v)] = \prod_{i=1}^{k} \Pr_{\mathcal{S}_2}\left[\overline{\mathrm{CO}}_i(\sigma, v)\right] \tag{4.9}$$

Therefore,

$$\begin{aligned}
\Pr_{\mathcal{S}_1}[\mathrm{C}(\sigma)] &= N \min_{v \in V}\left(\Pr_{\mathcal{S}_1}[\mathrm{CO}(\sigma, v)]\right) && (4.10)\\
&\leq N \min_{v \in V}\left(\psi_{\mathcal{T}} \Pr_{\mathcal{S}_2}[\mathrm{CO}(\sigma, v)]\right) && (4.11)\\
&= N \min_{v \in V}\left(\psi_{\mathcal{T}} \prod_{i=1}^{k} \Pr_{\mathcal{S}_2}\left[\overline{\mathrm{CO}}_i(\sigma, v)\right]\right) && (4.12)\\
&\leq \frac{N}{N^k} \cdot \psi_{\mathcal{T}} \prod_{i=1}^{k}\left(\sum_{v \in V} \Pr_{\mathcal{S}_2}\left[\overline{\mathrm{CO}}_i(\sigma, v)\right]\right) && (4.13)\\
&= \frac{\psi_{\mathcal{T}}}{N^{k-1}} \prod_{i=1}^{k} \Pr_{\mathcal{S}_2}\left[\overline{\mathrm{C}}_i(\sigma)\right] && (4.14)\\
&= \frac{\psi_{\mathcal{T}}}{N^{k-1}} \Pr_{\mathcal{S}_2}[\mathrm{C}(\sigma)] && (4.15)
\end{aligned}$$

Here, (4.10) follows from Corollary 3.13, (4.11) follows from (4.7), (4.12) follows from (4.9), (4.13) follows from Lemma 4.11, (4.14) follows from the fact that $\overline{\mathrm{C}}_i(\sigma)$ is the disjoint union over $v \in V$ of $\overline{\mathrm{CO}}_i(\sigma, v)$, and (4.15) follows from (4.8).

Summing over all complete conversations yields

$$1 = \sum_{\sigma \in \mathrm{cc}} \Pr_{\mathcal{S}_1}[\mathrm{C}(\sigma)] \leq \sum_{\sigma \in \mathrm{cc}} \frac{\psi_{\mathcal{T}}}{N^{k-1}} \Pr_{\mathcal{S}_2}[\mathrm{C}(\sigma)] = \frac{\psi_{\mathcal{T}}}{N^{k-1}}$$

Since N is an integer, it follows immediately that

$$N \leq \left\lfloor \sqrt[k-1]{\psi_{\mathcal{T}}} \right\rfloor$$

as desired. ∎

We say a source $\mathcal{T}$ is *uniform* if $\mathrm{feas}\left(U_i^{\mathcal{T}}\right) = U_i^{\mathcal{T}}$ for $i \in K$ and $\Pr_{\mathcal{T}}(u) = \Pr_{\mathcal{T}}(u')$ for all $u, u' \in \mathrm{feas}\left(U^{\mathcal{T}}\right)$. (Note that this implies that $U^{\mathcal{T}} = \mathrm{feas}\left(U_1^{\mathcal{T}}\right) \times \cdots \times \mathrm{feas}\left(U_k^{\mathcal{T}}\right)$. However, it does not imply $\mathrm{feas}\left(U^{\mathcal{T}}\right) = U^{\mathcal{T}}$, since there may be inputs that are not feasible even though each player's component of the input is feasible.)

Corollary 4.13 *Let $\mathcal{T}$ be a uniform source. Then*

$$\mathrm{pcap}(\mathcal{T}) \leq \left\lfloor \sqrt[k-1]{\frac{|U^{\mathcal{T}}|}{|\mathrm{feas}\left(U^{\mathcal{T}}\right)|}} \right\rfloor$$

Proof: Let $\mathcal{T}$ be a uniform source. Then $\Pr(q) = 1/|\mathrm{feas}\left(U^{\mathcal{T}}\right)|$ for all $q \in \mathrm{feas}\left(U^{\mathcal{T}}\right)$. Hence,

$$\psi_{\mathcal{T}} = \prod_{i=1}^{k} |\mathrm{feas}\left(U_i^{\mathcal{T}}\right)| \cdot \max_{q \in \mathrm{feas}(U^{\mathcal{T}})} \Pr(q) = \frac{|U^{\mathcal{T}}|}{|\mathrm{feas}\left(U^{\mathcal{T}}\right)|}$$

By Theorem 4.12, $\mathrm{pcap}(\mathcal{T}) \leq \left\lfloor \sqrt[k-1]{|U^{\mathcal{T}}|/|\mathrm{feas}\left(U^{\mathcal{T}}\right)|} \right\rfloor$. ∎

It follows from a result of Maurer ([Mau93], Corollary 1) that for sources $\mathcal{T}$ with two team players, $\mathrm{pcap}(\mathcal{T}) \leq \left\lfloor 2^{I(U_1^{\mathcal{T}};U_2^{\mathcal{T}})} \right\rfloor$, where $I(X;Y) = -\sum_x \Pr(x) \log \Pr(x) - \sum_{x,y} \Pr(x,y) \log \Pr(x|y)$ denotes the mutual information of X and Y. (See [Gal68] for a general reference on information theory and mutual information.) For the case $k=2$, the bound given by Corollary 4.13 for uniform source coincides with Maurer's result since in this case the mutual information $I(U_1^{\mathcal{T}}; U_2^{\mathcal{T}}) = \log(|U^{\mathcal{T}}|/|\mathrm{feas}\left(U^{\mathcal{T}}\right)|)$. For nonuniform sources with two team players, Maurer's result generally provides a tighter bound than Theorem 4.12. (For example, for the source $\mathcal{S}_3$ of Example 2, Theorem 4.12 yields the bound $\mathrm{pcap}(\mathcal{S}_3) \leq \left\lfloor 2 \cdot 2 \cdot \frac{2}{3} \right\rfloor = 2$, while Maurer's result yields the bound $\mathrm{pcap}(\mathcal{S}_3) \leq \left\lfloor 2^{(-\frac{1}{3}\log\frac{1}{3} - \frac{2}{3}\log\frac{2}{3})} \right\rfloor = 1$.) In addition, for two team players, Maurer's result is more general than our result because it takes into account the information available to Eve and can handle imperfect secrecy as well as perfect secrecy. It is not clear, however, how to generalize Maurer's work to more than two team players because it is not clear how to generalize the concept of mutual information, which is symmetric for the two players, to more than two players.

CHAPTER 5: CARD GAMES

5.1 Card Game Systems

In this section, we formalize the use of deals of cards as correlated random variables. A *deck* Δ is a finite set, whose elements we call *cards*; a *hand* is subset of Δ. A *deal* $\delta = (h_1, \ldots, h_k)$ is a sequence of hands, one for each player. (Note that there may be cards in the deck that don't appear in any hand.) The deal δ is *legal* if $h_i \cap h_j = \emptyset$ for $i \neq j$. A deal that may or may not be legal is called a *general* deal. In the real world, where all hands are typically dealt from a single deck of cards, all deals are legal. General deals are of interest to us because they can arise when legal deals are interpolated.

A *signature*[1] $(s_1, s_2, \ldots, s_k; d)$, where $s_1, \ldots, s_k$ and d are nonnegative integers such that $d \geq \sum s_i$, describes the number k of players, the size s_i of each player's hand, and the number d of cards in the deck. If all k team players have the same hand size s in the signature ξ, we write $\xi = (s^{(k)}; d)$. Given a signature, we define $e = d - \sum s_i$, so e is the number of cards not dealt to any player in a legal deal.

Let $\xi = (s_1, \ldots, s_k; d)$ be a signature. Without loss of generality, we always fix the deck $\Delta^\xi = \{1, \ldots, d\}$. A ξ*-deal* is a deal $\delta = (h_1, \ldots, h_k)$ such that $|h_i| = s_i$ for $i \in K$. We define $H_i^\xi = \{h : h \subseteq \Delta^\xi \text{ and } |h| = s_i\}$, so H_i^ξ is the set of possible hands for P_i. We write L^ξ to denote the set of legal ξ-deals and D^ξ to denote the set of general ξ-deals, so $L^\xi \subseteq D^\xi = H_1^\xi \times \cdots \times H_k^\xi$.

A source $\mathcal{T}$ for D^ξ is *legal* if $\Pr_{\mathcal{T}}(\delta) = 1/|L^\xi|$ for $\delta \in L^\xi$, and $\Pr_{\mathcal{T}}(\delta) = 0$ for $\delta \in D^\xi - L^\xi$. Hence $\mathcal{T}$ is a legal source if it assigns zero probability to all illegal deals and equal probability to all legal deals. We say $\mathcal{S}$ is a *card game source* if $\mathcal{S}$ is a randomized extension of some legal source for D^ξ. We say a protocol $\mathcal{P}$ is a ξ*-protocol* if $U^{\mathcal{P}} = (R_1 \times H_1^\xi) \times \cdots \times (R_k \times H_k^\xi)$ for some sets $R_1, \ldots, R_k$. We say a system $\mathcal{P}_{\mathcal{S}}$ is a *card game* ξ*-system* if $\mathcal{S}$ is a card game source and $\mathcal{P}$ is a ξ-protocol. If $U = D^\xi$ (i.e. all the R_i are singleton sets), we say $\mathcal{P}_{\mathcal{S}}$ is *deterministic*.

For $i \in K$ and $x \in \Delta$, we define the event

$$\mathrm{B}_i(x) = \{\omega : x \in \mathrm{corr}(\theta_i(\omega))\}$$

Thus $\mathrm{B}_i(x)$ is the event that card x belongs to P_i.

[1]This term is borrowed from algebra, and is not intended to have any connection to digital signatures.

Since all legal sources for D^ξ are team-equivalent, it follows from Proposition 4.2 that they all have the same perfect capacity. We denote this capacity by $\mathrm{pcap}(\xi)$. Similarly, we denote the weak capacity of all legal sources for D^ξ by $\mathrm{wcap}(\xi)$. Hence, if $N \leq \mathrm{pcap}(\xi)$ (respectively, if $N \leq \mathrm{wcap}(\xi)$) then there exists a card game ξ-system $\mathcal{P}_\mathcal{S}$ for N-valued perfect (respectively, weak) secret key exchange.

By Proposition 4.1 and Theorem 4.5, $1 \leq \mathrm{pcap}(\xi) \leq \mathrm{wcap}(\xi) \leq \min_{i \in K} |H_i^\xi|$. Since a legal source is uniform, we have the following improved bound on perfect capacity.

Theorem 5.1 *Let* $\xi = (s_1, \ldots, s_k; d)$. *Then* $\mathrm{pcap}(\xi) \leq \left\lfloor \sqrt[k-1]{|D^\xi|/|L^\xi|} \right\rfloor$.

Proof: Let $\mathcal{T}$ be a legal source for D^ξ. Then $\mathcal{T}$ is uniform since $\mathrm{feas}\left(U_i^\mathcal{T}\right) = H_i^\xi = U_i^\mathcal{T}$ for $i \in K$ and $\Pr_\mathcal{T}(u) = \Pr_\mathcal{T}(u') = 1/|L^\xi|$ for all $u, u' \in \mathrm{feas}\left(U^\mathcal{T}\right)$. By Corollary 4.13, $\mathrm{pcap}(\mathcal{T}) \leq \left\lfloor \sqrt[k-1]{|U^\mathcal{T}|/|\mathrm{feas}\left(U^\mathcal{T}\right)|} \right\rfloor = \left\lfloor \sqrt[k-1]{|D^\xi|/|L^\xi|} \right\rfloor$. ∎

Calculating $|D^\xi|/|L^\xi|$, we can apply Theorem 5.1 to obtain an upper bound on the perfect capacity of any signature. An example follows.

Corollary 5.2 $\mathrm{pcap}(2,2,2;7) \leq 3$.

Proof: Let $\xi = (2,2,2;7)$. Then $k = 3$, $|D^\xi| = \binom{7}{2}^3 = 9261$ and $|L^\xi| = \binom{7}{2}\binom{5}{2}\binom{3}{2} = 630$. By Theorem 5.1, $\mathrm{pcap}(2,2,2;7) \leq \left\lfloor \sqrt[2]{9261/630} \right\rfloor = 3$. ∎

In Chapter 8, we obtain tight bounds on $\mathrm{wcap}(1^{(k)}; k)$ and $\mathrm{pcap}(1^{(k)}; k)$ for $k \geq 3$ by considering the particular structure of the set of $(1^{(k)}; k)$-deals.

In the remainder of this dissertation, we will be interested in two views for Eve: the empty view, as already defined, and the *extra cards view*, defined for any randomized extension $\mathcal{S}$ of a source for D^ξ as $\theta_e(\omega) = \Delta^\xi - \bigcup_{i \in K} \mathrm{corr}(\theta_i(\omega))$. Thus, the extra cards view gives Eve the cards that are not dealt to any team player. Hence in a legal deal, Eve is dealt $d - \sum s_i$ cards; we always denote this quantity by e. Since Eve's view in the extra cards view depends only on the players' inputs, secrecy against the extra cards view is not affected by the underlying source, but only by the distribution on the players' inputs. Hence the following proposition holds.

Proposition 5.3 *Perfect secrecy and weak secrecy against the extra cards view respect team-equivalence.*

5.2 Card Game Techniques

There are several techniques we will use in designing card game protocols based on simulation of other protocols or recursion to smaller cases. These techniques result in a "new" deal of a "new" signature dealt from a "new" deck. During the execution of a protocol, we say a set S of cards is *discarded* from the deck if all team players agree to play as if the cards in S are no longer part of the deck. This can be carried out only if the set S and the number of cards each team player holds in S are common knowledge among the team players. Generally, common knowledge will be achieved by some player announcing the set S and each team player announcing the number of cards she holds in S. The resulting deck has the cards from S removed, the resulting signature has each player's hand size decreased by the number of cards she holds in S, and each player's hand has any cards from S removed. All the team players continue as if their new hands were dealt from the new deck, yielding a deal of the new signature.

Another useful technique is that of *dropping out*. In this case, a team player P_i announces all of the cards in her current hand h_i of size s_i. These cards are discarded. The resulting deck has the set h_i removed and the new signature has one fewer player (P_i is no longer there). P_i no longer participates in the protocol.

Finally, additional cards can be *added* to the deck. Effectively, these cards are added to Eve's hand. If the deck consists of the cards $1, \ldots, d$, the cards $d+1, \ldots, d'$ can be added to the deck to increase the deck size to d'. Then, in order to preserve the uniformity of the deal, the cards are renamed according to a random permutation. In this case, the resulting deck is $\{1, \ldots, d'\}$. No team player's hand size changes.

Let $\xi = (s_1, \ldots, s_k; d)$ and $\xi' = (s'_1, \ldots, s'_{k'}; d')$ be two signatures, and let $e = d - \sum s_i$ and $e' = d' - \sum s'_i$. We say ξ' *dominates* ξ if $k = k'$, $e' \leq e$, and there is a permutation π on $\{1, \ldots, k\}$ such that $s'_i \geq s_{\pi(i)}$ for each i. Note that this does not imply an ordering on d and d'. We write $\xi' \trianglerighteq \xi$.

The following theorem captures the intuition that giving the team players additional cards or giving Eve fewer cards can only help the team.

Theorem 5.4 (Domination Theorem) *Let $\xi' \trianglerighteq \xi$. If there is an N-valued card game ξ-system for perfect (weak) secret key exchange against the extra cards view, then there is an N-valued card game ξ'-system for perfect (weak) secret key exchange against the extra cards view.*

Proof (sketch): Let $\xi = (s_1, \ldots, s_k; d)$ and $\xi' = (s'_1, \ldots, s'_{k'}; d')$ be two signatures, and let $e = d - \sum s_i$ and $e' = d' - \sum s'_i$. Suppose that $\xi' \trianglerighteq \xi$ via some permutation π and suppose that there is a card game ξ-system $\mathcal{P}_{\mathcal{S}}$ for N-valued perfect (weak) secret key exchange against the extra cards view.

We construct a card game ξ'-system $\mathcal{P}'_{\mathcal{S}'}$ for N-valued perfect (weak) secret key exchange against the extra cards view as follows. First, each team player P_i randomly chooses $s'_i - s_{\pi(i)}$ cards from her hand and announces these cards, which are then discarded from the deck. Second, $e - e'$ additional cards are added to the deck. The players then simulate $\mathcal{P}_{\mathcal{S}}$ with the resulting deal, where player P_i in $\mathcal{P}'_{\mathcal{S}'}$ simulates player $P_{\pi(i)}$ in $\mathcal{P}_{\mathcal{S}}$.

In the simulation of $\mathcal{P}_{\mathcal{S}}$, Eve has been given some additional information beyond what she would normally be given as her view of the input to $\mathcal{P}_{\mathcal{S}}$. Namely, she has heard the team players discard some of their cards. However, since this information is independent of the resulting input, it follows from Theorem 3.7 that the simulation results in an N-valued perfect (weak) secret key. Hence, $\mathcal{P}'_{\mathcal{S}'}$ is the desired card game ξ'-system for N-valued perfect (weak) secret key exchange against the empty view. ■

5.3 Example 1, Revisited

In Example 1 (§1.1), we looked at the deck $\{\mathrm{J}, \mathrm{Q}, \mathrm{K}, \mathrm{A}\}$ and the signature $(2, 1; 4)$. In this section, we formalize the results described there. Formally, Alice will be P_1 and Bob will be P_2. We identify the deck $\{\mathrm{J}, \mathrm{Q}, \mathrm{K}, \mathrm{A}\}$ with the deck $\Delta = \{1, 2, 3, 4\}$ by considering $\mathrm{J} = 1$, $\mathrm{Q} = 2$, $\mathrm{K} = 3$, and $\mathrm{A} = 4$. We denote the hand $\{x\}$ containing the single card x simply by x. Let L denote the set of legal $(2, 1; 4)$-deals of Δ, let D denote the set of general $(2, 1; 4)$-deals of Δ, and let H_i denote the set of possible hands for P_i . Then $D = H_1 \times H_2$. Let ϵ denote the empty conversation. The message $\emptyset$ will be used as a meaningless message, and can be thought of as saying "pass."

Recall the deterministic protocol $\mathcal{P}_1$, in which Alice says whichever of the following statements is true:

- "I hold $\{\mathrm{J}, \mathrm{Q}\}$ or $\{\mathrm{K}, \mathrm{A}\}$."
- "I hold $\{\mathrm{J}, \mathrm{K}\}$ or $\{\mathrm{Q}, \mathrm{A}\}$."
- "I hold $\{\mathrm{J}, \mathrm{A}\}$ or $\{\mathrm{Q}, \mathrm{K}\}$."

Then Alice and Bob both output the truth value of the statement "Alice holds J."

Let the messages $c_1 =$ "I hold $\{\mathrm{J}, \mathrm{Q}\}$ or $\{\mathrm{K}, \mathrm{A}\}$," $c_2 =$ "I hold $\{\mathrm{J}, \mathrm{K}\}$ or $\{\mathrm{Q}, \mathrm{A}\}$," and $c_3 =$"I hold $\{\mathrm{J}, \mathrm{A}\}$ or $\{\mathrm{Q}, \mathrm{K}\}$." Denote the conversation $((c_i, \emptyset))$ by σ_i for $1 \leq i \leq 3$. The protocol $\mathcal{P}_1$ can be defined formally as

$$\mathcal{P}_1 = (2, 1, D, \{c_1, c_2, c_3, \emptyset\}, \{0, 1\}, (\mu_1, \mu_2), (\nu_1, \nu_2))$$

The message functions μ_1 and μ_2 are defined by

$$\begin{aligned} \mu_1(\{\mathrm{J}, \mathrm{Q}\}, \epsilon) &= \mu_1(\{\mathrm{K}, \mathrm{A}\}, \epsilon) &= c_1 \\ \mu_1(\{\mathrm{J}, \mathrm{K}\}, \epsilon) &= \mu_1(\{\mathrm{Q}, \mathrm{A}\}, \epsilon) &= c_2 \\ \mu_1(\{\mathrm{J}, \mathrm{A}\}, \epsilon) &= \mu_1(\{\mathrm{Q}, \mathrm{K}\}, \epsilon) &= c_3 \end{aligned}$$

$$\mu_2(\mathrm{J}, \epsilon) = \mu_2(\mathrm{Q}, \epsilon) = \mu_2(\mathrm{K}, \epsilon) = \mu_2(\mathrm{A}, \epsilon) = \emptyset$$

and the relevant portions of the output functions ν_1 and ν_2 are defined by

$$\begin{aligned} \nu_1(\{\mathrm{J}, \mathrm{Q}\}, \sigma_i) &= \nu_1(\{\mathrm{J}, \mathrm{K}\}, \sigma_i) = \nu_1(\{\mathrm{J}, \mathrm{A}\}, \sigma_i) = 1 \quad \text{for } 1 \leq i \leq 3 \\ \nu_1(\{\mathrm{Q}, \mathrm{K}\}, \sigma_i) &= \nu_1(\{\mathrm{Q}, \mathrm{A}\}, \sigma_i) = \nu_1(\{\mathrm{K}, \mathrm{A}\}, \sigma_i) = 0 \quad \text{for } 1 \leq i \leq 3 \end{aligned}$$

$$\begin{aligned} \nu_2(\mathrm{J}, \sigma_1) &= \nu_2(\mathrm{Q}, \sigma_1) = \nu_2(\mathrm{J}, \sigma_2) = 0 \\ \nu_2(\mathrm{K}, \sigma_2) &= \nu_2(\mathrm{J}, \sigma_3) = \nu_2(\mathrm{A}, \sigma_3) = 0 \\ \nu_2(\mathrm{K}, \sigma_1) &= \nu_2(\mathrm{A}, \sigma_1) = \nu_2(\mathrm{Q}, \sigma_2) = 1 \\ \nu_2(\mathrm{A}, \sigma_2) &= \nu_2(\mathrm{Q}, \sigma_3) = \nu_2(\mathrm{K}, \sigma_3) = 1 \end{aligned}$$

Note that to be strictly complete, it would be necessary to define the output functions for the infeasible conversations, such as $((c_1, c_2))$, and for infeasible combinations of hand and conversation, such as (K, σ_1) . However, since the output values resulting for infeasible inputs do not affect the secrecy, agreement, and uniformity conditions, we have left them out.

Let $\mathcal{S}_1$ be the canonical legal source for D.

Theorem 5.5 *$(\mathcal{P}_1)_{(\mathcal{S}_1)}$ is a deterministic card game $(2, 1; 4)$-system for perfect secret bit exchange against the empty view.*

Proof: Since $U^{\mathcal{P}_1} = D$, $(\mathcal{P}_1)_{(\mathcal{S}_1)}$ is a deterministic card game $(2, 1, ; 4)$-system. Since $|V^{\mathcal{P}_1}| = 2$, $(\mathcal{P}_1)_{(\mathcal{S}_1)}$ is 2-valued. Figure 5.1 depicts the input space D for $(\mathcal{P}_1)_{(\mathcal{S}_1)}$ and the outcomes for each feasible input. Each entry of the matrix corresponds to a $(2, 1; 4)$-deal of Δ, where the rows are labeled by P_1's hands and the columns are labeled by P_2's hands. An entry is marked by a dash if the pair of hands does not

$h_1 \setminus h_2$	J	Q	K	A
{J, Q}	–	–	$\sigma_1, 1, 1$	$\sigma_1, 1, 1$
{J, K}	–	$\sigma_2, 1, 1$	–	$\sigma_2, 1, 1$
{J, A}	–	$\sigma_3, 1, 1$	$\sigma_3, 1, 1$	–
{Q, K}	$\sigma_3, 0, 0$	–	–	$\sigma_3, 0, 0$
{Q, A}	$\sigma_2, 0, 0$	–	$\sigma_2, 0, 0$	–
{K, A}	$\sigma_1, 0, 0$	$\sigma_1, 0, 0$	–	–

Figure 5.1: Possible outcomes of $(\mathcal{P}_1)_{(\mathcal{S}_1)}$

form a feasible input (i.e. if the hands do not form a legal deal). Otherwise, the entry corresponds to a feasible input u, and is marked by the conversation $\mathrm{conv}(u)$, P_1's output $\mathrm{out}_1(u)$, and P_2's output $\mathrm{out}_2(u)$. There are 12 equally likely feasible inputs.

Inspection of the possible outcomes shows that $\Pr[\mathrm{C}(\sigma)] = \frac{1}{3}$ for all feasible $\sigma \in \mathsf{cc}$ and $\Pr[\mathrm{C}(\sigma) \cap \mathrm{O}(v)] = \frac{1}{6}$ for all feasible $\sigma \in \mathsf{cc}$ and $v \in \{0,1\}$. Hence, let $\sigma \in \mathsf{cc}$ and $v \in \{0,1\}$ if $\sigma \in \mathsf{cc}$ is feasible, then $2 \cdot \Pr[\mathrm{C}(\sigma) \cap \mathrm{O}(v)] = \frac{1}{3} = \Pr[\mathrm{C}(\sigma)]$. If σ is not feasible, then $2 \cdot \Pr[\mathrm{C}(\sigma) \cap \mathrm{O}(v)] = 0 = \Pr[\mathrm{C}(\sigma)]$. In either case $2 \cdot \Pr[\mathrm{C}(\sigma) \cap \mathrm{O}(v)] = \Pr[\mathrm{C}(\sigma)]$, so by Corollary 3.13, $(\mathcal{P}_1)_{(\mathcal{S}_1)}$ is a 2-valued perfect secret key exchange system. It follows that $(\mathcal{P}_1)_{(\mathcal{S}_1)}$ is a deterministic card game $(2,1;4)$-system for secret bit exchange against the empty view. ■

Corollary 5.6 $\mathrm{pcap}(2,1;4) = 2$.

Proof: It follows from Theorem 5.5 that $\mathrm{pcap}(2,1;4) \geq 2$. $|D^{(2,1;4)}| = \binom{4}{2}\binom{4}{1} = 24$ and $|L^{(2,1;4)}| = \binom{4}{2}\binom{2}{1} = 12$, so by Theorem 5.1, $\mathrm{pcap}(2,1;4) \leq 24/12 = 2$. Hence $\mathrm{pcap}(2,1;4) = 2$. ■

In the randomized protocol $\mathcal{P}_1'$, Alice chooses a random card in her hand and a random card not in her hand and announces these two cards. Bob then says whether or not he holds one of the two cards. If Bob holds one of the two announced cards, then Alice and Bob both know which of the two cards Alice holds and which Bob holds, and they can use this information to agree on a secret bit. If Bob does not hold one of the two cards, then Alice and Bob know which of the two remaining cards Alice holds and which Bob holds, and again they can agree on a secret bit. In order to choose a card in her hand and a card not in her hand, Alice needs to be given as part of her input the outcome of two coin flips.

$$
\begin{aligned}
c_0 &= \text{“So do I.”} \\
c_1 &= \text{“I hold exactly one of the cards in the set } \{\mathrm{J}, \mathrm{Q}\}.\text{”} \\
c_2 &= \text{“I hold exactly one of the cards in the set } \{\mathrm{J}, \mathrm{K}\}.\text{”} \\
c_3 &= \text{“I hold exactly one of the cards in the set } \{\mathrm{J}, \mathrm{A}\}.\text{”} \\
c_4 &= \text{“I hold exactly one of the cards in the set } \{\mathrm{Q}, \mathrm{K}\}.\text{”} \\
c_5 &= \text{“I hold exactly one of the cards in the set } \{\mathrm{Q}, \mathrm{A}\}.\text{”} \\
c_6 &= \text{“I hold exactly one of the cards in the set } \{\mathrm{K}, \mathrm{A}\}.\text{”}
\end{aligned}
$$

Interpretations of messages of $\mathcal{P}_1'$

$$
\begin{aligned}
\sigma_{1a} &= ((c_1, \emptyset), (\emptyset, c_0)) & \sigma_{1r} &= ((c_1, \emptyset), (\emptyset, \emptyset)) \\
\sigma_{2a} &= ((c_2, \emptyset), (\emptyset, c_0)) & \sigma_{2r} &= ((c_2, \emptyset), (\emptyset, \emptyset)) \\
\sigma_{3a} &= ((c_3, \emptyset), (\emptyset, c_0)) & \sigma_{3r} &= ((c_3, \emptyset), (\emptyset, \emptyset)) \\
\sigma_{4a} &= ((c_4, \emptyset), (\emptyset, c_0)) & \sigma_{4r} &= ((c_4, \emptyset), (\emptyset, \emptyset)) \\
\sigma_{5a} &= ((c_5, \emptyset), (\emptyset, c_0)) & \sigma_{5r} &= ((c_5, \emptyset), (\emptyset, \emptyset)) \\
\sigma_{6a} &= ((c_6, \emptyset), (\emptyset, c_0)) & \sigma_{6r} &= ((c_6, \emptyset), (\emptyset, \emptyset))
\end{aligned}
$$

Feasible conversations of $\mathcal{P}_1'$

Figure 5.2: Messages and conversations of $\mathcal{P}_1'$

Let $M = \{c_i : 1 \le i \le k\} \cup \{\emptyset\}$. The protocol $\mathcal{P}_1'$ can be formally defined as

$$\mathcal{P}_1' = (2, 2, (\{00, 01, 10, 11\} \times H_1) \times H_2, M, \{0, 1\}, (\mu_1, \mu_2), (\nu_1, \nu_2))$$

The interpretations of the messages in M and the conversations that will be feasible when $\mathcal{Q}_1$ is used with a card game source are shown in Figure 5.2.

If τ is any conversation of length one and $h_1 \in H_1$ is a hand for P_1, then $\mu_1((r, h_1), \tau) = \emptyset$; $\mu_1((r, h_1), \epsilon)$ is defined as shown in Figure 5.3. Each row is labeled by P_1's hand h_1, each column is labeled by an outcome r of P_1's coin flips, and entry is $\mu_1((r, h_1), \epsilon)$. If $h_2 \in H_2$, then $\mu_2(h_2, \epsilon) = \emptyset$; $\mu_2(h_2, \tau)$ for $|\tau| = 1$ is defined as shown in Figure 5.3. Each row is labeled by P_2's hand h_2, each column is labeled by a message c, and the entry is $\mu_2(h_2, ((c, \emptyset)))$. The relevant portions of ν_1 and ν_2 are shown in Figure 5.3. An entry is marked by a dash if the specified hand and conversation cannot simultaneously occur when $\mathcal{P}_1'$ is used with a card game source.

Let $\mathcal{R}_1$ be the uniform distribution over $\{00, 01, 10, 11\}$, and let $\mathcal{S}_1'$ be $\mathcal{S}_1$ extended by $\mathcal{R}_1$ (where $\mathcal{S}_1$ is the canonical card game source for D as above).

$h_1 \setminus r$	00	01	10	11
{J, Q}	c_2	c_3	c_4	c_5
{J, K}	c_1	c_3	c_4	c_6
{J, A}	c_1	c_2	c_5	c_6
{Q, K}	c_1	c_5	c_2	c_6
{Q, A}	c_1	c_4	c_3	c_6
{K, A}	c_2	c_4	c_3	c_5

Messages sent by P_1: $\mu_1((r, h_1), \epsilon)$.

$h_2 \setminus c$	c_1	c_2	c_3	c_4	c_5	c_6
J	c_0	c_0	c_0	$\emptyset$	$\emptyset$	$\emptyset$
Q	c_0	$\emptyset$	$\emptyset$	c_0	c_0	$\emptyset$
K	$\emptyset$	c_0	$\emptyset$	c_0	$\emptyset$	c_0
A	$\emptyset$	$\emptyset$	c_0	$\emptyset$	c_0	c_0

Messages sent by P_2: $\mu_2(h_2, ((c, \emptyset)))$.

$h_1 \setminus \sigma$	σ_{1a}	σ_{1r}	σ_{2a}	σ_{2r}	σ_{3a}	σ_{3r}	σ_{4a}	σ_{4r}	σ_{5a}	σ_{5r}	σ_{6a}	σ_{6r}
{J, Q}	–	–	1	1	1	1	1	1	1	1	–	–
{J, K}	1	1	–	–	1	0	0	1	–	–	1	1
{J, A}	1	0	1	0	–	–	–	–	0	1	0	1
{Q, K}	0	1	0	1	–	–	–	–	1	0	1	0
{Q, A}	0	0	–	–	0	1	1	0	–	–	0	0
{K, A}	–	–	0	0	0	0	0	0	0	0	–	–

P_1's outputs: $\nu_1((h_1, r), \sigma)$ (for all $r \in R_1$).

$h_2 \setminus \sigma$	σ_{1a}	σ_{1r}	σ_{2a}	σ_{2r}	σ_{3a}	σ_{3r}	σ_{4a}	σ_{4r}	σ_{5a}	σ_{5r}	σ_{6a}	σ_{6r}
J	0	–	0	–	0	–	–	0	–	0	–	0
Q	1	–	–	0	–	0	0	–	0	–	–	1
K	–	0	1	–	–	1	1	–	–	1	0	–
A	–	1	–	1	1	–	–	1	1	–	1	–

P_2's outputs: $\nu_2(h_2, \sigma)$.

Figure 5.3: Message and output functions of $\mathcal{P}_1'$

$\theta_e(u) = \mathrm{J}$			
$\langle u\rangle_1 \setminus \langle u\rangle_2$	Q	K	A
$\{\mathrm{Q},\mathrm{K}\},00$	–	–	$\sigma_{1r},1,1$
$\{\mathrm{Q},\mathrm{K}\},01$	–	–	$\sigma_{5a},1,1$
$\{\mathrm{Q},\mathrm{K}\},10$	–	–	$\sigma_{2r},1,1$
$\{\mathrm{Q},\mathrm{K}\},11$	–	–	$\sigma_{6a},1,1$
$\{\mathrm{Q},\mathrm{A}\},00$	–	$\sigma_{1r},0,0$	–
$\{\mathrm{Q},\mathrm{A}\},01$	–	$\sigma_{4a},1,1$	–
$\{\mathrm{Q},\mathrm{A}\},10$	–	$\sigma_{3r},1,1$	–
$\{\mathrm{Q},\mathrm{A}\},11$	–	$\sigma_{6a},0,0$	–
$\{\mathrm{K},\mathrm{A}\},00$	$\sigma_{2r},0,0$	–	–
$\{\mathrm{K},\mathrm{A}\},01$	$\sigma_{4a},0,0$	–	–
$\{\mathrm{K},\mathrm{A}\},10$	$\sigma_{3r},0,0$	–	–
$\{\mathrm{K},\mathrm{A}\},11$	$\sigma_{5a},0,0$	–	–

$\theta_e(u) = \mathrm{Q}$			
$\langle u\rangle_1 \setminus \langle u\rangle_2$	J	K	A
$\{\mathrm{J},\mathrm{K}\},00$	–	–	$\sigma_{1r},1,1$
$\{\mathrm{J},\mathrm{K}\},01$	–	–	$\sigma_{3a},1,1$
$\{\mathrm{J},\mathrm{K}\},10$	–	–	$\sigma_{4r},1,1$
$\{\mathrm{J},\mathrm{K}\},11$	–	–	$\sigma_{6a},1,1$
$\{\mathrm{J},\mathrm{A}\},00$	–	$\sigma_{1r},0,0$	–
$\{\mathrm{J},\mathrm{A}\},01$	–	$\sigma_{2a},1,1$	–
$\{\mathrm{J},\mathrm{A}\},10$	–	$\sigma_{5r},1,1$	–
$\{\mathrm{J},\mathrm{A}\},11$	–	$\sigma_{6a},0,0$	–
$\{\mathrm{K},\mathrm{A}\},00$	$\sigma_{2a},0,0$	–	–
$\{\mathrm{K},\mathrm{A}\},01$	$\sigma_{4r},0,0$	–	–
$\{\mathrm{K},\mathrm{A}\},10$	$\sigma_{3a},0,0$	–	–
$\{\mathrm{K},\mathrm{A}\},11$	$\sigma_{5r},0,0$	–	–

$\theta_e(u) = \mathrm{K}$			
$\langle u\rangle_1 \setminus \langle u\rangle_2$	J	Q	A
$\{\mathrm{J},\mathrm{Q}\},00$	–	–	$\sigma_{2r},1,1$
$\{\mathrm{J},\mathrm{Q}\},01$	–	–	$\sigma_{3a},1,1$
$\{\mathrm{J},\mathrm{Q}\},10$	–	–	$\sigma_{4r},1,1$
$\{\mathrm{J},\mathrm{Q}\},11$	–	–	$\sigma_{5a},1,1$
$\{\mathrm{J},\mathrm{A}\},00$	–	$\sigma_{1a},1,1$	–
$\{\mathrm{J},\mathrm{A}\},01$	–	$\sigma_{2r},0,0$	–
$\{\mathrm{J},\mathrm{A}\},10$	–	$\sigma_{5a},0,0$	–
$\{\mathrm{J},\mathrm{A}\},11$	–	$\sigma_{6r},1,1$	–
$\{\mathrm{Q},\mathrm{A}\},00$	$\sigma_{1a},0,0$	–	–
$\{\mathrm{Q},\mathrm{A}\},01$	$\sigma_{4r},0,0$	–	–
$\{\mathrm{Q},\mathrm{A}\},10$	$\sigma_{3a},0,0$	–	–
$\{\mathrm{Q},\mathrm{A}\},11$	$\sigma_{6r},0,0$	–	–

$\theta_e(u) = \mathrm{A}$			
$\langle u\rangle_1 \setminus \langle u\rangle_2$	J	Q	K
$\{\mathrm{J},\mathrm{Q}\},00$	–	–	$\sigma_{2a},1,1$
$\{\mathrm{J},\mathrm{Q}\},01$	–	–	$\sigma_{3r},1,1$
$\{\mathrm{J},\mathrm{Q}\},10$	–	–	$\sigma_{4a},1,1$
$\{\mathrm{J},\mathrm{Q}\},11$	–	–	$\sigma_{5r},1,1$
$\{\mathrm{J},\mathrm{K}\},00$	–	$\sigma_{1a},1,1$	–
$\{\mathrm{J},\mathrm{K}\},01$	–	$\sigma_{3r},0,0$	–
$\{\mathrm{J},\mathrm{K}\},10$	–	$\sigma_{4a},0,0$	–
$\{\mathrm{J},\mathrm{K}\},11$	–	$\sigma_{6r},1,1$	–
$\{\mathrm{Q},\mathrm{K}\},00$	$\sigma_{1a},0,0$	–	–
$\{\mathrm{Q},\mathrm{K}\},01$	$\sigma_{5r},0,0$	–	–
$\{\mathrm{Q},\mathrm{K}\},10$	$\sigma_{2a},0,0$	–	–
$\{\mathrm{Q},\mathrm{K}\},11$	$\sigma_{6r},0,0$	–	–

Figure 5.4: Possible outcomes of $(\mathcal{P}_1')_{(\mathcal{S}_1')}$

Theorem 5.7 $(\mathcal{P}_1')_{(\mathcal{S}_1')}$ *is a card game* $(2,1;4)$*-protocol for perfect secret bit exchange against the extra cards view.*

Proof: Let θ_e be the extra cards view. Figure 5.4 depicts the input space $U^{\mathcal{P}_1'}$ for $(\mathcal{P}_1')_{(\mathcal{S}_1')}$ and the possible outcomes. Each row is labeled by P_1's hand and coin flips; each column is labeled by P_2's hand. Eve's hand is shown at the top of each table. An entry is marked by a dash if it does not correspond to a feasible input (i.e. if the hands do not form a legal deal). Otherwise, the entry corresponds to a feasible input u, and is marked by the conversation $\mathrm{conv}(u)$, P_1's output $\mathrm{out}_1(u)$, and P_2's output $\mathrm{out}_2(u)$). There are 48 equally likely feasible inputs.

Case by case examination shows each conversation occurs at four feasible inputs, which are divided such that Eve has one view for two of the inputs and another view for the remaining two inputs. Of each of the two inputs with the same conversation and the same view for Eve, one has output 0 for both players and the other has output 1 for both players. That is, let $v \in \{0, 1\}$. If the conversation σ and the view V for Eve are simultaneously possible, then

$$2 \cdot \Pr\left[\mathrm{E}(V) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)\right] = \frac{2}{48} = \Pr\left[\mathrm{E}(V) \cap \mathrm{C}(\sigma)\right]$$

If σ and V are not simultaneously possible, then

$$2 \cdot \Pr\left[\mathrm{E}(V) \cap \mathrm{C}(\sigma) \cap \mathrm{O}(v)\right] = 0 = \Pr\left[\mathrm{E}(V) \cap \mathrm{C}(\sigma)\right]$$

By Lemma 3.12 (part 3), $(\mathcal{P}_1')_{(\mathcal{S}_1')}$ is a 2-valued perfect secret key exchange system against θ_e. It follows that $(\mathcal{P}_1')_{(\mathcal{S}_1')}$ is a card game $(2, 1; 4)$-protocol for perfect secret bit exchange against the extra cards view. ■

Protocol $\mathcal{P}_1'$ is a simplified example of a *key set protocol.* Key set protocols are a family of secret bit exchange protocols that will be presented in §6.1.

Finally, we show there is no deterministic card game $(2, 1; 4)$-system for secret bit exchange against the extra cards view.

Theorem 5.8 *There is no deterministic card game* $(2, 1; 4)$*-system for weak secret bit exchange against the extra cards view.*

Proof: Let θ_e be the extra cards view and suppose by way of contradiction that $\mathcal{P}_\mathcal{S}$ is a deterministic card game $(2, 1; 4)$-system for weak secret bit exchange against θ_e. Assume without loss of generality that $\mathcal{S}$ is canonical. Since $\mathcal{S}$ is a deterministic legal $(2, 1; 4)$-source for D, there are twelve feasible inputs $u_1 = (\{\mathrm{K}, \mathrm{A}\}, \mathrm{Q})$, $u_2 = (\{\mathrm{Q}, \mathrm{A}\}, \mathrm{K})$, $u_3 = (\{\mathrm{Q}, \mathrm{K}\}, \mathrm{A})$, $u_4 = (\{\mathrm{K}, \mathrm{A}\}, \mathrm{J})$, $u_5 = (\{\mathrm{J}, \mathrm{A}\}, \mathrm{K})$, $u_6 = (\{\mathrm{J}, \mathrm{K}\}, \mathrm{A})$, $u_7 = (\{\mathrm{Q}, \mathrm{A}\}, \mathrm{J})$, $u_8 = (\{\mathrm{J}, \mathrm{A}\}, \mathrm{Q})$, $u_9 = (\{\mathrm{J}, \mathrm{Q}\}, \mathrm{A})$, $u_{10} = (\{\mathrm{Q}, \mathrm{A}\}, \mathrm{J})$, $u_{11} = (\{\mathrm{J}, \mathrm{A}\}, \mathrm{Q})$, and $u_{12} = (\{\mathrm{J}, \mathrm{Q}\}, \mathrm{A})$. The possible outcomes of $\mathcal{P}_\mathcal{S}$ are shown in Figure 5.5, where an entry is marked by a dash if it does not correspond to a feasible input. Each feasible entry u_i is marked by the conversation $\sigma_i = \mathrm{conv}(u_i)$, P_1's output $v_i = \mathrm{out}_1(u_i)$, and P_2's output $y_i = \mathrm{out}_2(u_i)$. It follows that for $1 \leq i \leq 12$, $\Pr(u_i) = 1/12$, σ_i is a feasible complete conversation, $v_i \in V$, and $y_i \in V$. Furthermore, since agreement is satisfied and $\Pr\left[\mathrm{O}_1(v_i) \cap \mathrm{O}_2(y_i)\right] > 0$, it follows that $v_i = y_i$.

$\theta_e(u) = \mathrm{J}$			
$h_1 \setminus h_2$	Q	K	A
$\{\mathrm{Q},\mathrm{K}\}$	–	–	σ_3, v_3, y_3
$\{\mathrm{Q},\mathrm{A}\}$	–	σ_2, v_2, y_2	–
$\{\mathrm{K},\mathrm{A}\}$	σ_1, v_1, y_1	–	–

$\theta_e(u) = \mathrm{Q}$			
$h_1 \setminus h_2$	J	K	A
$\{\mathrm{J},\mathrm{K}\}$	–	–	σ_6, v_6, y_6
$\{\mathrm{J},\mathrm{A}\}$	–	σ_5, v_5, y_5	–
$\{\mathrm{K},\mathrm{A}\}$	σ_4, v_4, y_4	–	–

$\theta_e(u) = \mathrm{K}$			
$h_1 \setminus h_2$	J	Q	A
$\{\mathrm{J},\mathrm{Q}\}$	–	–	σ_9, v_9, y_9
$\{\mathrm{J},\mathrm{A}\}$	–	σ_8, v_8, y_8	–
$\{\mathrm{Q},\mathrm{A}\}$	σ_7, v_7, y_7	–	–

$\theta_e(u) = \mathrm{A}$			
$h_1 \setminus h_2$	J	Q	K
$\{\mathrm{J},\mathrm{Q}\}$	–	–	$\sigma_{12}, v_{12}, y_{12}$
$\{\mathrm{J},\mathrm{K}\}$	–	$\sigma_{11}, v_{11}, y_{11}$	–
$\{\mathrm{Q},\mathrm{K}\}$	$\sigma_{10}, v_{10}, y_{10}$	–	–

Figure 5.5: Possible outcomes of $\mathcal{P}_{\mathcal{S}}$

Claim 1: $\sigma_i = \sigma_{i+1} = \sigma_{i+2}$ for $i = 1, 4, 7, 10$. To see this, let $i \in \{1, 4, 7, 10\}$ and suppose that σ_{i+1}, σ_{i+2}, and σ_{i+3} are not all equal. Then there is some $j \in \{i, i+1, i+2\}$ such that $\sigma_j \neq \sigma_{j'}$ for $j' \in \{i, i+1, i+2\} - \{j\}$. Note that $\theta_e(u_i) = \theta_e(u_{i+1}) = \theta_e(u_{i+2})$. Let $v = \mathrm{out}_1(u_j)$, and let $v' = 1 - v$. Then $\Pr[\mathrm{E}(\theta_e(u_j)) \cap \mathrm{C}(\sigma_j) \cap \mathrm{O}_1(v')] = 0$. However, $\Pr[\mathrm{E}(\theta_e(u_j)) \cap \mathrm{C}(\sigma_j)] = 1/12 > 0$ and, since weak uniformity is satisfied, $\Pr[\mathrm{O}_1(v')] > 0$, a contradiction to the weak secrecy condition. Hence Claim 1 is true.

Claim 2: $\sigma_i = \sigma_j$ for $1 \leq i, j \leq 12$. To see this, suppose as an induction hypothesis that at a given round $\ell < t$, τ is the only feasible conversation. Then τ is a prefix of each σ_i. (Note that this is true in the base case $\ell = 0$, when the empty conversation ϵ is the only feasible conversation). We will show that there is a statement m such that τm is the only feasible conversation at round $\ell + 1$. By Claim 1, $\langle(\sigma_i)_{\ell+1}\rangle_j = \langle(\sigma_{i+1})_{\ell+1}\rangle_j = \langle(\sigma_{i+2})_{\ell+1}\rangle_j$ for $i = 1, 4, 7, 10$ and $j = 1, 2$.

Since there is only one feasible partial conversation at round ℓ, Alice must send the same message in round $\ell + 1$ whenever she holds the same hand. In particular, $\langle(\sigma_1)_{\ell+1}\rangle_1 = \mu_1(\{\mathrm{K},\mathrm{A}\}, \tau) = \langle(\sigma_4)_{\ell+1}\rangle_1$, $\langle(\sigma_2)_{\ell+1}\rangle_1 = \mu_1(\{\mathrm{Q},\mathrm{A}\}, \tau) = \langle(\sigma_7)_{\ell+1}\rangle_1$, and $\langle(\sigma_3)_{\ell+1}\rangle_1 = \mu_1(\{\mathrm{Q},\mathrm{K}\}, \tau) = \langle(\sigma_{10})_{\ell+1}\rangle_1$. It follows from Claim 1 and transitivity that there is some $c_1 \in M$ such that $\langle(\sigma_i)_{\ell+1}\rangle_1 = c_1$ for $1 \leq i \leq 12$. Similarly, Bob must send the same message in round $\ell + 1$ whenever he holds the same hand, so $\langle(\sigma_1)_{\ell+1}\rangle_2 = \langle(\sigma_8)_{\ell+1}\rangle_2 = \langle(\sigma_{11})_{\ell+1}\rangle_2 = \mu_2(\{\mathrm{Q}\}, \tau)$ and $\langle(\sigma_2)_{\ell+1}\rangle_2 = \langle(\sigma_5)_{\ell+1}\rangle_2 = \langle(\sigma_{12})_{\ell+1}\rangle_2 = \mu_2(\{\mathrm{K}\}, \tau)$, so there is some $c_2 \in M$ such that $\langle(\sigma_i)_{\ell+1}\rangle_2 = c_2$ for $1 \leq i \leq 12$. Thus, $\tau \cdot (c_1, c_2)$ is the only feasible conversation at round $\ell + 1$. Thus by induction, there is exactly one feasible conversation at each round, and in particular, there is only one feasible complete conversation σ, so Claim 2 is proved.

It follows that $y_1 = y_8 = y_{11} = \nu_2(\{\mathrm{Q}\}, \sigma)$. Similarly, $y_2 = y_5 = y_{12}$, $y_3 = y_6 = y_9$, $y_4 = y_7 = y_{10}$, $v_1 = v_4$, $v_2 = v_7$, and $v_3 = v_{10}$. Since in addition $v_i = y_i$ for $1 \leq i \leq 12$, it follows by transitivity that there is some $v \in V$ such that $v_i = v_j = v$ for all $1 \leq i, j \leq 12$. Let v' be the other element of V. Then $\Pr[\mathrm{O}_1(v')] = 0$, a contradiction to weak uniformity. It follows that there is no deterministic weak secret bit exchange $(2, 1; 4)$-protocol. ∎

CHAPTER 6: KEY SET PROTOCOLS

In this chapter, we present and analyze the class of key set protocols. Key set protocols, defined in §6.1, are simple 1-bit perfect secret key exchange protocols against the extra cards view. We use a game-theoretic argument to exhibit an optimal key set protocol in §6.2

6.1 The Class of Key Set Protocols

The building block of a key set protocol is the notion of an *opaque key set*, originally developed by Winkler [Win81a,Win81b,Win83] for use in the game of bridge, and later used by Fischer, Paterson, and Rackoff [FPR91] for secret bit exchange between two players. A key set X consists of two cards, one held by a team player P, the other held by a *different* team player Q. We say X is *shared* by P and Q. A key set $X = \{x, y\}$ is *opaque* if, given the information available to Eve, it is equally likely that P holds x and Q holds y or that P holds y and Q holds x. We will define opaque key sets formally in §6.1.2. Once P and Q determine a shared opaque key set X, they can use it as a one-time pad to transmit a bit secretly. For example, to transmit a bit v to Q, P says whichever of the following statements is true.

- "I hold the lower card in X if and only if $v = 0$."
- "I hold the lower card in X if and only if $v = 1$."

During the run of the protocol, the team players are thought of as the nodes of a labeled graph called the *key set graph*. Two players P and Q are connected by an edge labeled X if P and Q share the opaque key set X. The protocol obtains opaque key sets between pairs of team players until the key set graph is connected. In particular, the key set graph will be a spanning tree. Once the key set graph is connected, P_1 chooses a bit v randomly. Using the key sets as one-time pads as described above, P_1 sends v secretly to each of the team players she shares a key set with. In turn, each player learning v sends v to each of the other players she shares a key set with, until all team players learn v. Clearly v satisfies agreement and uniformity. Secrecy is satisfied because each one-time pad preserves secrecy. The idea of using a spanning tree of 2-player secret keys to communicate a single secret key is very useful in multiparty protocols, and will be used again in Chapter 7.

Using some of the techniques described in §5.2, the basic key set protocol repeatedly carries out the following strategy.

1. A team player P is chosen to be the *proposer*.

2. The proposer P chooses a random card x contained in her hand and a random card y not in her hand and proposes $X = \{x, y\}$ as a key set.

3. If another team player Q holds y, she knows that X is a key set, so she *accepts* X. Q is called the *accepter*. An edge labeled $\{x, y\}$ between P and Q is added to the key set graph and the cards x and y are discarded. Whichever player of P and Q holds fewer cards exposes the remaining cards in her hand, which are discarded, and drops out of the protocol. The remaining team players go back to step 1, using the remaining cards in their hands as a "new" deal.

4. If none of the team players holds y, then X is *rejected*. In this case, x and y are discarded, and the players go back to step 1.

By having the proposer or accepter of each obtained key set drop out, this strategy maintains the property that in the key set graph, each connected component is a tree and has exactly one active player. Hence, when there is only one team player left, the key set graph is connected and is a spanning tree.

6.1.1 The Protocol

In this section, we describe how the team players carry out the above strategy. By considering different rules for choosing the proposer in step 1, we obtain the class of *key set protocols*. We call such a rule a *proposing rule*. Formally, a proposing rule is a function $\mathcal{V}$ that takes any signature and returns an index $i \in K \cup \{0\}$. If $i \in K$ is returned, then P_i is the proposer chosen by $\mathcal{V}$. If 0 is returned, no proposer is chosen (and the protocol terminates).

For simplicity, we are not concerned here with conserving independent randomness. We give each team player P_i a sufficient number of rolls of j-sided dice of various sizes j for P_i to choose random cards from sets of various sizes without running out of dice rolls during the course of the protocol. Each die roll is used at most once.

The $\mathcal{V}$ key set ξ-protocol $\mathcal{P}$ operates in two parts. The first part obtains key sets using the strategy outlined above; the second part uses the key sets obtained in the first phase to distribute a secret bit to the team players. The first part is divided into phases, each of which attempts to obtain a key set. A successful key set phase takes three communication rounds, while a failed key set phase takes only two communication rounds. The protocol maintains a current deck and signature that change throughout the course of the protocol by players discarding cards and/or dropping out. In addition, the key set graph is maintained. Initially, the graph consists of a node for each team player and no edges. Each key set phase operates as follows.

1. Each player determines whether or not she is the proposer by applying $\mathcal{V}$ to the current signature. The player P who determines she is the proposer uses her random inputs to choose a card x in her hand and a card y not in her hand. P sends the message "I hold exactly one of the cards $\{x, y\}$." All other players send the pass message $\emptyset$.

2. Each player except P checks whether she holds one of $\{x, y\}$ in her current hand. If any player Q does hold one of $\{x, y\}$, she sends the message "So do I." (Note that there is at most one such player Q.) All other players, including P, send the pass message $\emptyset$.

3. If some team player Q sends the message "So do I," then

   (a) The set $\{x, y\}$ is discarded. (This requires no further communication). The resulting signature has one fewer card for each of P and Q, and two fewer cards in the deck. The edge $\{P, Q\}$, labeled $\{x, y\}$, is added to the graph.

   (b) After $\{x, y\}$ has been discarded, the player of P and Q holding fewer cards sends a message announcing her remaining cards and drops out of the protocol. The resulting signature contains one fewer player and s fewer cards, where s is the number of cards currently held by the player who drops out.

4. If no team player sends the message "So do I," then the set $\{x, y\}$ is discarded. (Again, this requires no further communication). The resulting signature has one fewer card for each of P and Eve, and two fewer cards in the deck.

These steps are repeated until either there are not enough cards left to complete steps 1 and 2, or until only one team player is left. In the second case, all the team players, including those who dropped out earlier, use their key sets to agree on a secret bit as described above, and the run of the protocol is said to *succeed*. Otherwise, it *fails*.

A source $\mathcal{S}$ for $U^{\mathcal{P}}$ is a *key set* source if $\mathcal{S}$ is a card game source and the distribution induced on every die roll is uniform. In this case, we call $\mathcal{P}_{\mathcal{S}}$ a $\mathcal{V}$ key set ξ-system.

6.1.2 Analysis

Fix a signature ξ and a proposing rule $\mathcal{V}$, and let $\mathcal{P}$ be the $\mathcal{V}$ key set ξ-protocol and let $\mathcal{S}$ be a source for $U^{\mathcal{P}}$. Let $\tau \in \mathsf{c}$. We say a key set X shared by P_{i_0} and P_{i_1} is *opaque with respect to* τ if

$$\Pr\left[\mathrm{E}(a_e) \cap \mathrm{C}(\tau) \cap \mathrm{B}_i(x)\right] = \Pr\left[\mathrm{E}(a_e) \cap \mathrm{C}(\tau)\right] \cdot \frac{1}{2}$$

for all $a_e \in U_e$, $i \in \{i_0, i_1\}$, and $x \in X$.

For $\sigma \in \mathsf{c}$, we define keysets(σ) to be the set of key sets that label the edges of the graph that results when the conversation is σ. Let L be the set of legal ξ-deals, and let δ be a ξ-deal. If $\mathcal{P}_{\mathcal{S}}$ succeeds for every possible run with the deal δ, then we say $\mathcal{P}_{\mathcal{S}}$ *succeeds* for δ. If $\mathcal{P}_{\mathcal{S}}$ succeeds for every $\delta \in L$, we say $\mathcal{P}_{\mathcal{S}}$ *succeeds for* ξ, or simply *succeeds*. If $\mathcal{P}_{\mathcal{S}}$ does not succeed, then $\mathcal{P}_{\mathcal{S}}$ is said to *fail*.

Theorem 6.1 *Let $\mathcal{P}_{\mathcal{S}}$ be a $\mathcal{V}$ key set ξ-system. If $\mathcal{P}_{\mathcal{S}}$ succeeds for ξ, then $\mathcal{P}_{\mathcal{S}}$ is a card game ξ-system for perfect secret bit exchange against the extra cards view.*

Proof (sketch): Let $\mathcal{P}_{\mathcal{S}}$ be a $\mathcal{V}$ key set ξ-system and suppose that $\mathcal{P}_{\mathcal{S}}$ succeeds for ξ. Then every run of $\mathcal{P}_{\mathcal{S}}$ succeeds. Agreement is satisfied because for each feasible point, P_1 chooses the output value v, which each team player successively learns. Since P_1 chooses v uniformly from $\{0, 1\}$, perfect uniformity is satisfied.

To see that perfect secrecy against the extra cards view is satisfied, we examine the key sets obtained during the run of the system. First, if the conversation is τ and a key set has just been accepted, then the newly obtained key set $\{x, y\}$ shared by P and Q is opaque with respect to τ since out of all the inputs where τ is the conversation and Eve holds a particular hand, P holds x in exactly half of them. It can be inductively shown that if a key set X is opaque with respect to a feasible conversation τ, then X is opaque with respect to every feasible extension of τ. This is because x and y are discarded from the deck, and none of what follows depends on the cards x and y or which of P and Q held x. Hence, for every feasible $\sigma \in$ cc, $a_e \in U_e$, $i \in K$, $X \in$ keysets(σ), and $x \in X$, X is opaque with respect to σ, so

$$\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\tau) \cap \mathrm{B}_i(x)] = \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\tau)] \cdot \frac{1}{2}$$

Since $\mathrm{C}(\sigma) \cap \mathrm{O}_i(v) = \mathrm{C}(\sigma) \cap \mathrm{B}_i(x)$ for some x and X such that $x \in X \in$ keysets(σ),

$$\begin{aligned} \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)] &= \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{B}_i(x)] \\ &= \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\tau)] \cdot \frac{1}{2} \\ &= \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] \cdot \Pr[\mathrm{O}_i(v)] \end{aligned}$$

for every feasible complete conversation σ, $a_e \in U_e$, $i \in K$, and $v \in V$. Since, in addition, for nonfeasible $\sigma \in$ cc, $\Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma) \cap \mathrm{O}_i(v)] = \Pr[\mathrm{E}(a_e) \cap \mathrm{C}(\sigma)] \cdot \Pr[\mathrm{O}_i(v)] = 0$, it follows that perfect secrecy against the extra cards view is satisfied. ∎

We say a team player P_i is *viable* if $s_i > 1$, or if $s_i = 1$, $e = 0$, and $s_j \neq 1$ for all $j \neq i$. Otherwise P_i is *nonviable*. P_i is called a *singleton* if $s_i = 1$. Define the function size$(\xi) = k + e$. An inductive argument shows the following.

Theorem 6.2 *Let $\mathcal{V}$ be a proposing rule that always chooses a viable player whenever some player is viable. Let $\xi = (s_1, ..., s_k; d)$. Let $\mathcal{P}_S$ be a $\mathcal{V}$ key set ξ-system. If $s_i \geq 1$ for $1 \leq i \leq k$ and $\max s_i + \min s_i \geq$ size(ξ), then $\mathcal{P}_S$ succeeds for ξ.*

Proof: Suppose the conditions of the theorem and suppose that $s_i \geq 1$ for $1 \leq i \leq k$ and $\max s_i + \min s_i \geq$ size(ξ). We show by induction on size(ξ) that the $\mathcal{P}_S$ succeeds for ξ. If $k = 1$, the system succeeds trivially before it starts. Otherwise $k \geq 2$. If size$(\xi) = 2$ then, since $k \geq 2$, we have $k = 2$ and $e = 0$, and thus the signature is $(s_1, s_2; s_1 + s_2)$, for some $s_1, s_2 \geq 1$. A proposed key set is guaranteed to be accepted, leaving one team player.

Let $t \geq 2$. Inductively assume that the theorem holds for all ξ such that size$(\xi) = t$, and suppose that size$(\xi) = t + 1$.

Let $M = \max s_i$, and $m = \min s_i$. By assumption, $M + m \geq$ size$(\xi) > 2$. Then since $M \geq m$, we have $M \geq 2$. Hence there is at least one viable player (namely the one who holds M cards), so the proposer chosen is viable. After step 2 of the protocol, execution goes to either step 3 or step 4.

If step 4 is taken, then the proposed key set cards are discarded and execution returns to step 1 of the protocol with a "new" deal of some signature ξ'. The signature

ξ' has the same number of team players as ξ, and Eve loses a card. Hence size(ξ') = size$(\xi) - 1$. Let M' and m' denote the maximum and the minimum of the team players' hand sizes in ξ', respectively. In going from ξ to ξ', only one team player's hand changes size, and that hand loses only one card. Hence at most one of the maximum or the minimum decreases, and if either decreases, it decreases by at most one. Hence $M' + m' \geq M + m - 1 \geq \text{size}(\xi) - 1 = \text{size}(\xi') = t$. In addition, since the proposer is feasible, the team player that loses a card must have at least two cards, so in the new signature each player still has at least one card. By induction, the protocol finishes for ξ' and we are done.

Suppose instead step 3 is taken and execution returns to step 1 with a new deal of some signature ξ'. If $k = 2$, then the new signature has only one team player and the system succeeds. Otherwise, $k > 2$. In ξ', there are $k - 1$ team players and the number of cards held by Eve is the same as in ξ, so size(ξ') = size$(\xi) - 1$. As before let M' and m' be the size of the largest and smallest hand in ξ'. In this case, in going from ξ to ξ', one team player drops out and one team player discards a card. Furthermore, the team player who discards a card has at least as large a hand as the player who drops out. Hence removing the player (but not yet discarding the card from the other player's hand) does not decrease the maximum, and can only increase the minimum. Finally, discarding one card decreases the maximum plus the minimum by at most one. To see this, note that if all the s_i are equal, then discarding one card decreases the minimum by one and does not change the maximum. Hence $M = m \geq 2$ and so $M' = M \geq 2$, and $m' = M - 1 \geq 1$, so $M' + m' = M + m - 1 \geq \text{size}(\xi) - 1 = \text{size}(\xi')$, and we are done. If the s_i are not all equal, then discarding one card decreases the maximum by at most one, and the minimum does not decrease. Hence $M' \geq M - 1 \geq 1$, and $m' \geq m \geq 1$. Thus again $M' + m' \geq \text{size}(\xi')$, and we are done. ■

Corollary 6.3 *Let* $\xi = (s_1, ..., s_k; d)$*. If* $s_i \geq 1$ *for* $1 \leq i \leq k$ *and* $\max s_i + \min s_i \geq \text{size}(\xi)$*, then* $\text{pcap}(\xi) \geq 2$*.*

Proof: Let $\xi = (s_1, ..., s_k; d)$ and suppose that $s_i \geq 1$ for $1 \leq i \leq k$ and $\max s_i + \min s_i \geq \text{size}(\xi)$. Let $\mathcal{V}$ be any proposing rule that always chooses a viable player whenever some player is viable and let $\mathcal{P_S}$ be a $\mathcal{V}$ key set ξ-system. By Theorem 6.2, $\mathcal{P_S}$ succeeds for ξ. Hence, by Theorem 6.1, $\mathcal{P_S}$ is a card game ξ-system for perfect secret bit exchange against the extra cards view. Since the extra cards view is a refinement of the empty view, it follows from Lemma 3.4 that $\mathcal{P}_{\mathcal{S}}$ is a card game ξ-system for perfect secret bit exchange against the empty view. Hence, $\mathrm{pcap}(\xi) \geq 2$. ■

6.2 An Optimal Key Set Protocol

The converse to Theorem 6.2 does not in general hold. For example, consider the *smallest viable player* proposing rule (SVP): If some team player is viable, the viable player with the smallest hand is chosen. (Ties are broken in favor of the lower-numbered player.) If no team player is viable, the lowest-numbered team player with a nonempty hand is chosen, if there is one. The signature $\xi = (3, 2, 2, 2, 1; 10)$ does not satisfy the conditions of the Theorem 6.2 since $3+1 < 5+0$, but any SVP key set ξ-system succeeds for ξ. In fact, an SVP key set ξ-system succeeds for any signature ξ in which Eve has no cards, each team player has at least one card, and there is no nonviable singleton.

Theorem 6.4 *Let $\xi = (s_1, \ldots, s_k; d)$ and let $e = d - \sum s_i$ and let $\mathcal{P}_{\mathcal{S}}$ be a SVP key set ξ-system. If $e = 0$, $s_i \geq 1$ for $1 \leq i \leq k$, and there is at most one $i_0 \in \{1, \ldots, k\}$ such that $s_{i_0} = 1$, then $\mathcal{P}_{\mathcal{S}}$ succeeds for ξ.*

Proof: Suppose the conditions of the theorem. We show by induction on k that $\mathcal{P}_{\mathcal{S}}$ succeeds for ξ. If $k = 1$, the system succeeds trivially before it starts. Otherwise $k \geq 2$. Inductively assume that the theorem holds for all signatures with less than k players and suppose that ξ has $k \geq 2$ players.

Since $e = 0$, $s_i \geq 1$ for $1 \leq i \leq k$, and there is at most one $i_0 \in \{1, \ldots, k\}$ such that $s_{i_0} = 1$, every player is viable. Thus, in step 1 of the protocol, the player P_i with the smallest hand is chosen. Since $e = 0$, execution of the protocol, goes to step 3 after step 2. Since P_i holds the smallest hand, P_i drops out of the protocol. The player P_j who holds the second key set card y discards y. The remaining players return to step 1 with the resulting deal of some signature ξ'.

In going from ξ to ξ', player P_i has dropped out. Of the remaining players, only player P_j's hand changed size, from s_j to $s_j - 1$. Therefore, by induction, it suffices to show that $s_j - 1 \geq 1$ and $s_\ell > 1$ for all $\ell \neq i, j$. Since P_i was chosen by SVP, it follows that $s_\ell \geq s_i$ for all $\ell \neq i$. Furthermore, since $s_\ell \geq 1$ for all ℓ and there is at most one i_0 such that $s_{i_0} = 1$, it follows that $s_\ell > 1$ for all $\ell \neq i$. Thus $s_j - 1 \geq 1$ and $s_\ell > 1$ for all $\ell \neq i, j$. By induction, $\mathcal{P}_{\mathcal{S}}$ succeeds for ξ'. It follows that $\mathcal{P}_{\mathcal{S}}$ succeeds for ξ. ■

There are other signatures, such as $(3, 3, 2, 1; 10)$, for which the SVP key set system succeeds that satisfy the conditions neither for Theorem 6.2 nor the conditions for Theorem 6.4. We have been unable to find an exact characterization of the signatures for which some key set system succeeds. Nevertheless, it is possible to show that the SVP key set protocol is optimal for the class of key set protocols. That is, we show for any signature ξ and proposing rule $\mathcal{V}$ that if the $\mathcal{V}$ key set system succeeds for ξ, then the SVP key set system also succeeds for ξ. To prove this, we look at a simple combinatorial stick game between a team and an adversary. The stick game abstracts the important aspects of the key set protocol.

6.2.1 The Stick Game

The stick game is a game between two players, the *chooser* and the *adversary.* It is played on $k+1$ piles of sticks: k chooser piles P_1 through P_k and a single adversary pile E. Pile P_i contains $|P_i|$ sticks, and pile E contains $|E|$ sticks. On the chooser's turn, the chooser chooses one of her piles, say P_i, that contains at least one stick. On the adversary's turn, the adversary chooses any nonempty pile other than the pile just chosen by the chooser. At this point, the following occurs.

1. If the pile chosen by the adversary is another chooser pile P_j, the smaller of P_i and P_j is removed entirely, and one stick is removed from the larger pile.
2. If the pile chosen by the adversary is the adversary pile E, one stick is removed from each of P_i and E.

Note that removing a pile is not the same as removing all the sticks in the pile; the latter leaves an empty pile. The chooser always moves first.

A configuration of the stick game is described by a tuple $(p_1, \ldots, p_k; e; I)$, where $p_i = |P_i|$ and $e = |E|$. I is called the *index component.* On the chooser's turn, I has the value 0. On the adversary's turn, I is the index of the pile designated by the chooser on her previous turn. We generally talk about a chooser move as a *choice* of P_i, and an adversary move as a *match*, either with some P_j or with E.

Play ends when the player whose turn it is cannot move. If there are one or zero chooser piles left, the chooser *wins.* Otherwise the chooser *loses.* We call the stick game starting from configuration C the C *stick game.*

A *strategy* for the chooser, or chooser strategy, is a function that, given a configuration with index component 0 in which a chooser move is possible, specifies the next chooser move. Similarly, an adversary strategy is a function that, given a configuration with index component $I \neq 0$ in which an adversary move is possible, specifies the next adversary move.

We say a chooser strategy $\mathcal{W}$ *succeeds on* configuration C if the chooser wins when playing the C stick game according to $\mathcal{W}$, regardless of the moves chosen by the adversary. We call C *winning* if some chooser strategy succeeds on C. A chooser strategy $\mathcal{W}$ is *optimal* if it succeeds on every winning configuration. Similarly, we say an adversary strategy $\mathcal{A}$ *succeeds on* configuration C if the chooser loses when the adversary plays the C stick game according to $\mathcal{A}$, regardless of the moves chosen by the chooser. We call C *losing* if some adversary strategy succeeds on C. An adversary strategy $\mathcal{A}$ is *optimal* if it succeeds on every losing configuration.

Note that if C is a configuration in which it is the chooser's turn, then C is winning if and only if there exists a chooser move that results in a winning configuration. If C is a configuration in which it is the adversary's turn, then C is winning if and only if all adversary moves result in a winning configuration.

The stick game is a finite game, since every adversary turn decreases the total number of sticks by at least two. Furthermore, it is a game of complete information, since the chooser and the adversary take turns and all information about the state is known to both the chooser and the adversary. Hence a well known game theoretic result [BCG82] tells us that every configuration is either winning or losing, and there exist both an optimal chooser strategy $\mathcal{W}$ and an optimal adversary strategy $\mathcal{A}$.

The stick game is not as simple as it might at first seem. Clearly, the chooser's goal is to get rid of all but one of her piles, after which a win is guaranteed. The adversary's goal is to cause the game to terminate while the chooser still has two or more piles remaining. Let k be the number of chooser piles and e be the size of the adversary's pile. The adversary can make at most e matches with E, after which E is empty. In order to defeat the chooser, the adversary can make at most $k-2$ matches with chooser piles, for each such match reduces by one the number of chooser piles, and after $k-1$ such matches, the chooser is guaranteed to win.

The chooser is faced with the dilemma of whether to choose a large pile, possibly giving the adversary the opportunity to choose and eliminate another large pile, or to choose a small pile, which may give the adversary the opportunity to produce a still smaller pile. (Note that if any chooser pile becomes empty before the game ends, then the chooser will eventually lose, so it is to the adversary's advantage to make chooser piles smaller.)

The adversary in turn must decide whether to match with a chooser pile or with his own pile, and if he matches with a chooser pile, whether to choose a larger or smaller pile than the one chosen by the chooser. The advantage of choosing a larger pile is that more sticks get removed from the game when the chooser's designated pile is eliminated. The advantage of choosing a smaller pile is that a stick gets removed from the chooser's designated pile, reducing its size.

6.2.2 Chooser Strategies

Given a configuration $C = (p_1, \ldots, p_k; e; I)$, we define the functions $\text{size}(C) = k + e$ and $\text{piles}(C) = k$. The following proposition is an immediate consequence of the fact that the stick game is a finite game and the number of chooser piles never increases during play.

Proposition 6.5 *If* $\text{piles}(C) \leq 1$*, then* C *is winning, and every chooser strategy succeeds on* C*.*

The following proposition is easily proved by an exhaustive case analysis of possible moves.

Proposition 6.6 *If* $C = (1, 1; 0; 0)$*, then* C *is winning, and every chooser strategy succeeds on* C*.*

We say a chooser pile P_i is *viable* if either

- $|P_i| \geq 2$, or
- $|P_i| = 1, |E| = 0$ and for all $j \neq i$, $|P_j| \geq 2$.

Note that if all chooser piles are nonempty and the chooser chooses a viable pile, then, regardless of the match made by the adversary, all chooser piles in the resulting configuration are nonempty. This is obvious when a pile of size 2 or greater is chosen. When a viable singleton is chosen, all other chooser piles have size at least 2, and the adversary has no choice but to match with a chooser pile.

6.2.3 Optimality of the SVP Stick Game Strategy

In order to show the optimality of the SVP strategy, it suffices to show the following: for any winning configuration C in which it is the chooser's turn to move next, if the chooser moves according to SVP, then the resulting configuration C' is also winning. We use what is known as a strategy stealing argument to show that C winning implies C' is winning.

The basic idea of the argument is to simultaneously play two stick games, one on C and one on C'. We call the two games G and G', respectively. In game G, the chooser plays using an optimal chooser strategy and we specify the adversary moves. Since the chooser plays optimally, if C is winning, then all resulting configurations in G will be winning, regardless of the moves we specify for the adversary. In G', the adversary plays using an optimal adversary strategy and we specify the chooser moves. Since the adversary plays optimally, if C' is losing, then all resulting configurations in G' are losing, regardless of the moves we specify for the chooser.

In more detail, let $C_0 = C$ and $C'_0 = C'$. We call the C stick game G, and we refer to the G chooser and the G adversary. Similarly, we have the G' game, chooser and adversary. We play G in order to develop a sequence of configurations $C_0, C_1, C_2, \ldots$ that describe a play of the C stick game, and we play G' in order to develop a sequence of configurations $C'_0, C'_1, C'_2, \ldots$ that describe a play of the C' stick game. If C is winning, then all C_i are winning since the chooser plays optimally in G. If C'_0 is losing, then all C'_i are losing since the adversary plays optimally in G'. Hence, if any C'_j can be shown to be winning from the assumption that some C_i is winning, it will follow that $C = C_0$ winning implies $C' = C'_0$ winning. This is illustrated in Figure 6.1.

$$
\begin{array}{ccc}
C_0 \text{ winning} & & C'_0 \text{ winning} \\
\Downarrow & & \Uparrow \\
\vdots & & \vdots \\
\Downarrow & & \Uparrow \\
C_i \text{ winning} & \Longrightarrow & C'_j \text{ winning}
\end{array}
$$

Figure 6.1: The strategy stealing argument

In our strategy stealing argument, adversary moves in G and chooser moves in G' will generally be based on *mirroring*. G and G' will be played simultaneously, with moves alternately chosen from the two games. The G adversary selects its next move by mirroring the G' adversary, matching with the same pile that the G' adversary previously matched with, if possible. Similarly, the G' chooser selects its next move by mirroring the G chooser, choosing the same pile that the G chooser previously selected, if possible. Where mirroring cannot be applied (because, for example, the corresponding pile does not exist), we will be able to either choose an alternate move or prove on the basis of the play constructed so far that C winning implies C' winning. In this latter case, the construction can stop.

Since we do not know the optimal strategy of either chooser or adversary, we must consider all possible moves in our analysis and show for each that the argument works, under the assumption that the given move really was optimal. A large number of cases can result during this process. In order to make the case analysis manageable, we have broken the proof of the optimality of SVP into several lemmas. Lemmas 6.7, 6.9, and 6.10 formalize some useful and intuitively easy facts. Lemmas 6.8, 6.11 and 6.12 start with two configurations C and C' that have a known relationship and show that if C is winning then C' is also winning. These are each proved using a strategy stealing argument. The proof of each Lemma is inductive, so often it is only necessary to play G and G' for one or two steps in order to obtain a pair (C_i, C'_j) to which the inductive hypothesis applies. Lemma 6.13 provides an exact characterization of winning configurations with only two piles. Finally, Theorem 6.14 shows the optimality of SVP.

We start with the intuitively obvious fact that it is impossible for the chooser to win the stick game when there is a chooser pile containing no sticks unless it is the only chooser pile.

Lemma 6.7 *Let $C = (p_1, \ldots, p_k; e; 0)$. If $k \geq 2$ and $p_i = 0$ for some i, then C is losing.*

Proof: We proceed by induction. Let $C = (p_1, \ldots, p_k; e; 0)$, let $k \geq 2$, and suppose, without loss of generality, that $p_1 = 0$. Inductively assume that the lemma holds for all $\overline{C}$ such that $\text{size}(\overline{C}) < \text{size}(C)$. We show that C is losing.

If there is no nonempty chooser pile in C, then the chooser cannot move and the game ends. The chooser loses since $k \geq 2$. Otherwise, at least one chooser pile is nonempty and the chooser has a move. We show that no matter which move the chooser makes, either the game ends and the chooser loses, or there is an adversary response that results in a losing configuration. Hence, C is also losing.

Let P_i be the nonempty pile selected by the chooser, so the resulting configuration is $C_1 = (0, p_2, \ldots, p_k; e; i)$ for some $i \neq 1$. If there is no adversary move, then the game ends and the chooser loses. Otherwise, the adversary must choose a nonempty pile to match with. Let C_2 be the resulting configuration. If the adversary matches with E, then C_2 has at least two piles, P_1 and P_i, and $\text{size}(C_2) = k + (e-1) = \text{size}(C) - 1$. If the adversary matches with P_j, then $j \notin \{1, i\}$, so $k \geq 3$. Hence, C_2 has at least two piles, and $\text{size}(C_2) = (k-1) + e = \text{size}(C) - 1$. In either case, $\text{piles}(C_2) \geq 2$, C_2 has a pile of size 0, and $\text{size}(C_2) = \text{size}(C) - 1$. Therefore C_2 is losing by induction. ∎

We extend the notion of domination to configurations. Let $C = (p_1, \ldots, p_k; e; 0)$ and $C' = (p'_1, \ldots, p'_{k'}; e'; 0)$ be two configurations. We say C' *dominates* C if $k = k'$, $e' \leq e$ and there is a permutation π on $\{1, \ldots, k\}$ such that $p'_i \geq p_{\pi(i)}$ for each i. We write $C' \trianglerighteq C$.

Lemma 6.8 (Stick Game Domination Lemma) *Let $C' \trianglerighteq C$. If C is winning, then C' is winning.*

Proof: Let C and C' be two configurations such that $C' \trianglerighteq C$, and suppose that C is winning. Inductively assume that the lemma holds for all $\overline{C}$ and $\overline{C}'$ such that $\text{size}(\overline{C}) < \text{size}(C)$. We use a strategy stealing argument to show that C' is winning by playing the C stick game G and the C' stick game G', as shown in Figure 6.1. Let $C_0 = C = (p_1, \ldots, p_k; e; 0)$ and let $C'_0 = C' = (p'_1, \ldots, p'_{k'}; e'; 0)$. We show that C'_0 is winning.

If $k \leq 1$, then by Proposition 6.5, C'_0 is winning. Otherwise, $\text{size}(C_0) \geq \text{size}(C'_0) \geq k \geq 2$.

The G chooser chooses a pile, say P_i. Let $C_1 = (p_1, \ldots, p_k; e; i)$ be the resulting configuration. In game G', we consider the corresponding chooser choice $P_{\pi^{-1}(i)}$, which must be possible because $p'_{\pi^{-1}(i)} \geq p_i$, by the definition of domination. Hence, the resulting configuration is $C'_1 = (p'_1, \ldots, p'_k; e'; \pi^{-1}(i))$. The G' adversary now makes a move, resulting in configuration C'_2.

We construct configuration C_2 by having the G adversary mirror the G' adversary, if possible. If the G' adversary matches with E, then the G adversary can match with E, since $e \geq e'$, by the definition of domination. Then $C'_2 \trianglerighteq C_2$, and so by induction, C'_2 is winning, and hence C'_0 is winning.

Otherwise the G' adversary matches with some P_j. Because C_0 is winning and $k \geq 2$, Lemma 6.7 implies that $p_\ell \geq 1$ for every ℓ. Hence the G adversary can match with $P_{\pi(j)}$. In this case, the resulting C_2 is the same as C_0 except that in place of P_i and $P_{\pi(j)}$ there is a single pile of size $\max(p_i, p_{\pi(j)}) - 1$. Similarly, C_2' differs from C_0' only in that piles $P_{\pi^{-1}(i)}$ and P_j are replaced by a single pile of size $\max(p'_{\pi^{-1}(i)}, p'_j) - 1$. Since $C_0' \trianglerighteq C_0$ (via π), we have that $p'_{\pi^{-1}(i)} \geq p_i$ and $p'_j \geq p_{\pi(j)}$. Hence $\max(p'_{\pi^{-1}(i)}, p'_j) \geq \max(p_i, p_{\pi(j)})$. Therefore, modifying π so that the two new piles correspond, we see that $C_2' \trianglerighteq C_2$. Also, $\mathrm{size}(C_2) = \mathrm{size}(C_0) - 1$. Hence by induction, C_2' is winning, so C_0' is winning. ■

A special case of the domination lemma is that if two configurations differ only by a permutation of the pile sizes, they are either both winning or both losing. We sometimes say two such configurations are *equal up to permutation.*

We call a chooser pile of size 1 a *singleton*. Lemma 6.9 says that whenever a singleton is viable in a given configuration, the configuration is winning.

Lemma 6.9 *Let $C = (p_1, \ldots, p_k; e; 0)$. If for some i, P_i is viable and $p_i = 1$, then C is winning.*

Proof: Let $C = (p_1, \ldots, p_k; e; 0)$, and without loss of generality, assume that P_1 is viable and $p_1 = 1$. Inductively assume that the lemma holds for all configurations $\overline{C}$ such that $\mathrm{size}(\overline{C}) < \mathrm{size}(C)$, and let $C_0 = C$.

If $k \leq 1$ then C is winning by Proposition 6.5. Otherwise $k \geq 2$, and since P_1 is viable and $p_1 = 1$, then $e = 0$ and for every $\ell \geq 2$, $p_\ell \geq 2$. Let the chooser choose pile P_1, so $C_1 = (1, p_2, \ldots, p_k; 0; 1)$ is the resulting configuration. To show that C_0 is winning, it suffices to show that C_1 is winning.

Since $p_\ell \geq 2$ for every $\ell \geq 2$, the adversary does have some move. Because $e = 0$, the adversary must match with some P_j. Without loss of generality, let $j = 2$, so the resulting configuration is $C_2 = (p_2 - 1, p_3, \ldots, p_k; 0; 0)$. Let $\overline{C} = (1, p_3, \ldots, p_k; 0; 0)$. Since $\mathrm{size}(\overline{C}) = \mathrm{size}(C_0) - 1$, then $\overline{C}$ is winning by induction. Since $C_2 \trianglerighteq \overline{C}$, then C_2 is winning by Lemma 6.8. Hence, C_1, and therefore C_0, are both winning. ■

The proof of Lemma 6.9 shows that if a viable singleton is chosen in a given configuration, the resulting configuration is winning. Lemma 6.10 provides a partial converse to this result. Namely, if a nonviable singleton is chosen in a configuration with at least three piles, then the resulting configuration is losing.

Lemma 6.10 *Let $k \geq 3$ and $p_i = 1$ for $i \in K$, and let $C = (p_1, \ldots, p_k; e; 0)$. Let $C_1 = (p_1, \ldots, p_k; e; i)$. If P_i is not viable in C, then C_1 is losing.*

Proof: Without loss of generality, let $C = (1, p_2, \ldots, p_k; e; 0)$, where $k \geq 3$, and let $C_1 = (1, p_2, \ldots, p_k; e; 1)$. Assume P_1 is not viable in C, and let $C_0 = C$. We show that C_1 is losing.

By Lemma 6.7, if any $p_j = 0$, then C_0, and hence C_1, are both losing. Otherwise, assume all $p_j \geq 1$. Since P_1 is not viable, then either $e \geq 1$ or $p_j = 1$ for some $j \neq 1$. If $e \geq 1$, then the adversary can match with E yielding a configuration $C_2 = (0, p_2, \ldots, p_k; e-1; 0)$, which is losing by Lemma 6.7. Similarly, if $p_j = 1$ for some $j \neq 1$, then the adversary can match with P_j. Without loss of generality, $j = 2$ and the resulting configuration is $C_2 = (0, p_3, \ldots, p_k; e; 0)$. Since C_2 has $k - 1 \geq 2$ piles, it is losing by Lemma 6.7. ■

Lemma 6.11 *Let $C = (p_1, \ldots, p_{k-1}, x; e; 0)$ and $C' = (p_1, \ldots, p_{k-1}, x-1; e-1; 0)$, where $e \geq 1$, $x \geq 2$, and each p_i satisfies $p_i \geq x$ or $p_i = 1$. Then if C is winning, C' is winning.*

Proof: Let $C = (p_1, \ldots, p_{k-1}, x; e; 0)$ and $C' = (p_1, \ldots, p_{k-1}, x-1; e-1; 0)$, and suppose that C is winning, $e \geq 1$, $x \geq 2$, and each p_i satisfies $p_i \geq x$ or $p_i = 1$. Inductively assume that the lemma holds for all $\overline{C}$ and $\overline{C}'$ such that size($\overline{C}$) $<$ size(C). We use a strategy stealing argument to show that C' is winning by playing the C stick game G and the C' stick game G', as shown in Figure 6.1. Let $C_0 = C$ and $C'_0 = C'$. We show that C'_0 is winning.

If $k \leq 1$, then C'_0 is winning by Proposition 6.5. Otherwise $k \geq 2$, so the G chooser has not already won in C_0. We assume the G chooser is using an optimal strategy which selects, say, pile P_i. If $p_i = 1$, then the G adversary can match with E. In this case the resulting configuration C_2 still has $k \geq 2$ piles and so is losing by Lemma 6.7, a contradiction. Hence $p_i > 1$, and therefore $p_i \geq x$.

If $p_i = x$, then we have the G adversary match with E. Then the resulting configuration C_2 is equal up to permutation to $(p_1, \ldots, p_{k-1}, x-1; e-1; 0) = C'_0$, so C'_0 is winning. Otherwise $p_i > x$, and without loss of generality, $i = 1$. Then C_1 is $(p_1, \ldots, p_{k-1}, x; e; 1)$. We mirror the G chooser move to G', so $C'_1 = (p_1, \ldots, p_{k-1}, x - 1; e-1; 1)$.

Given C'_1, the G' adversary has some possible move, because $x \geq 2$ and therefore the pile containing $x - 1$ sticks is nonempty. We mirror the G' adversary move to G to obtain C_2. We consider three types of adversary moves separately. In each case we show that C'_2 is winning.

- G' adversary matches with E. Then $e-1 \geq 1$, and

$$\begin{aligned} C_2' &= (p_1-1, p_2, \ldots, p_{k-1}, x-1; e-2; 0) \\ C_2 &= (p_1-1, p_2, \ldots, p_{k-1}, x; e-1; 0) \end{aligned}$$

  Furthermore, $p_1 > x$ implies $p_1 - 1 \geq x$, so by induction, C_2' is winning.

- G' adversary matches with P_j such that $p_j \neq x-1$. Without loss of generality $j=2$, so

$$\begin{aligned} C_2' &= (\max(p_1,p_2)-1, p_3, \ldots, p_{k-1}, x-1; e-1; 0) \\ C_2 &= (\max(p_1,p_2)-1, p_3, \ldots, p_{k-1}, x; e; 0) \end{aligned}$$

  Since $p_1 > x$, we have that $\max(p_1,p_2) - 1 \geq p_1 - 1 \geq x$. By induction, C_2' is winning.

- G' adversary matches with P_j such that $|P_j| = x-1$ in C_1'. Without loss of generality, $j = k$. Since $p_1 > x$, we have

$$\begin{aligned} C_2' &= (p_1-1, p_2, \ldots, p_{k-1}; e-1; 0) \\ C_2 &= (p_1-1, p_2, \ldots, p_{k-1}; e; 0) \end{aligned}$$

  Then $C_2' \trianglerighteq C_2$, so C_2' is winning.

In each case, C_2' is winning. Hence, C_0' is winning, as desired. ■

Lemma 6.12 *Let $C = (p_1, \ldots, p_{k-2}, x, y; e; 0)$ and $C' = (p_1, \ldots, p_{k-2}, y'; e; 0)$, where $k \geq 2$, $y'+1 \geq y \geq x$, $y' \geq 1$, each p_i satisfies $p_i \geq x$ or $p_i = 1$, and either $y' > 1$ or there is no viable singleton in C. Then if C is winning, C' is winning.*

Proof: Let $C = (p_1, \ldots, p_{k-2}, x, y; e; 0)$ and $C' = (p_1, \ldots, p_{k-2}, y'; e; 0)$, where $k \geq 2$, $y'+1 \geq y \geq x$, $y' \geq 1$, each p_i satisfies $p_i \geq x$ or $p_i = 1$, and either $y' > 1$ or there is no viable singleton in C. Suppose further that C is winning. Inductively assume that the lemma holds for all $\overline{C}$ and $\overline{C}'$ such that size$(\overline{C}) <$ size(C). We use a strategy stealing argument to show that C' is winning by playing the C stick game G and the C' stick game G', as shown in Figure 6.1. Let $C_0 = C$ and $C_0' = C'$. We show that C_0' is winning.

If $k = 2$, then piles$(C_0') = 1$, so C_0' is winning by Proposition 6.5. Otherwise $k \geq 3$, and Lemma 6.7 implies each pile of C_0 must be of size at least 1. In particular, $y \geq x \geq 1$. Hence the G chooser has a move, and must choose a pile P_i.

1. G chooser chooses P_i such that $|P_i| = x$ (or y). In this case the G adversary matches with a pile P_j with $p_j = y$ (or x). Without loss of generality, $i = k-1$ and $j = k$. Then $C_2 = (p_1, \ldots, p_{k-2}, y-1; e; 0)$. Hence $C_0' \trianglerighteq C_2$, so C_0' is winning.

2. G chooser chooses P_i such that $|P_i| > x \geq 1$ and $|P_i| \neq y$. Without loss of generality, let $i = 1$. We mirror the G chooser move to G' and consider possible G' adversary moves. Note that there is some possible G' adversary move, since the G' adversary can match with pile P_{k-1}, which has size $y' \geq 1$.

   - G' adversary matches with E. Then $e \geq 1$. If $y' > 1$ or $e > 1$, we mirror the match with E back to G. Then

     $$\begin{aligned} C_2' &= (p_1 - 1, p_2, \ldots, p_{k-2}, y'; e-1; 0) \\ C_2 &= (p_1 - 1, p_2, \ldots, p_{k-2}, x, y; e-1; 0) \end{aligned}$$

     If $y' > 1$, then C_2' is winning by induction. If $e > 1$, then C_2 has no viable singletons, so C_2' is winning by induction.

     Otherwise $y' = 1$ and $e = 1$. In this case, we have the G adversary match with P_{k-1} (of size x). Then

     $$\begin{aligned} C_2' &= (p_1 - 1, p_2, \ldots, p_{k-2}, 1; 0; 0) \\ C_2 &= (p_1 - 1, p_2, \ldots, p_{k-2}, y; 1; 0) \end{aligned}$$

     We have $2 = y' + 1 \geq y \geq 1$. If $y = 1$, then $C_2' \trianglerighteq C_2$, and hence C_2' is winning by Lemma 6.8. If $y = 2$, then since $p_1 - 1 \geq x$, the configurations C_2 and C_2' satisfy Lemma 6.11 and hence C_2' is winning.

   - G' adversary matches with P_j, where $|P_j| \neq y'$. In this case, we mirror the G' adversary move to G. Without loss of generality $j = 2$, so

     $$\begin{aligned} C_2' &= (\max(p_1, p_2) - 1, p_3, \ldots, p_{k-2}, y'; e; 0) \\ C_2 &= (\max(p_1, p_2) - 1, p_3, \ldots, p_{k-2}, x, y; e; 0) \end{aligned}$$

     Hence $p_1 > x$ implies $\max(p_1, p_2) - 1 \geq x$. If $y' > 1$ or there is no viable singleton in C_2, then by induction C_2' is winning. Otherwise $y' = 1$ and there is a viable singleton in C_2. We show that this is impossible. We have $e = 0$, $p_1 > x \geq 1$ and $p_2 \neq y' = 1$. Hence, neither of the piles that changed from C_0 to C_2 were singletons in C_0, so C_0 has at most one singleton. By assumption and the fact that $y' = 1$, there is no viable singleton in C_0, which, together with $e = 0$, implies that there is no singleton at all in C_0. In particular, we have $|P_{k-1}| = x \geq 2$, so since $2 = y' + 1 \geq x$, we conclude that $x = 2$. Furthermore, $p_1 > x$, so $\max(p_1, p_2) - 1 \geq 2$. Hence there are no singletons in C_2, so in particular there is no viable singleton in C_2.

- G' adversary matches with P_j such that $|P_j| = y'$. Without loss of generality, $j = k - 1$. Then the G adversary can match with P_k, so up to permutation,

$$\begin{aligned} C_2' &= (p_2, \ldots, p_{k-2}, \max(p_1, y') - 1; e; 0) \\ C_2 &= (p_2, \ldots, p_{k-2}, x, \max(p_1, y) - 1; e; 0) \end{aligned}$$

We have $p_1 > x$, so $\max(p_1, y') \geq \max(p_1, y) - 1 \geq x$. Also, $\max(p_1, y') - 1 \geq p_1 - 1 \geq x \geq 1$. If $\max(p_1, y') - 1 > 1$, or if there is no viable singleton in C_2, then by induction, C_2' is winning. Otherwise $\max(p_1, y') - 1 = 1$ and there is some viable singleton in C_2. In this case, $p_1 \leq 2$, and thus $x = 1$, hence P_{k-1} is the unique viable singleton in C_2. Then $e = 0$, $p_1 = 2$, and $p_i \geq 2$ for $2 \leq i \leq k - 2$. Then the last pile of C_2' is a viable singleton, so by Lemma 6.9, C_2' is winning.

Since C_2' is winning in each case, C_0' is winning.

3. G chooser chooses P_i such that $|P_i| = 1$ for some $i \leq k - 2$. Without loss of generality, let $i = 1$. Since $k \geq 3$, P_1 is viable in C_0 by Lemma 6.10 and the fact that C_0 is winning. Therefore, $y' > 1$ and $e = 0$, so P_1 is viable in C_0'. By Lemma 6.9, C_0' is winning. ■

Lemma 6.13 *If $C = (p_1, p_2; e; 0)$, then $p_1 + p_2 \geq 2 + e$ if and only if C is winning.*

Proof: Let $C = (p_1, p_2; e; 0)$. It is easy to see that if $p_1 + p_2 \geq 2 + e$, then C is winning, since as long as the chooser never chooses a nonviable singleton, there can be at most e matches with E, after which the adversary has no choice but to match with the other chooser pile.

For the converse, we show by induction on size(C) that if $p_1 + p_2 < 2 + e$ then C is losing. Assume $p_1 + p_2 < 2 + e$, and assume the induction hypothesis holds for all configurations $\overline{C}$ for which size($\overline{C}$) < size(C). Let $C_0 = C$.

If $p_1 = p_2 = 0$, then the chooser cannot move and therefore loses, so C is losing. Otherwise at least one chooser pile is nonempty, and so the chooser must choose a pile P_i. Without loss of generality, let $i = 1$, so the resulting configuration is $C_1 = (p_1, p_2; e; 1)$. If $e = 0$, then $p_1 + p_2 < 2$, so $p_2 = 0$. Hence the adversary cannot move and again C is losing. Otherwise $e > 0$, so the adversary can match with E resulting in configuration $C_2 = (p_1 - 1, p_2; e - 1; 0)$. Then $\text{size}(C_2) < \text{size}(C)$ and $p_1 - 1 + p_2 < 2 + e - 1$, so by induction, C_2 is losing. Hence C is losing. ∎

Thus SVP is optimal when there are only two chooser piles since SVP succeeds on $(p_1, p_2; e; 0)$ if $p_1 + p_2 \geq 2 + e$. We can now show the optimality of SVP in general.

Theorem 6.14 *The SVP chooser strategy is an optimal chooser strategy.*

Proof: By Proposition 6.5 and Lemma 6.13, it remains only to show that SVP is optimal for configurations with at least 3 chooser piles.

Suppose $k \geq 3$ and $C_0 = (p_1, \ldots, p_k; e; 0)$ is winning. Suppose SVP chooses P_i for some i. We may assume without loss of generality that $i = 1$, so the resulting configuration is $C_1' = (p_1, \ldots, p_k; e; 1)$. We use a strategy stealing argument to show that C_1' is winning by playing the C_0 stick game G and the C_1' stick game G'. This shows the optimality of the SVP chooser strategy.

We assume the G chooser is using an optimal strategy which selects, say, pile P_j. Let $C_1 = (p_1, \ldots, p_k; e; j)$ be the resulting configuration. If $p_j = p_1$, then up to permutation $C_1' = C_1$, so C_1' is winning, and we are done. Hence we need only consider $p_j \neq p_1$, in which case $j \neq 1$. We may assume without loss of generality that $j = 2$. Also, since $k \geq 3$, Lemma 6.7 implies $p_\ell \geq 1$ for every ℓ.

Suppose $p_1 = 1$. Then $p_2 \geq 2$, so there is a viable pile in C_0, and P_1 is viable in C_0. Then by Lemma 6.9, C_0' is winning and we are done. Hence, we may assume $p_1 \geq 2$.

Suppose $p_2 = 1$. Then P_2 is not viable, or it would have been chosen by SVP. But this contradicts Lemma 6.10, since C_1 is winning. Hence, $p_2 \geq 2$. Furthermore, because p_1 is chosen according to the *smallest* viable pile rule, it must be the case that $p_2 \geq p_\ell$ or $p_\ell = 1$ for all $\ell \geq 2$. Since also $p_2 \neq p_1$, we have $p_2 > p_1 \geq 2$.

As before, we distinguish the C_1 stick game and the C_1' stick game by calling them the G game and the G' game. We consider possible G' adversary moves from C_1' and show that for every possible G' adversary move from C_1', the resulting configuration C_2' is winning, and hence C_1' is winning.

1. G' adversary matches with E. Then

$$\begin{aligned} C_2' &= (p_1-1, p_2, \ldots, p_k; e-1; 0) \\ C_0 &= (p_1, p_2, \ldots, p_k; e; 0) \end{aligned}$$

Hence C_0 and C_2' satisfy the conditions of Lemma 6.11, so C_2' is winning.

2. G' adversary matches with P_2. Then

$$C_2' = (p_2-1, p_3, \ldots, p_k; e; 0)$$

We have the G adversary match with P_1. Since P_2 matches with P_1 in both games, $C_2 = C_2'$, so C_2' is winning.

3. G' adversary matches with some P_a, where $a \geq 3$. We may assume without loss of generality that $a = 3$, so the resulting configuration is $C_2' = (\max(p_1, p_3) - 1, p_2, p_4, \ldots, p_k; e; 0)$, up to permutation. (For the remainder of this proof, we omit the phrase "up to permutation.") We consider the possible relationships between p_1 and p_3 separately.

   - If $p_1 < p_3$, then we mirror the G' adversary move to G. Then

     $$\begin{aligned} C_2' &= (p_3-1, p_2, p_4, \ldots, p_k; e; 0) \\ C_2 &= (p_1, \max(p_2, p_3) - 1, p_4, \ldots, p_k; e; 0) \end{aligned}$$

     If $p_2 \geq p_3$, then obviously $C_2' \trianglerighteq C_2$. If $p_2 < p_3$, then by flipping the first two piles of C_2, it becomes apparent that $C_2' \trianglerighteq C_2$. In either case, C_2' is winning by Lemma 6.8.

   - If $p_3 < p_1$, we have the G adversary match with P_1. Then

     $$\begin{aligned} C_2' &= (p_1-1, p_2, p_4, \ldots, p_k; e; 0) \\ C_2 &= (p_2-1, p_3, p_4, \ldots, p_k; e; 0) \end{aligned}$$

     Since $p_1 - 1 \geq p_3$ and $p_2 \geq p_2 - 1$, we have $C_2' \trianglerighteq C_2$, and C_2' is winning by Lemma 6.8.

   - If $p_1 = p_3$, then

     $$\begin{aligned} C_2' &= (p_1-1, p_2, p_4, \ldots, p_k; e; 0) \\ C_0 &= (p_1, p_2, p_1, p_4, \ldots, p_k; e; 0) \end{aligned}$$

     We have $p_1 \geq 2$, and because P_1 was chosen by SVP, C_0 has no viable singleton. Hence Lemma 6.12 applies to C_0 and C_2', so C_2' is winning.

Since C_2' is winning in each case, C_1' is winning. ■

6.2.4 Optimality of the SVP Key Set Protocol

There is a close connection between the stick game and key set systems. Namely, a configuration $(s_1, \ldots, s_k; e; 0)$ in which it is the chooser's turn corresponds to a signature $(s_1, \ldots, s_k; e + \sum s_i)$. A configuration in which it is the adversary's turn corresponds to a signature plus a choice of proposer. A pile P_i corresponds to team player P_i, and the pile E corresponds to Eve. A chooser's strategy $\mathcal{W}$ in the stick game corresponds to a proposing rule $\mathcal{V}$ for the key set protocol. That is, the pile P_i is chosen in the stick game if and only if player P_i proposes in the $\mathcal{V}$ key set protocol. The adversary's response in the stick game corresponds to naming the holder of the second card of the proposed key set. Finally, a winning end configuration in the stick game corresponds to a signature in which the key set system succeeds (since the team is of size at most one), and a losing end configuration in the stick game corresponds a configuration in which no key set is possible (since the team has size two or greater and all but possibly one player has run out of cards).

We show that a chooser strategy $\mathcal{W}$ succeeds on a given configuration in the stick game if and only if the $\mathcal{V}$ key set system succeeds for the corresponding signature. We first show that we can restrict out attention in the key set system to any particular legal ξ-deal. To show that a given $\mathcal{V}$ key set system succeeds for a given signature ξ, it is sufficient to show that every run on every legal ξ-deal succeeds. The following lemma shows that one can in fact restrict attention to the runs that result from a single arbitrary legal ξ-deal. The proof uses a simple strategy stealing argument.

Lemma 6.15 *Let ξ be a signature, let δ be a legal ξ-deal, let $\mathcal{V}$ be a proposing rule, and let $\mathcal{P}_\mathcal{S}$ be a $\mathcal{V}$ key set system. If $\mathcal{P}_\mathcal{S}$ succeeds for δ, then it succeeds for ξ.*

Proof: Suppose the conditions of the lemma, and suppose that $\mathcal{P}_\mathcal{S}$ succeeds for δ. We show that $\mathcal{P}_\mathcal{S}$ also succeeds for δ'. Since δ' is an arbitrary legal ξ-deal, it will follow that $\mathcal{P}_\mathcal{S}$ succeeds for ξ.

Inductively assume that for all legal $\overline{\xi}$-deals $\overline{\delta}$ and $\overline{\delta}'$ such that size$(\overline{\xi}) <$ size(ξ), if the $\mathcal{P}_\mathcal{S}$ succeeds for $\overline{\delta}$, then it succeeds for $\overline{\delta}'$. Suppose that δ is dealt from deck Δ, and δ' is dealt from Δ', and let f be a player-preserving bijection from Δ' to Δ, where a bijection $f\colon \Delta' \to \Delta$ is said to be *player-preserving* if $x \in h_i' \Rightarrow f(x) \in h_i$.

We simultaneously construct two runs of the $\mathcal{V}$ key set protocol, one on δ and one on δ'. To distinguish the two, we refer to them as G and G' respectively. If the G' team cannot propose a key set, then G' terminates. Since G and G' have the same signature, G also terminates, and either both succeed or both fail. Otherwise, a key set $\{x, y\}$ is proposed in G' by some player, say P_i. Since a proposing rule is a function only of the current signature, P_i is also chosen to propose in G. Moreover, it is possible for the proposed key set in G to be $\{f(x), f(y)\}$ since, in the deal δ, $f(x)$ is held by P_i and $f(y)$ is not. Furthermore, since the same player holds y in δ' as holds $f(y)$ in δ, the deals δ_1 and δ'_1 that result from one iteration of the $\mathcal{V}$ key set system in G and G', respectively, both have the same signature, and the function f restricted to the remaining cards is still a player-preserving bijection.

Since $\mathcal{P}_\mathcal{S}$ succeeds for δ, it must succeed for δ_1. By induction, $\mathcal{P}_\mathcal{S}$ therefore succeeds for δ'_1. Thus, since we considered an arbitrary choice of key set in G', $\mathcal{P}_\mathcal{S}$ succeeds for δ'. ■

Lemma 6.16 (Stick Game Lemma) *Let ξ be a signature and let $\mathcal{V}$ be a proposing rule. Let the stick game strategy $\mathcal{W}$ correspond to the proposing rule $\mathcal{V}$ and let $\mathcal{P}_\mathcal{S}$ be a $\mathcal{V}$ key set ξ-system. Then $\mathcal{P}_\mathcal{S}$ succeeds for $\xi = (s_1, \ldots, s_k; d)$ if and only if $\mathcal{W}$ succeeds on the configuration $(s_1, \ldots, s_k; d - \sum s_i; 0)$.*

Proof: It is obvious that if the chooser wins in the stick game, then the corresponding key set system succeeds. In the other direction, there is a complication in that the adversary in the stick game has complete control over which pile to match, whereas in the key set system, the hand containing the matching card is determined by the deal. However, since the key set system is required to succeed on *all* runs of the given deal, it must in particular succeed when the matching card of the proposed key set is in the hand corresponding to the pile chosen by the adversary in the stick game.

Specifically, let $\xi = (s_1, \ldots, s_k; d)$, let $\mathcal{P}_\mathcal{S}$ be a $\mathcal{V}$ key set ξ-system, let δ be a ξ-deal, and suppose that $\mathcal{P}_\mathcal{S}$ succeeds for ξ. Let $e = d - \sum s_i$ and let $C = (s_1, \ldots, s_k; e; 0)$, so C is the stick game configuration corresponding to ξ. Inductively assume that the result holds for all configurations of size smaller than size(C).

Since $\mathcal{P}_\mathcal{S}$ succeeds for δ, either $k \leq 1$ or it is possible for the team to propose a key set. If $k \leq 1$, then by Proposition 6.5, C is winning. Otherwise, let P_i be the player chosen by $\mathcal{V}$, and consider the corresponding configuration $C_1 = (s_1, \ldots, s_k; e; i)$. Since it is possible for P_i to propose a key set, this implies that either $s_j > 0$ for some

$j \neq i$, or $e > 0$. Therefore, the stick game adversary can make a move, and matches with some pile, either P_j for $j \neq i$ or E, resulting in configuration C_2.

In either case, it is possible in the key set protocol for P_i to propose a key set in which the location of the second card corresponds to the matching pile, resulting in some deal δ' of signature ξ'. Since $\mathcal{P}_\mathcal{S}$ succeeds for ξ, it must succeed for all runs of all ξ-deals. In particular, it succeeds for all runs of δ, and hence must succeed for δ'. By Lemma 6.15, $\mathcal{P}_\mathcal{S}$ therefore succeeds for ξ'. The signature ξ' corresponds to C_2, so by induction, C_2 is winning, and hence C is also winning. ■

By Lemma 6.16, $\mathcal{P}_\mathcal{S}$ succeeds whenever the SVP stick game strategy is a successful chooser strategy for the corresponding stick game. In addition, the optimality of the SVP stick game strategy implies the optimality of the SVP key set protocol within the class of all key set protocols.

Corollary 6.17 *The SVP key set protocol is an optimal key set protocol.*

Proof: Let $\xi = (s_1, \ldots, s_k; d)$ and let $\mathcal{V}$ be a proposing rule. let $\mathcal{P}_\mathcal{S}$ be a SVP key set ξ-system, let $\mathcal{P'}_{\mathcal{S'}}$ be a $\mathcal{V}$ key set ξ-system, and suppose that $\mathcal{P'}_{\mathcal{S'}}$ succeeds for $\xi = (s_1, \ldots, s_k; d)$. Then by Lemma 6.16 (Stick Game Lemma), $C = (s_1, \ldots, s_k; d - \sum s_i; 0)$ is a winning configuration. Then by Theorem 6.14, the SVP strategy for the stick game succeeds on C. Finally, by Lemma 6.16, $\mathcal{P}_\mathcal{S}$ succeeds for ξ. ■

CHAPTER 7: THE TRANSFORMATION PROTOCOL

The family of key set protocols presented in Chapter 6 has two limitations: first, to guarantee success, the key set protocol requires that the team hold more than half the cards in the deck; second, it only provides a 1-bit secret key. Moreover, it is not obvious how to modify the protocol to overcome these limitations. For example, the protocol cannot be repeated to obtain additional key bits since players drop out and expose their remaining cards during execution.

We present a system, the transformation system, consisting of the transformation protocol and a uniform card game source, that overcomes both these limitations. Specifically, for teams of two players who each receive a fixed fraction β of the cards, the transformation system performs n-bit perfect secret key exchange against the extra cards view, provided the deck is sufficiently large. Using the transformation system, we show there is a function $d_0(n, \beta) = O(n(1/\beta)^{2.71})$ such that for any real number $\beta > 0$ and signature $\xi = (s_1, \ldots, s_k; d)$, if $d \geq d_0(n, \beta)$ and each $s_i \geq \lfloor \beta d \rfloor$, there exists a card game ξ-system for perfect n-bit secret key exchange against the extra cards view. It follows that $\mathrm{pcap}(\xi) \geq 2^n$.

Fischer, Paterson, and Rackoff [FPR91] exhibit a protocol for the case that $k = 2$ and $n = 1$, but their required deck size grows super-polynomially in $1/\beta$. Furthermore, although their solution performs weak secret bit exchange, it does not perform perfect secret bit exchange as they claimed because of a subtle bug in their construction. (Perfect secrecy was violated by the use of key sets that in some cases would not be opaque.) Our formal model led to the discovery of their bug.

Let ξ be a signature. The transformation ξ-system is presented in §7.1 and analyzed in §7.2. The analysis is based on a nontrivial potential argument and shows that the system performs n-bit perfect secret key exchange against the extra cards view whenever the potential of ξ is sufficiently large.

In §7.3, we consider applications of the transformation system. We show that for any $n \geq 1$ and $0 < \beta \leq 1/k$, if $d \geq d_0(n, \beta)$ then the transformation system is a card game $(\lfloor \beta d \rfloor, \lfloor \beta d \rfloor; d)$-system for n-bit perfect secret bit exchange against the extra cards view. It follows that for any $n \geq 1$ and $0 < \beta \leq 1/k$, if $d \geq d_0(n, \beta)$, then $\mathrm{pcap}(\lfloor \beta d \rfloor, \lfloor \beta d \rfloor; d) \geq 2^n$.

Second, we show a general reduction of the multiparty case to the two player case. Applying this reduction to the transformation system yields a system that performs n-bit perfect secret key exchange for teams of arbitrary size k, where each team player receives fraction β of the cards, provided the deck is sufficiently large. The required deck size is again only $O(n(1/\beta)^{2.71})$. Applying this to the case where the deck is initially divided evenly between m players, the deck size needed to guarantee that any team that forms will be able to obtain an n-bit perfect secret key is $O(nm^{2.71})$.

7.1 The Protocol

Consider a team of two players, Alice (P_1) and Bob (P_2). Fix a signature $\xi = (a, b; d)$ and a deck Δ, and let Alice and Bob be dealt a random ξ-deal δ of Δ. We extend the concept of key sets to larger sets, which we call portions. Specifically, a set of cards $S \subseteq \Delta$ is called an (s, i, j)*-portion* (relative to δ) if $|S| = s$ and S contains exactly i cards from Alice's hand and exactly j cards from Bob's hand. The remaining $s - i - j$ cards belong to Eve. We sometimes refer to an $(s, 1, 1)$-portion as an s-portion, and to any (s, i, j)-portion simply as a portion. Hence a 2-portion is a key set. An (s, i, j)-portion is *useful* if $i, j \geq 1$.

An (s, j_1, j_2)-portion S is opaque if Eve does not know anything about the location of the cards in S that she does not hold, other than the information provided by the fact that S is an (s, j_1, j_2)-portion. More formally, let $\tau \in \mathrm{c}$. Recall that since we are dealing with the extra cards view, Eve's view a_e is a hand. We say S is *opaque with respect to* τ if

$$\Pr\left[\mathrm{E}(a_e) \cap \mathrm{C}(\tau) \cap \mathrm{B}_i(x)\right] = \Pr\left[\mathrm{E}(a_e) \cap \mathrm{C}(\tau)\right] \cdot \frac{j_i}{j_1 + j_2}$$

for all $a_e \in U_e$, $i \in K$, and $x \in S - a_e$.

The *transformation protocol* maintains a collection $\mathcal{C}$ of pairwise disjoint, useful portions that are opaque with respect to the conversation so far. The portions in $\mathcal{C}$ are common knowledge to Alice, Bob, and Eve at all times. The deck Δ is a (d, a, b)-portion relative to δ since δ is an $(a, b; d)$-deal. The initial collection, $\mathcal{C}_0$, is the singleton collection $\{\Delta\}$.

A step of the protocol modifies the current collection $\mathcal{C}$ by removing one or more portions from $\mathcal{C}$ and then adding one or more smaller new portions to $\mathcal{C}$ according to a rule called a *transformation.* A transformation is *applicable* to $\mathcal{C}$ if $\mathcal{C}$ contains portions satisfying the preconditions of the transformation. A collection $\mathcal{C}$ is *terminal* if no transformation is applicable to it.

The protocol operates in phases, where each phase consists of two or three rounds of communication that apply a transformation to $\mathcal{C}$. For any conversation τ in which a phase has just been completed, collection(τ) is the collection that results from $\mathcal{C}_0$ by applying the transformations carried out by the phases of τ. The protocol runs for $t = 3(d-1)$ rounds. If a terminal collection is reached before t rounds or if there are not enough rounds left to complete the next phase, dummy rounds are added until t rounds are executed. If more than one transformation is applicable at a given round, the players can agree which transformation to apply using any prearranged deterministic function of the signature and the current conversation.

Once t rounds have been executed, the team players label the 2-portions in $\mathcal{C}$, in lexicographic order, $X_1, X_2, \ldots, X_n$. Each 2-portion X_i is used to determine a perfectly secret bit v_i by the truth value of the statement "Alice holds the lower card in X_i." The secret bits are concatenated together to form a single n-bit perfectly secret key $v_1 v_2 \cdots v_n$. Note that no further communication is required for this step.

The two types of transformations are *splitting* and *combining*. Splitting replaces an (s,i,j)-portion in $\mathcal{C}$ with several smaller portions, each of which contains exactly one of Alice's cards if $i \geq j$, and exactly one of Bob's cards if $i < j$. Combining replaces two $(s,1,1)$-portions by a single $(s',1,1)$-portion for some $s' < s$.

Splitting: An (s,i,j)-portion S in $\mathcal{C}$ can be split if $i + j \geq 3$. If $i \geq j$, the splitting proceeds as described below. If $i < j$, the roles of Alice and Bob are reversed.

1. S is removed from $\mathcal{C}$.
2. Alice randomly partitions S into i sets, each of size $\lfloor s/i \rfloor$ or $\lceil s/i \rceil$, such that she holds exactly one card in each set, and announces the sets.
3. Bob says how many cards he holds in each set announced by Alice.
4. Each set in which Bob holds at least one card is added to $\mathcal{C}$.

Combining: Two $(s,1,1)$-portions S_1 and S_2 can be combined if $s \geq 3$.

1. S_1 and S_2 are removed from $\mathcal{C}$.
2. Alice randomly chooses $p \in \{1,2\}$. Let $q = 3 - p$.
3. Alice constructs and announces a new set T consisting of her card from S_p, $\lfloor s/3 \rfloor - 1$ randomly chosen cards that are not hers from S_p, and $\lfloor s/3 \rfloor$ randomly chosen cards that are not hers from S_q.

4. Bob announces how many cards he holds in T.
   (a) If Bob holds no cards in T, then Alice announces the set difference $S_q - T$, which is added to $\mathcal{C}$.
   (b) If Bob holds one card in T, then T is added to $\mathcal{C}$.
   (c) If Bob holds two cards in T, then Alice announces $S_p \cap T$, which is added to $\mathcal{C}$.

Note that each transformation can be carried out using three or fewer rounds of communication.

Lemma 7.1 *Suppose $a, b \geq 1$ and let σ be a feasible complete conversation of the transformation system. Then* collection(σ) *is a collection of disjoint, useful portions that are opaque with respect to σ.*

Proof (sketch): Since the initial collection $\mathcal{C}_0$ is a singleton collection containing a single useful portion that is opaque with respect to the empty conversation, it suffices to show that each phase preserves the disjointness, usefulness and opaqueness of the portions in a collection with respect to the conversation so far. Let τ be a feasible conversation in which a phase has just been completed. Suppose that $\mathcal{C}$ = collection(τ) is a collection of disjoint, useful portions that are opaque with respect to τ and consider the phase τ' that performs the next transformation, resulting in the conversation $\tau \cdot \tau'$. First note that the portions not removed from $\mathcal{C}$ are disjoint, useful, and opaque with respect to $\tau \cdot \tau'$ because τ' is completely independent of those portions.

Suppose τ' results in an (s, i, j)-portion S in $\mathcal{C}$ being split. Splitting preserves the disjointness property because S is removed and the new portions added are disjoint subsets of S. Each portion added to $\mathcal{C}$ is useful because, if $i \geq j$, then Alice holds exactly one card in each set she announces, and only those sets in which Bob holds at least one card are added. Similarly, if $i < j$, then Bob holds exactly one card and Alice holds at least one card in each portion added to $\mathcal{C}$. To see that splitting produces only portions that are opaque with respect to $\tau \cdot \tau'$, suppose that $i \geq j$, and consider an $(s', 1, j')$-portion S' that is added to $\mathcal{C}$. Since S is opaque with respect to τ and Alice randomly chose the partition for S among all partitions in which she holds one card in each subset, Alice's card is equally likely to be any of the $j' + 1$ cards not held by Eve in S', given the conversation $\tau \cdot \tau'$. It follows that also Bob holds any one of the $j'+1$ cards not held by Eve with probability j'. It follows that S' is opaque with respect to $\tau \cdot \tau'$.

Now consider the combining of two i-portions S_1 and S_2, resulting in the portion S'. Combining preserves the disjointness property because all the elements in S' are in either S_1 or S_2, both of which have been removed from $\mathcal{C}$. To see that S' is useful, we must consider the possible outcomes. Let p be as chosen by Alice in the process of combining, let $q = 3 - p$, and let T be the new set constructed by Alice. Then Alice and Bob both hold one card in each of S_p and S_q, and in particular, Alice holds one card in $S_p \cap T$ and one card in $S_q - T$. If Bob holds no cards in T, Bob's card in S_q must lie in $S_q - T$. Therefore $S_q - T$, which is added to $\mathcal{C}$, is useful. If Bob holds one card in T, then T, which is added to $\mathcal{C}$, is useful. If Bob holds two cards in T, then one of them must lie in $S_p \cap T$ and one in $S_q \cap T$. Therefore $S_p \cap T$, which is added to $\mathcal{C}$, is useful. To see that S' is opaque with respect to $\tau \cdot \tau'$, suppose that Alice holds x in S', Bob holds y in S', Alice holds card x_1 in S_1 and card x_2 in S_2, and Bob holds card y_1 in S_1 and card y_2 in S_2, where $x \in \{x_1, x_2\}$ and $y \in \{y_1, y_2\}$. Then the conversation $\tau \cdot \tau'$ can occur only in the cases that (1) Alice holds x and Bob holds y or (2) Alice holds y and Bob holds x. In the first case, Alice holds x in S', while in the second case, Bob holds x in S'. Since these two cases occur with equal probability, it follows that S' is opaque with respect to $\tau \cdot \tau'$. ∎

7.2 Analysis

We use a "potential" argument to analyze how many 2-portions are produced by the transformation system. Given a collection $\mathcal{C}$ of useful portions, we define a quantity $\phi(\mathcal{C})$, called its *potential*. We show that if $\mathcal{C}'$ results from $\mathcal{C}$ by any transformation, then $\phi(\mathcal{C}') \geq \phi(\mathcal{C})$. Thus, if $\mathcal{C}'$ results from $\mathcal{C}$ via any sequence of transformations, then $\phi(\mathcal{C}') \geq \phi(\mathcal{C})$. We define a constant W, and we show that if $\phi(\mathcal{C}) > W + 2p$, where p is the number of 2-portions in $\mathcal{C}$, then at least one transformation is applicable to $\mathcal{C}$. It follows from the above claims that if $\mathcal{C}$ is terminal and $\phi(\mathcal{C}) > W + 2(n-1)$ then the number of 2-portions in $\mathcal{C}$ is at least n. Finally, we show that every trace from a finite collection of portions is bounded in length since each transformation reduces the difference between total size of all portions and the number of portions. Since t is large enough, it follows that every run of the transformation system results in a terminal collection. Hence, if $a, b \geq 1$ and $d \geq a + b$, the transformation system is a card game $(a, b; d)$-system for $\lceil(\phi(\{\Delta\}) - W)/2\rceil$-bit perfect secret key exchange against the extra cards view. The remainder of this section defines $\phi(\mathcal{C})$ and proves the above claims.

The constant $c = \log_{3/2} 2 = 1/\log_2(3/2) < 1.7096$ is used throughout the analysis. Note that $(2/3)^{-c} = 2$ and $1 - 1/c > 0$. Given an (s,i,j)-portion S, we recursively define $\phi(S) = \phi(s,i,j) =$

$$\begin{cases} 2 & \text{if } s = 2,\ i = j = 1 \\ (s-2)^{-c} & \text{if } s \geq 3,\ i = j = 1 \\ j\,\phi(\lceil s/i \rceil, 1, 1) & \text{if } i \geq j,\ i \geq 2 \\ \phi(s,j,i) & \text{if } i < j \end{cases}$$

Hence, $\phi(s,i,j)$ is symmetric in its last two arguments, and $\phi(s,1,1)$ is monotonically decreasing in s for all integers $s \geq 2$. We extend the potential function to collections $\mathcal{C}$ of useful portions by defining $\phi(\mathcal{C}) = \sum_{S \in \mathcal{C}} \phi(S)$.

Fact 7.2 *Let y, z be integers, $z \neq 0$. Then $\lceil y/z \rceil \leq (y-1)/z + 1$.*

Proof: We have $y = qz + r$ for integers q and r such that $0 < r \leq z$, so $\lceil y/z \rceil = q + 1 \leq q + (r-1)/z + 1 = (qz + r - 1)/z + 1 = (y-1)/z + 1$. ∎

In analyzing the splitting transformation, we will need the following lemma relating the potential of an s-portion to the potential of a $\lceil s/b \rceil$-portion.

Lemma 7.3 *Let b be an integer such that $1 \leq b \leq s - 1$. Then $b\,\phi(s,1,1) \leq \phi(\lceil s/b \rceil, 1, 1)$.*

Proof: Let b be an integer such that $1 \leq b \leq s-1$. If $b = 1$, then trivially $b\,\phi(s,1,1) = \phi(\lceil s/b \rceil, 1, 1)$.

Otherwise, $2 \leq b \leq s - 1$, and thus $\phi(s,1,1) = (s-2)^{-c}$ and $\lceil s/b \rceil \geq 2$. If $\lceil s/b \rceil = 2$ then

$$\begin{aligned} b\,\phi(s,1,1) &= b(s-2)^{-c} \\ &\leq (s-1)(s-2)^{-c} \\ &\leq 2(s-2)^{(1-c)} \\ &\leq 2 \\ &= \phi\left(\left\lceil \frac{s}{b} \right\rceil, 1, 1\right) \end{aligned}$$

as desired. Otherwise, $\lceil s/b \rceil \geq 3$, so

$$\begin{aligned} b\,\phi(s,1,1) &= b(s-2)^{-c} & (7.1)\\ \phi(\lceil s/b \rceil, 1, 1) &= (\lceil s/b \rceil - 2)^{-c} & (7.2)\end{aligned}$$

Since $(1/c) - 1 < 0 < c$, we have $b^{(1/c)-1} < 1 < b^{1/c}$. Also $s - 1 > b$ since $\lceil s/b \rceil \geq 3$, so

$$\begin{aligned} b(s-2)^{-c} &= b((s-1)-1)^{-c} \\ &< b\left(b^{(1/c)-1}(s-1) - b^{1/c}\right)^{-c} \\ &= \left(\frac{s-1}{b} - 1\right)^{-c} & (7.3)\end{aligned}$$

By Fact 7.2,

$$\left(\frac{s-1}{b} - 1\right)^{-c} \leq \left(\left\lceil \frac{s}{b} \right\rceil - 2\right)^{-c} \tag{7.4}$$

Combining (7.1) through (7.4) yields the desired result. ∎

Lemma 7.4 *Suppose C' results from C by a splitting transformation. Then $\phi(C') \geq \phi(C)$.*

Proof: It suffices to show that the potential of the portion to be split is no more than the total potential of the resulting portions. Let S be an (s,x,y)-portion, and suppose without loss of generality that $x \geq y$ (the case $x < y$ is symmetric). Let $S_1, S_2, \ldots, S_\ell$ be the portions added to C' as a result of splitting S, where S_i is an $(s_i, 1, b_i)$-portion.

Since each S_i was added to C', it follows that $1 \leq b_i \leq s_i - 1$ and $s_i \geq 2$. Also, $s_i \in \{\lfloor s/x \rfloor, \lceil s/x \rceil\}$, so $2 \leq s_i \leq \lceil s/x \rceil$. Thus, by the monotonicity of $\phi(s,1,1)$ and Lemma 7.3,

$$\begin{aligned} b_i\,\phi\left(\left\lceil \frac{s}{x} \right\rceil, 1, 1\right) &\leq b_i\,\phi(s_i, 1, 1) \\ &\leq \phi\left(\left\lceil \frac{s_i}{b_i} \right\rceil, 1, 1\right) \\ &= \phi(s_i, 1, b_i)\end{aligned}$$

Therefore,

$$\begin{aligned} \phi(S) &= \phi(s,x,y) \\ &= y\,\phi\left(\left\lceil \frac{s}{x} \right\rceil, 1, 1\right)\end{aligned}$$

$$\begin{aligned}
&= \sum_{i=1}^{\ell} b_i \phi\left(\left\lceil \frac{s}{x} \right\rceil, 1, 1\right) \\
&\leq \sum_{i=1}^{\ell} \phi(s_i, 1, b_i) \\
&= \sum_{i=1}^{\ell} \phi(S_i)
\end{aligned}$$

as desired. ■

Lemma 7.5 *Suppose $\mathcal{C}'$ results from $\mathcal{C}$ by a combining transformation. Then $\phi(\mathcal{C}') \geq \phi(\mathcal{C})$.*

Proof: As before, we need only compare the potential of the portions that are combined to the potential of the resulting portion.

Let S_1 and S_2 be s-portions, and suppose S' is an s'-portion resulting from combining S_1 and S_2. In order for combining to be possible, we must have $s \geq 3$. Hence $\phi(S_1) + \phi(S_2) = 2(s-2)^{-c} \leq 2$.

Let T be the new set constructed by Alice. Then $|T| = 2\lfloor s/3 \rfloor$. If Bob holds no cards in T, then an $(s - \lfloor s/3 \rfloor)$-portion is added to $\mathcal{C}$. If Bob holds one card in T, then a $(2\lfloor s/3 \rfloor)$-portion is added to $\mathcal{C}$. If Bob holds two cards in T, then a $(\lfloor s/3 \rfloor)$-portion is added to $\mathcal{C}$. In all cases, we have $s' \leq \lceil 2s/3 \rceil$.

If $s' = 2$, then $\phi(S') = 2$, so $\phi(S_1) + \phi(S_2) \leq \phi(S')$, as desired. Otherwise, $2 < s' \leq \lceil 2s/3 \rceil$. Using this and Fact 7.2, we have

$$\begin{aligned}
\phi(S') &= (s'-2)^{-c} \\
&\geq (\lceil 2s/3 \rceil - 2)^{-c} \\
&\geq (2s/3 + 2/3 - 2)^{-c} \\
&= 2(s-2)^{-c} \\
&= \phi(S_1) + \phi(S_2)
\end{aligned}$$

completing the proof. ■

Let $W = \sum_{s=3}^{\infty}(s-2)^{-c}$. Since $c > 1$, this series converges and W is finite. Numerical analysis shows that $2.0356 < W < 2.0358$. Given a collection $\mathcal{C}$, we define $\pi(\mathcal{C})$ to be the number of 2-portions in $\mathcal{C}$.

Lemma 7.6 *If $\mathcal{C}$ is terminal, then $\phi(\mathcal{C}) \leq W + 2\pi(\mathcal{C})$.*

Proof: Let $\mathcal{C}$ be a collection of portions such that $\phi(\mathcal{C}) > W + 2\pi(\mathcal{C})$. We show that $\mathcal{C}$ is not terminal. Since $\phi(\mathcal{C}) > 0$, $\mathcal{C}$ is nonempty. If $\mathcal{C}$ contains an (s, i, j)-portion such that $i + j \geq 3$, then splitting is possible. Otherwise, each portion S_i in $\mathcal{C}$ is an s_i-portion for some $s_i \geq 2$. In order to satisfy $\phi(\mathcal{C}) > W + 2\pi(\mathcal{C})$, it must be the case that there are two s-portions in $\mathcal{C}$ for some $s \geq 3$, since $W + 2\pi(\mathcal{C})$ is the potential of a collection containing $\pi(\mathcal{C})$ 2-portions and one s-portion for every $s \geq 3$. Thus combining can be applied. ■

A sequence $T = \mathcal{C}_0, \mathcal{C}_1, \mathcal{C}_2, \ldots$ of collections of portions is called a *trace* from $\mathcal{C}_0$ if for each $i \geq 1$, $\mathcal{C}_i$ can result from $\mathcal{C}_{i-1}$ by some transformation. We define $|T|$ to be the number of collections in the sequence T. For any collection $\mathcal{C}$, we define

$$M(\mathcal{C}) = \sum_{S \in \mathcal{C}} (|S| - 1) = \sum_{S \in \mathcal{C}} |S| - |\mathcal{C}|$$

Lemma 7.7 *Let $\mathcal{C}$ be a finite nonempty collection of useful portions and let T be a trace from $\mathcal{C}$. Then $|T| \leq M(\mathcal{C})$.*

Proof: Let $\mathcal{C}$ be a finite collection of useful portions. If $\mathcal{C}$ is terminal, then no transformations are applicable to $\mathcal{C}$, so $|T| = 1$. Since $\mathcal{C}$ is nonempty and every portion in $\mathcal{C}$ is useful (and hence contains at least two cards), $M(\mathcal{C}) \geq 1$. Hence $|T| \leq M(\mathcal{C})$.

Otherwise, some transformation is applicable to $\mathcal{C}$. Suppose $\mathcal{C}'$ is the next collection in the trace from $\mathcal{C}$. Then $\mathcal{C}'$ results from $\mathcal{C}$ by one transformation. Inductively, it suffices to show that $M(\mathcal{C}') \leq M(\mathcal{C}) - 1$. We denote $\sum_{S \in \mathcal{C}} |S|$ by s, $\sum_{S \in \mathcal{C}'} |S|$ by s', $|\mathcal{C}|$ by c, and $|\mathcal{C}'|$ by c'. Hence $M(\mathcal{C}) = s - c$ and $M(\mathcal{C}') = s' - c'$.

Suppose $\mathcal{C}'$ results from $\mathcal{C}$ by a splitting transformation of an (s, i, j)-portion S with $i \geq j$. (The case $i < j$ is symmetric). In a splitting transformation, one portion gets removed and at least one portion gets added, so $c' \geq c$. The portions added are disjoint subsets of the removed portion, so $s' \leq s$. In order for S to be split, $i + j \geq 3$, so $i \geq 2$. If $c' = c$, then $s' \leq s - 1$, since in this case at least one of Alice's $i \geq 2$ cards is discarded. If $s' = s$, then $c' \geq c + 1$, since in this case all $i \geq 2$ portions are returned to the collection. It follows that $M(\mathcal{C}') \leq M(\mathcal{C}) - 1$.

If $\mathcal{C}'$ results from $\mathcal{C}$ by a combining transformation, then two portions are removed and one is added, so $c' = c - 1$. Also $s' \leq s - 2$ since $\bigcup_{S \in \mathcal{C}} S - \bigcup_{S \in \mathcal{C}'} S$ contains one card from each of Alice's and Bob's hand. Thus $s' - c' \leq s - c - 1$, so $M(\mathcal{C}') \leq M(\mathcal{C}) - 1$. ■

Since each transformation can be implemented using three rounds of communication, and there are $3(d-1) = 3 \cdot M(\mathcal{C}_0)$ rounds, it follows from Lemma 7.7 that if σ is a feasible complete conversation of the transformation system, then collection(σ) is terminal.

Lemma 7.8 *Let $\mathcal{C}$ be a terminal collection of useful portions. If $\phi(\mathcal{C}) > W+2(n-1)$, then $\pi(\mathcal{C}) \geq n$.*

Proof: Let $\mathcal{C}$ be a terminal collection of useful portions such that $\phi(\mathcal{C}) > W+2(n-1)$. Since $\mathcal{C}$ is terminal, Lemma 7.6 implies $\phi(\mathcal{C}) \leq W + 2\pi(\mathcal{C})$. Thus $W + 2(n-1) < W + 2\pi(\mathcal{C})$, and hence $\pi(\mathcal{C}) > n - 1$, so $\pi(\mathcal{C}) \geq n$. ■

Theorem 7.9 *Let $a, b \geq 1$, $d \geq a+b$, and $\phi(d, a, b) > W+2(n-1)$. Then the transformation system is a card game $(a, b; d)$-system for n-bit perfect secret key exchange against the extra cards view.*

Proof: Assume the conditions of the theorem. Consider a feasible complete conversation σ of the transformation system and let $\mathcal{C}' =$ collection(σ). The initial collection $\mathcal{C}_0$ contains the single (d, a, b)-portion Δ. Thus, $\phi(\mathcal{C}_0) = \phi(d, a, b) > W + 2(n-1)$. By Lemma 7.1, $\mathcal{C}'$ is a collection of disjoint, useful, opaque portions with respect to σ. By Lemmas 7.4 and 7.5, $\phi(\mathcal{C}') \geq \phi(\mathcal{C}_0)$. Since $a, b \geq 1$, Δ is useful. Hence, by Lemma 7.7, $\mathcal{C}'$ is terminal. Thus, $\phi(\mathcal{C}') > W + 2(n-1)$. Hence by Lemma 7.8, $\pi(\mathcal{C}') \geq n$. Since, in particular, all the 2-portions in $\mathcal{C}'$ are opaque with respect to σ, the output sequence is a perfect secret key of length at least n. ■

Corollary 7.10 *Let $a, b \geq 1$ and let $d \geq a + b$. Then the transformation system is a card game $(a, b; d)$-system for $\lceil(\phi(d, a, b) - W)/2\rceil$-bit perfect secret key exchange against the extra cards view.*

Proof: Assume the conditions of the corollary and let $\phi_0 = \phi(d, a, b)$. It follows that $n = \lceil(\phi(d, a, b) - W)/2\rceil = \lceil(\phi_0 - W)/2\rceil$. By Fact 7.2, $n \leq (\phi_0 - W + 1)/2$. Thus $W + 2(n-1) \leq \phi_0 - 1 < \phi_0$. By Theorem 7.9, the transformation system performs n-bit perfect secret key exchange against the extra cards view, as desired. ■

Since the extra cards view is a refinement of the empty view, Lemma 3.4 yields the following corollary to Corollary 7.10.

Corollary 7.11 *Let $a, b \geq 1$ and let $d \geq a + b$. Then* $\text{pcap}(a, b; d) \geq 2^{\lceil(\phi(d,a,b)-W)/2\rceil}$.

Note that the bounds given by Theorem 7.9, Corollary 7.10, and Corollary 7.11 may not be tight since the inequalities in Fact 7.2, Lemma 7.4, Lemma 7.5, and Lemma 7.8 may not be tight. Even if these inequalities are tight, it might be possible to get a tighter bound on $\mathrm{pcap}(a, b; d)$ by finding another potential function ϕ' and constant W' such that the potential ϕ' does not decrease when any transformation is applied, the potential $\phi'(\mathcal{C}) \leq W' + 2\pi(\mathcal{C})$ for any terminal collection $\mathcal{C}$, and $\phi'(d, a, b) - W' > \phi(d, a, b) - W$.

The results of Theorem 7.9 and Corollary 7.10 are based on analysis of the worst case. That is, if $a, b \geq 1$, $d \geq a+b$ and $n = \lceil(\phi(d, a, b) - W)/2\rceil$, then the transformation system is guaranteed to yield n secret bits when the deal is a legal $(a, b; d)$-deal. In general, the expected or best case number of bits may be higher than the worst case. However, due to the symmetric nature of combining, which takes two portions of size s and yields in the worst case a portion of size approximately $(2/3)s$ and in the best case a portion of size approximately $(1/3)s$, we conjecture that the best and average cases are within a constant factor of the worst case.

7.3 Applications

In this section, we present two applications of the transformation protocol. The first obtains a much improved bound for a problem studied in [FPR91] in which each player holds a constant fraction of the cards. The second uses the transformation protocol as a building block for constructing a multiparty perfect secret key exchange system.

7.3.1 Two Players Each Holding a Fraction of the Cards

We consider the situation in which each of two players receives a constant fraction β of the cards in the deck, and the remainder go to Eve. This situation arises naturally with $\beta = 1/m$, for example, in systems where the deck is dealt out evenly to m players. We are interested in how large the deck must be in order for the transformation protocol to work in this situation. We use the following in our analysis.

Fact 7.12 *Let x be a positive integer and let $\beta \geq 1/x$ be a real number. Then $\lceil x/\lfloor\beta x\rfloor\rceil < 2/\beta + 1$.*

Proof: Let $\ell = \lfloor\beta x\rfloor$. Then $1 \leq \ell \leq \beta x < \ell + 1$, so $x < (\ell+1)/\beta$ and $(\ell+1)/\ell \leq 2$. Thus $x/\ell \leq (\ell+1)/(\beta\ell) \leq 2/\beta$. It follows that $\lceil x/\lfloor\beta x\rfloor\rceil = \lceil x/\ell\rceil \leq \lceil 2/\beta\rceil < 2/\beta+1$.

■

The following shows that the transformation system performs perfect n-bit secret key exchange for two players each holding a fraction β of the cards if the deck is sufficiently large with respect to n and $1/\beta$. The required deck size is only $O(n(1/\beta)^{(c+1)})$, which is polynomial in $1/\beta$ and linear in n. Recall that $c = \log_{3/2} 2$ and $W = \sum_{s=3}^{\infty}(s-2)^{-c}$, and let $c_1 = 2^{(c+1)} < 6.5411$ and $c_2 = (W-2)/2 + 2^{-c}/c_1$. Calculation shows that $0.0645 < c_2 < 0.0647$.

Theorem 7.13 *Let $0 < \beta \leq 1/2$, $n \geq 1$, and suppose that $d \geq c_1(1/\beta)^{(c+1)}(n + c_2)$. Then the transformation system is a card game $(\lfloor \beta d \rfloor, \lfloor \beta d \rfloor; d)$-system for n-bit perfect secret key exchange against the extra cards view.*

Proof: Let β, d, and $\xi = (\lfloor \beta d \rfloor, \lfloor \beta d \rfloor; d)$ satisfy the conditions of the theorem. By Theorem 7.9, it suffices to show that $\lfloor \beta d \rfloor \geq 1$, $d \geq 2\lfloor \beta d \rfloor$ and $\phi(d, \lfloor \beta d \rfloor, \lfloor \beta d \rfloor) > W + 2(n-1)$.

Since $\beta \leq 1/2$, it follows that $2\lfloor \beta d \rfloor \leq 2\beta d \leq d$, as desired. Furthermore, $(2\beta)^{-c} \geq 1$ and $2/\beta \geq 4$. Since $n + c_2 \geq 1$, using the bound on d and the definition of c_1 gives

$$\begin{aligned}
\beta d &\geq c_1 \beta^{-c}(n + c_2) \\
&\geq 2\left(\frac{2}{\beta}\right)^{c} \\
&\geq 8
\end{aligned}$$

So $\lfloor \beta d \rfloor \geq 1$, as desired.

We now establish that $\phi(d, \lfloor \beta d \rfloor, \lfloor \beta d \rfloor) > W + 2(n-1)$. We begin by examining $\phi(\lceil d/\lfloor \beta d \rfloor \rceil, 1, 1)$. Note that we have $\lceil d/\lfloor \beta d \rfloor \rceil \geq d/\beta d = 1/\beta \geq 2$. If $\lceil d/\lfloor \beta d \rfloor \rceil = 2$, then $\phi(\lceil d/\lfloor \beta d \rfloor \rceil, 1, 1) = 2 > 1 > 3^{-c} \geq (2/\beta - 1)^{-c}$. If $\lceil d/\lfloor \beta d \rfloor \rceil > 2$, then

$$\begin{aligned}
\phi(\lceil d/\lfloor \beta d \rfloor \rceil, 1, 1) &= (\lceil d/\lfloor \beta d \rfloor \rceil - 2)^{-c} \\
&> (2/\beta - 1)^{-c}
\end{aligned}$$

by Fact 7.12. Hence, in either case, $\phi(\lceil d/\lfloor \beta d \rfloor \rceil, 1, 1) > (2/\beta - 1)^{-c}$. Thus

$$\begin{aligned}
\phi(d, \lfloor \beta d \rfloor, \lfloor \beta d \rfloor) &= \lfloor \beta d \rfloor \phi(\lceil d/\lfloor \beta d \rfloor \rceil, 1, 1) \\
&> (\beta d - 1)(2/\beta - 1)^{-c} \qquad (7.5)
\end{aligned}$$

Using the bound on d and the definitions of c_1 and c_2, we get

$$\begin{aligned}
\beta d - 1 &\geq c_1\beta^{-c}(n + c_2) - 1 \\
&= c_1\beta^{-c}\left(n + \frac{W-2}{2} + \frac{2^{-c}}{c_1}\right) - 1 \\
&\geq c_1\beta^{-c}\left(n + \frac{W-2}{2}\right) \\
&= (2/\beta)^c(W + 2(n-1)) \qquad (7.6)
\end{aligned}$$

Combining (7.5) and (7.6) yields

$$\begin{aligned}
\phi(d, \lfloor\beta d\rfloor, \lfloor\beta d\rfloor) &> \left(\frac{2/\beta}{2/\beta - 1}\right)^c (W + 2(n-1)) \\
&> W + 2(n-1)
\end{aligned}$$

Hence, by Theorem 7.9, the transformation system us a card game ξ-system for n-bit perfect secret key exchange for against the extra cards view. ■

It was claimed in [FPR91] that secret bit transmission is possible for two players each holding a fraction β of the cards, but the minimum deck size for their solution to work is super-polynomial in $1/\beta$. From Theorem 7.13 with $n = 1$, it follows that the transformation system can be used to solve this problem with a minimum deck size that is only $O((1/\beta)^{(c+1)})$.

The following corollary is a consequence of Theorem 7.13, Theorem 5.4, and Lemma 3.4.

Corollary 7.14 *Let* $0 < \beta \leq 1/2$, $n \geq 1$, *and suppose that* $d \geq c_1(1/\beta)^{(c+1)}(n + c_2)$ *and* $s_i \geq \lfloor\beta d\rfloor$ *for* $i \in \{1, 2\}$. *Then* $\mathrm{pcap}(s_1, s_2; d) \geq 2^n$.

7.3.2 Multiparty Secret Key Exchange

We reduce the problem of multiparty n-bit perfect secret key exchange to the problem of 2-party n-bit perfect secret key exchange by showing how to use an $(a, b; d)$-system $\mathcal{P}_\mathcal{S}$ to construct a $(s_1, \ldots, s_k; d)$-system $\mathcal{P}^*_{\mathcal{S}^*}$ for certain a, b, s_1, …, s_k, and d. The construction has the property that if $\mathcal{P}_\mathcal{S}$ is a card game $(a, b; d)$-system for n-bit perfect secret key exchange against the extra cards view, then $\mathcal{P}^*_{\mathcal{S}^*}$ is a card game $(s_1, \ldots, s_k; d)$-system for n-bit perfect secret key exchange against the extra cards view. Applying this construction to the 2-player transformation protocol yields an efficient multiparty n-bit perfect secret key exchange protocol.

The construction is a composition of $k-1$ simulations of $\mathcal{P}_{\mathcal{S}}$. Each team player P_i for $1 \le i \le k-1$ randomly chooses a subset h_i^a containing a of her cards. Each team player P_i for $2 \le i \le k$ randomly chooses a subset h_i^b containing b of her cards not in h_i^a. In the i^{th} simulation of $\mathcal{P}_{\mathcal{S}}$, neighbors P_i and P_{i+1} become the *active* players and simulate $\mathcal{P}_{\mathcal{S}}$ to establish an n-bit perfect secret key y_i that they share. Player P_i uses h_i^a as her hand to play the role of Alice in $\mathcal{P}_{\mathcal{S}}$. Player P_{i+1} uses h_{i+1}^b as her hand to play the role of Bob in $\mathcal{P}_{\mathcal{S}}$. The other players send only the message "pass".

After the $k-1$ simulations of $\mathcal{P}_{\mathcal{S}}$ are completed, player P_1 becomes the leader and randomly chooses an n-bit string v to be the team's secret key. When P_i learns v, she sends $z_i = v \oplus y_i$ to P_{i+1} publicly. P_{i+1} recovers v by computing $z_i \oplus y_i$. This is a simple instance of the idea, presented in Chapter 6, of using a spanning tree (here a chain) of 2-player secret keys to communicate a single secret key.

(Note that in order to decrease the total number of rounds of $\mathcal{P}^*$, it is possible to run each of the simulations of $\mathcal{P}_{\mathcal{S}}$ in parallel since they do not affect each other. In this case, $\mathcal{P}^*$ has $t^{\mathcal{P}} + k - 1$ rounds.)

Theorem 7.15 *Let $n \ge 1$ and $k \ge 2$, and let $\xi = (a, b; d)$ and $\xi^* = (s_1, \ldots, s_k; d)$ such that $s_1 \ge a$, $s_k \ge b$, and $s_i \ge a + b$ for all $2 \le i \le k-1$, and let $\mathcal{P}_{\mathcal{S}}$ be a card game ξ-system for n-bit perfect secret key exchange against the extra cards view. Then ${\mathcal{P}^*}_{\mathcal{S}^*}$ as constructed above is a card game ξ^*-system for n-bit perfect secret key exchange against the extra cards view.*

Proof (sketch): Assume the conditions of the theorem. Since $\mathcal{P}_{\mathcal{S}}$ achieves perfect n-bit secret key exchange against the extra cards view, each simulation of $\mathcal{P}_{\mathcal{S}}$ results in an n-bit perfect secret key for the two participants. Since each simulation uses different cards, later simulations do not affect the secrecy of earlier simulations. Hence, after all the simulations, the results from the simulations as one-time pads can successfully be used to transmit the randomly chosen n-bit string v secretly to all the team players. ∎

We can apply Theorem 7.15 to the transformation protocol to obtain an n-bit, k-player perfect secret key exchange system that requires the deck size to be only linear in n and polynomial in $1/\alpha$, where α is the fraction of the deck held by each team player. Recall that $c_1 = 2^{(c+1)}$ and $c_2 = (W-2)/2 + 2^{-c}/c_1$.

Corollary 7.16 *Let $k \geq 2$, $0 < \alpha \leq 1/k$, and suppose $d \geq (c_1)^2(1/\alpha)^{(c+1)}(n + c_2)$. Then ${\mathcal{P}^*}_{\mathcal{S}^*}$ is a card game $(\lfloor \alpha d \rfloor^{(k)}; d)$-system for perfect n-bit secret key exchange against the extra cards view.*

Proof: Suppose the conditions of the corollary, and let $\beta = \alpha/2$. Then since $(c_1)^2(1/\alpha)^{(c+1)} = c_1(2/\alpha)^{(c+1)}$, it follows from Theorem 7.13 that the transformation system is a card game $(\lfloor \beta d \rfloor, \lfloor \beta d \rfloor; d)$-system for n-bit perfect secret key exchange against the extra cards view. Since $2\lfloor \beta d \rfloor \leq \lfloor \alpha d \rfloor$, the conditions of Theorem 7.15 are satisfied, and hence ${\mathcal{P}^*}_{\mathcal{S}^*}$ is a card game $(\lfloor \alpha d \rfloor^{(k)}; d)$-system for n-bit perfect secret key exchange against the extra cards view. ■

Corollary 7.17 *Let $k \geq 2$, $0 < \alpha \leq 1/k$, and suppose $d \geq (c_1)^2(1/\alpha)^{(c+1)}(n + c_2)$ and $s_i \geq \lfloor \beta d \rfloor$ for $i \in K$. Then $\mathrm{pcap}(s_1, \ldots, s_k; d) \geq 2^n$.*

Corollary 7.18 *Assume $m \geq 2$ divides d, and let each of m players be dealt hands of size d/m from a deck of size $d \geq (c_1)^2 m^{(c+1)}(n + c_2)$. Then for any team of size $2 \leq k \leq m$ that forms, ${\mathcal{P}^*}_{\mathcal{S}^*}$ performs n-bit perfect secret key against the extra cards view.*

Taking $d = (c_1)^2 m^{(c+1)}(n + c_2)$, Corollary 7.18 says that for deals of cards, n-bit perfect secret key exchange is possible when each of m players is given $(c_1)^2 m^c(n + c_2)$ cards. Representing each card by $\log d$ bits, each hand requires $(c_1)^2 m^c(n + c_2)\log((c_1)^2 m^{(c+1)}(n + c_2)) = O(m^c n \log(mn))$ bits. The solution in which each of the m players is initially given an n-bit one-time pad for each of the other $m - 1$ players requires each player to be given only $(m - 1)n = O(mn)$ bits, but the initial data in this case does not enjoy the extreme simplicity of a card deal.

CHAPTER 8: THE CAPACITY OF $(1^{(k)}; k)$ IS ONE

It follows from the work of Fischer, Paterson and Rackoff [FPR91] that for teams of size two, 2-valued perfect secret key exchange is always possible when the team holds all the cards, provided that each player has at least one card. However, for larger teams, this is not the case. In particular, we show that secret key exchange is not possible when each of $k \geq 3$ team players holds one card from a k card deck. By Proposition 4.1 and Theorem 4.5, $1 \leq \text{pcap}(1^{(k)}; k) \leq \text{wcap}(1^{(k)}; k) \leq k$. By examining the structure of the set of legal $(1^{(k)}; k)$-deals, it is possible to show that $\text{pcap}(1^{(k)}; k) = \text{wcap}(1^{(k)}; k) = 1$ if $k \geq 3$. Fischer and Wright [FW92] showed this result for the case $k = 3$. The proof given here generalizes the proof in [FW92]. An alternate proof appears in [BHW93].

Let $k \geq 3$ and let $\xi = (1^{(k)}; k)$. Since there are k cards in ξ, $\Delta^{\xi} = K$. Let $\mathcal{P}_{\mathcal{S}}$ be a card game ξ-system and let L denote the set of legal ξ-deals. Since $\mathcal{P}_{\mathcal{S}}$ is a card game system, the set of feasible deals is L. (We assume throughout this section that the set of feasible deals is L.) We will denote by j the hand containing the single card j. A legal deal can be regarded as a permutation of K (and every permutation of K corresponds to a legal deal). Thus, if δ is legal, then for every $j \in K$, there is a unique $i \in K$ such that $\langle \delta \rangle_i = j$.

Lemma 8.1 *Suppose $\{\alpha, \beta\} \subseteq L$ is not coverable, and let $x, j \in K$. Then there exists $\delta \in L$ such that $\langle \delta \rangle_x = j$ and for all $i \neq x$, $\delta =_i \alpha$ or $\delta =_i \beta$.*

Proof: Suppose the conditions of the lemma. We will construct the desired deal δ. Figures 8.1–8.3 show the result of steps of the construction for the example $\alpha = (2, 1, 5, 4, 6, 3)$, $\beta = (5, 3, 6, 2, 1, 4)$, $x = 5$, and $j = 4$.

We begin by constructing a directed graph $G = (V, E)$ that represents the deals α and β, where $V = K$ and $E = \{(a, b) : \langle \alpha \rangle_a = \langle \beta \rangle_b\}$. We label the edge $(a, b) \in E$ by $\mathcal{L}(a, b) = \langle \alpha \rangle_a = \langle \beta \rangle_b$. Hence the vertices of G are coordinates and the edges of G are labeled by cards. (See Figure 8.1.)

Since α is a permutation, each vertex i has exactly one incoming edge, labeled $\langle \beta \rangle_i$, which we denote by incoming(i). Symmetrically, since β is a permutation, each vertex i has exactly one outgoing edge, labeled $\langle \alpha \rangle_i$, which we denote by outgoing(i). Hence G is a collection of disjoint cycles. Furthermore, since $\mathcal{L}(\text{outgoing}(i)) = \langle \alpha \rangle_i$ and α is a permutation, $\mathcal{L}(\text{outgoing}(V)) = \{\langle \alpha \rangle_i : i \in V\} = K$. Symmetrically, $\mathcal{L}(\text{incoming}(V)) = K$.

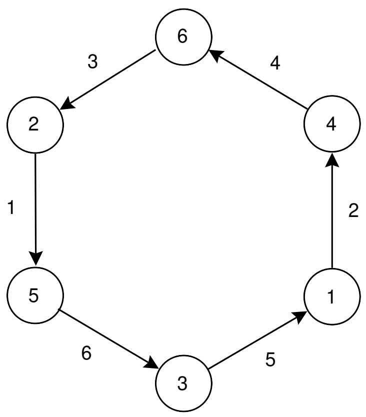

Figure 8.1: The graph G for $\alpha = (2,1,5,4,6,3)$ and $\beta = (5,3,6,2,1,4)$

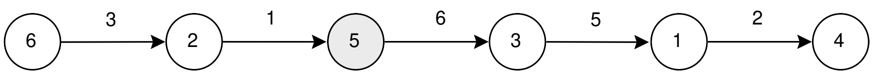

Figure 8.2: The graph G' for $x = 5$ and $j = 4$

We will now show that in fact G consists of exactly one cycle. Suppose not. Then G is the union of two disjoint graphs $G_1 = (V_1, E_1)$ and $G_2 = (V_2, E_2)$, where G_1 and G_2 are nonempty collections of disjoint cycles. Hence $\text{incoming}(V_1) = \text{outgoing}(V_1) = E_1$ and $\text{incoming}(V_2) = \text{outgoing}(V_2) = E_2$. Consider the deal γ defined by

$$\langle\gamma\rangle_i = \begin{cases} \mathcal{L}(\text{incoming}(i)) & \text{if } i \in V_1 \\ \mathcal{L}(\text{outgoing}(i)) & \text{if } i \in V_2 \end{cases}$$

The deal γ is a permutation (and hence feasible) because $\{\langle\gamma\rangle_i : i \in K\} = \mathcal{L}(E_1) \cup \mathcal{L}(E_2) = \mathcal{L}(E) = K$. Since $\mathcal{L}(\text{incoming}(i)) = \langle\beta\rangle_i$ and $\mathcal{L}(\text{outgoing}(i)) = \langle\alpha\rangle_i$, γ interpolates $\{\alpha, \beta\}$. Since V_1 and V_2 are nonempty, γ spans $\{\alpha, \beta\}$. It follows that γ covers $\{\alpha, \beta\}$, a contradiction. We conclude that G consists of a single cycle.

Let $G' = (V', E')$ be the directed graph obtained from G by removing the edge labeled j. Then G' consists of a single chain. (See Figure 8.2.)

Let $i, i' \in V'$. We write $i \xrightarrow{+} i'$ if there is a nonempty path in G' from i to i'. We define $\text{left}(i) = \{i' \in V' : i' \xrightarrow{+} i\}$ and $\text{right}(i) = \{i' \in V' : i \xrightarrow{+} i'\}$. Note that $V' = \{x\} \cup \text{left}(x) \cup \text{right}(x)$ and $E' = \text{outgoing}(\text{left}(x)) \cup \text{incoming}(\text{right}(x))$. We

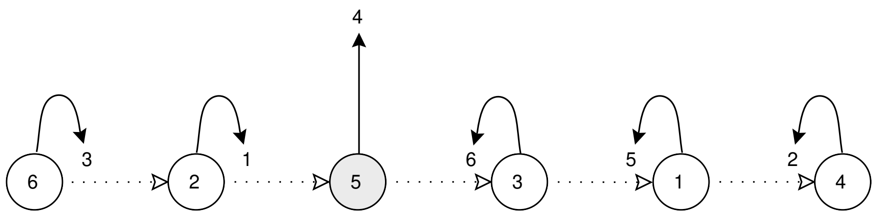

Figure 8.3: Use of G' to define the deal $\delta = (5, 1, 6, 2, 4, 3)$

construct the desired deal δ as follows.

$$\langle\delta\rangle_i = \begin{cases} j & \text{if } i = x \\ \mathcal{L}(\text{outgoing}(i)) & \text{if } i \in \text{left}(x) \\ \mathcal{L}(\text{incoming}(i)) & \text{if } i \in \text{right}(x) \end{cases}$$

(See Figure 8.3.) Then δ is a permutation because $\{\langle\delta\rangle_i : i \in K\} = \{j\} \cup \mathcal{L}(E') = K$. Hence $\delta \in L$. Clearly $\langle\delta\rangle_x = j$. Since $\mathcal{L}(\text{outgoing}(i)) = \langle\alpha\rangle_i$ and $\mathcal{L}(\text{incoming}(i)) = \langle\beta\rangle_i$, $\delta =_i \alpha$ or $\delta =_i \beta$ for all $i \neq x$. ■

Lemma 8.2 *Let $\alpha, \beta, \delta \in L$ and $x, y \in K$ such that $\langle\alpha\rangle_x = \langle\beta\rangle_y$ and for every $i \in K - \{x, y\}$, $\delta =_i \alpha$ or $\delta =_i \beta$. Then $\delta =_x \alpha$ or $\delta =_y \beta$.*

Proof: Suppose the conditions of the lemma and let $j = \langle\alpha\rangle_x = \langle\beta\rangle_y$. Let $i \in K - \{x, y\}$. Then $j \notin \{\langle\alpha\rangle_i, \langle\beta\rangle_i\}$. Hence since $\delta =_i \alpha$ or $\delta =_i \beta$, it follows that $\langle\delta\rangle_i \neq j$. Hence $\langle\delta\rangle_i \neq j$ for every $i \in K - \{x, y\}$. Since δ is a permutation, it follows that either $\langle\delta\rangle_x = j$ or $\langle\delta\rangle_y = j$, so $\delta =_x \alpha$ or $\delta =_y \beta$. ■

Together, Lemmas 8.1, 8.2, and 4.6 yield the following.

Lemma 8.3 *Let $\alpha, \beta, \gamma \in L$ such that $\{\alpha, \beta\}$ is not coverable and let $x \in K$. Then there exists a deal δ for which the following simultaneously hold.*

1. *δ is a feasible interpolant of $\{\alpha, \beta, \gamma\}$, and $\delta =_x \gamma$.*

2. *Let τ be a partial conversation compatible with α, β, and γ. Let m be a message vector such that α and β are both compatible with τm. Let m' be a message vector such that γ is compatible with $\tau m'$. Let the message vector $\widehat{m}$ be defined by $\langle\widehat{m}\rangle_x = \langle m'\rangle_x$ and $\langle\widehat{m}\rangle_i = \langle m\rangle_i$ for $i \neq x$. Then δ is compatible with $\tau\widehat{m}$.*

3. *Let y and z be such that $\langle\beta\rangle_y = \langle\alpha\rangle_x$ and $\langle\alpha\rangle_z = \langle\beta\rangle_x$. If $\gamma \neq_x \alpha$, then $\delta =_y \beta$. If $\gamma \neq_x \beta$, then $\delta =_z \alpha$.*

4. If $\gamma \neq_x \alpha$ and $\gamma \neq_x \beta$, then δ covers $\{\alpha, \beta, \gamma\}$, and hence $\{\alpha, \beta, \gamma\}$ is coverable.

Proof: Let $\alpha, \beta, \gamma \in L$ such that $\{\alpha, \beta\}$ is not coverable and let $x \in K$. Applying Lemma 8.1 to α, β, x, and $\langle\gamma\rangle_x$, we obtain a legal deal δ such that $\delta =_x \gamma$, and for all $i \neq x$, $\delta =_i \alpha$ or $\delta =_i \beta$.

1. Immediate by choice of δ.

2. Let τ, m, m' satisfy the conditions of part 2. We define deals $\delta_1, \ldots, \delta_k$ as follows: $\delta_x = \gamma$, and for $i \neq x$, $\delta_i = \alpha$ if $\delta =_i \alpha$, and $\delta_i = \beta$ if $\delta =_i \beta$. Then $(\delta_1, \ldots, \delta_k)$ is an $\{\alpha, \beta, \gamma\}$-derivation of δ. We similarly define message vectors $m_1, \ldots, m_k$ such that $m_x = m'$, and for $i \neq x$, $m_i = m$. This construction ensures that for all $i \in K$, δ_i is compatible with τm_i and $\widehat{m} = (\langle m_1\rangle_1, \langle m_2\rangle_2, \ldots, \langle m_k\rangle_k)$. By Lemma 4.6 (part 2), $\widehat{\delta}$ is compatible with $\widehat{m}$.

3. Let y and z be such that $\langle\beta\rangle_y = \langle\alpha\rangle_x$ and $\langle\alpha\rangle_z = \langle\beta\rangle_x$. (This is possible since α and β are permutations.) Suppose $\gamma \neq_x \alpha$. Since $\delta =_x \gamma$, we have $\delta \neq_x \alpha$. By Lemma 8.2 applied to α, β, δ, x, and y, we have $\delta =_x \alpha$ or $\delta =_y \beta$. Hence, $\delta =_y \beta$. Similarly, suppose $\gamma \neq_x \beta$. Since $\delta =_x \gamma$, we have $\delta \neq_x \beta$. By Lemma 8.2 applied to β, α, δ, x, and z, we have $\delta =_x \beta$ or $\delta =_z \alpha$. Hence, $\delta =_z \alpha$.

4. If $\gamma \neq_x \alpha$ and $\gamma \neq_x \beta$, then by part 3, $\delta =_y \beta$ and $\delta =_z \alpha$. Since $\delta =_x \gamma$, it follows that δ spans $\{\alpha, \beta, \gamma\}$. Since δ is a feasible interpolant of $\{\alpha, \beta, \gamma\}$, it follows that δ covers $\{\alpha, \beta, \gamma\}$, and hence $\{\alpha, \beta, \gamma\}$ is coverable. ■

By definition, a conversation is feasible if it is compatible with some legal deal. Lemma 8.4 shows that in fact each feasible complete conversation must be compatible with exactly two legal deals.

Lemma 8.4 *Let $k \geq 3$ and $N \geq 2$. Let $\mathcal{P}_S$ be a card game $(1^{(k)}; k)$-system for N-valued weak secret key exchange and let $\sigma \in$ cc be feasible. Then* compat(σ) *is not coverable and* $|\text{compat}(\sigma)| = 2$.

Proof: Suppose the conditions of the lemma. Let $v_0, v_1 \in V$ such that $v_0 \neq v_1$ and let $j \in \{0, 1\}$. By weak uniformity, $\Pr[\mathrm{O}_1(v_j)] > 0$. Since σ is feasible, $\Pr[\mathrm{C}(\sigma)] > 0$. Hence by Lemma 3.11 (part 4) and Proposition 3.3, $\Pr[\mathrm{CO}(\sigma, v_j)] =$

$\Pr[\mathrm{C}(\sigma) \cap \mathrm{O}_1(v_j)] > 0$. Thus $\mathrm{CO}(\sigma, v_j)$ contains a feasible point ω_j. Let $u_j = \theta(\omega_j)$. Then $\mathrm{conv}(u_j) = \sigma$ and $\mathrm{out}_i(u_j) = v_j$ for all $i \in K$. Let $\alpha = \mathrm{corr}(u_0)$, and let $\beta = \mathrm{corr}(u_1)$. Hence $u_0 \in \mathrm{F}(\alpha, \sigma)$ and $u_1 \in \mathrm{F}(\beta, \sigma)$. It follows that $\mathrm{compat}(\sigma) \supseteq \{\alpha, \beta\}$. By Lemma 4.3, $\mathrm{out}(\alpha, \sigma)$ and $\mathrm{out}(\beta, \sigma)$ are defined, so $\mathrm{out}(\alpha, \sigma) = v_0$ and $\mathrm{out}(\beta, \sigma) = v_1$. Since $\mathrm{out}(\alpha, \sigma) \neq \mathrm{out}(\beta, \sigma)$, it follows from Lemma 4.7 that Q is not coverable for any set Q such that $\{\alpha, \beta\} \subseteq Q \subseteq \mathrm{compat}(\sigma)$.

We complete the proof by showing that $\mathrm{compat}(\sigma) = \{\alpha, \beta\}$. Suppose by way of contradiction that $\gamma \in \mathrm{compat}(\sigma) - \{\alpha, \beta\}$. Since $\{\alpha, \beta\} \subseteq \{\alpha, \beta, \gamma\} \subseteq \mathrm{compat}(\sigma)$, the set $\{\alpha, \beta, \gamma\}$ is not coverable. The deal γ is feasible since $\gamma \in \mathrm{compat}(\sigma)$. By Proposition 2.4, there is a coordinate x such that $\gamma \neq_x \alpha$ and $\gamma \neq_x \beta$. By Lemma 8.3 (part 4), $\{\alpha, \beta, \gamma\}$ is coverable, a contradiction. We conclude that $\mathrm{compat}(\sigma) = \{\alpha, \beta\}$. ■

Theorem 8.5 *Let* $k \geq 3$. *Then* $\mathrm{pcap}(1^{(k)}; k) = \mathrm{wcap}(1^{(k)}; k) = 1$.

Proof: Let $k \geq 3$. Since $\mathrm{wcap}(\xi) \geq \mathrm{pcap}(\xi) \geq 1$ for any signature ξ, we need only show that $\mathrm{wcap}(1^{(k)}; k) < 2$. Suppose by way of contradiction that $\mathcal{P}_{\mathcal{S}}$ is a card game $(1^{(k)}; k)$-system for 2-valued weak secret key exchange. We construct a tree whose nodes are the feasible conversations of $\mathcal{P}_{\mathcal{S}}$. Two nodes τ and σ are connected by an edge if τ is a prefix of σ and $|\tau| + 1 = |\sigma|$. Thus, the internal nodes are partial conversations and the leaves are complete conversations. We often identify a node σ with the unique path from the root to σ. If τ is a prefix of σ, we say σ *passes through* τ.

If a deal δ is compatible with a conversation τ, then δ is compatible with every prefix of τ. Also, if τ is not complete, then δ is compatible with at least one extension of τ. Hence, δ is compatible with the parent of τ (provided τ is not empty), and δ is compatible with at least one child of τ (provided τ is not complete). It follows that if δ is compatible with σ, then δ is compatible with every node on the path to σ. Hence, every node on the path to σ is compatible with every deal $\delta \in \mathrm{compat}(\sigma)$. If σ is a leaf, then by Lemma 8.4, $\mathrm{compat}(\sigma)$ is not coverable and $|\mathrm{compat}(\sigma)| = 2$.

Every legal deal is compatible with the empty conversation and there are more than two legal $(1^{(k)}; k)$-deals. Hence, more than two deals are compatible with the root. Since there are exactly two deals compatible with every leaf and the tree is finite, there must be some node τ compatible with more than two deals, each of whose children is compatible with exactly two deals.

Since τ is compatible with more than two distinct deals and every deal compatible with τ is compatible with at least one of its children, τ must have two children τm and $\tau m'$ such that $\text{compat}(\tau m) \neq \text{compat}(\tau m')$. Let $\{\alpha, \beta\} = \text{compat}(\tau m)$ and let $\{\alpha', \beta'\} = \text{compat}(\tau m')$. By the above, $\alpha \neq \beta$ and $\{\alpha, \beta\}$ is not coverable. Similarly, $\alpha' \neq \beta'$ and $\{\alpha', \beta'\}$ is not coverable. Since $\{\alpha, \beta\} \neq \{\alpha', \beta'\}$, either $\alpha' \notin \{\alpha, \beta\}$ or $\beta' \notin \{\alpha, \beta\}$. We assume without loss of generality that $\alpha' \notin \{\alpha, \beta\}$. (Figure 8.4 illustrates this construction.)

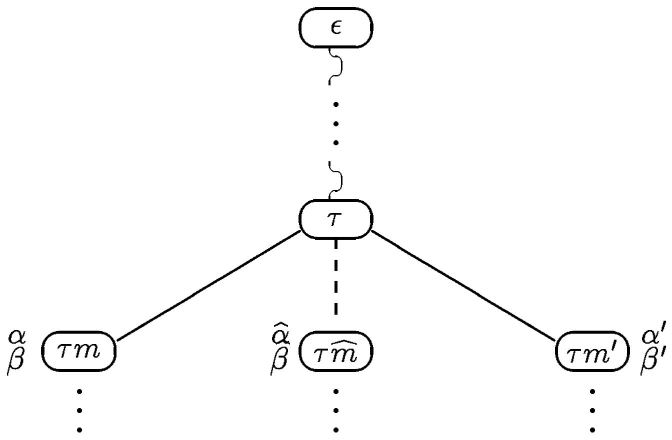

Figure 8.4: Tree of feasible conversations

By Proposition 2.4, there is a coordinate x such that $\alpha' \neq_x \alpha$ and $\alpha' \neq_x \beta$. Let y and z be such that $\langle\beta\rangle_y = \langle\alpha\rangle_x$ and $\langle\alpha\rangle_z = \langle\beta\rangle_x$. Let the message vector $\widehat{m}$ be defined by $\langle\widehat{m}\rangle_x = \langle m'\rangle_x$ and $\langle\widehat{m}\rangle_i = \langle m\rangle_i$ for $i \neq x$.

By Lemma 8.3 (parts 1–3) applied to α, β, α', and x, there is a deal $\widehat{\alpha}$ such that (by part 1) $\widehat{\alpha}$ is a feasible interpolant of $\{\alpha, \beta, \alpha'\}$,

$$\widehat{\alpha} =_x \alpha' \tag{8.1}$$

and (by part 2) $\widehat{\alpha}$ is compatible with $\tau\widehat{m}$. Since $\alpha' \neq_x \alpha$ and $\alpha' \neq_x \beta$, it follows (by part 3) that

$$\widehat{\alpha} =_y \beta \tag{8.2}$$

$$\widehat{\alpha} =_z \alpha \tag{8.3}$$

Similarly, by Lemma 8.3 (parts 1–3) applied to α, β, β', and x, there is a deal $\widehat{\beta}$ such that (by part 1) $\widehat{\beta}$ is a feasible interpolant of $\{\alpha, \beta, \beta'\}$,

$$\widehat{\beta} =_x \beta' \tag{8.4}$$

and (by part 2) $\widehat{\beta}$ is compatible with $\tau\widehat{m}$. By Proposition 2.3, since $\{\alpha, \beta\}$ is not coverable, α and β do not touch. In particular, $\alpha \neq_x \beta$. Hence $\widehat{\beta} \neq_x \alpha$ or $\widehat{\beta} \neq_x \beta$. It follows (by part 3) that

$$\widehat{\beta} =_y \beta \text{ or } \widehat{\beta} =_z \alpha \tag{8.5}$$

Since $\{\alpha', \beta'\}$ is not coverable, it follows from Proposition 2.3 that α' and β' do not touch. Thus, in particular, $\alpha' \neq_x \beta'$. Hence it follows from (8.1) and (8.4) that $\widehat{\alpha} \neq_x \widehat{\beta}$. Therefore $\widehat{\alpha}$ and $\widehat{\beta}$ are distinct. By (8.2), (8.3), and (8.5), $\widehat{\alpha} =_y \widehat{\beta}$ or $\widehat{\alpha} =_z \widehat{\beta}$. Thus $\widehat{\alpha}$ touches $\widehat{\beta}$ at y or z. By Proposition 2.3, $\{\widehat{\alpha}, \widehat{\beta}\}$ is coverable.

Since $\widehat{\alpha}$ and $\widehat{\beta}$ are compatible with $\tau\widehat{m}$, it follows that $\tau\widehat{m}$ is a feasible conversation and $\{\widehat{\alpha}, \widehat{\beta}\} \subseteq \text{compat}(\tau\widehat{m})$. Then $\tau\widehat{m}$ is a child of τ in the tree of feasible conversations, so $\text{compat}(\tau\widehat{m})$ is not coverable and $|\text{compat}(\tau\widehat{m})| = 2$. Since $\{\widehat{\alpha}, \widehat{\beta}\} \subseteq \text{compat}(\tau\widehat{m})$ and $\widehat{\alpha} \neq \widehat{\beta}$, it follows that $\text{compat}(\tau\widehat{m}) = \{\widehat{\alpha}, \widehat{\beta}\}$, a contradiction since $\{\widehat{\alpha}, \widehat{\beta}\}$ is coverable. ■

CHAPTER 9: CONCLUSIONS

A model is presented for multiparty communication among players receiving correlated inputs. Our model makes it possible to reason formally about intuitive concepts based on shared knowledge in a multiparty setting. We define the secret key exchange problem in this model and present several alternative characterizations of secret key exchange and the secret key exchange conditions.

We show that independent randomization does not help Eve to thwart secret key exchange, and that independent randomization alone does not allow a team to perform secret key exchange. In contrast, Example 1 (§5.3) exhibits a source $\mathcal{S}_1$ of correlated inputs that does not allow secret key exchange to be performed, while $\mathcal{S}_1$ augmented by independent randomization does allow secret key exchange to be performed.

We define the perfect capacity and weak capacity of a source and show that the perfect and weak capacities of every source are finite. In particular, for any source $\mathcal{T}$, we show that $\mathrm{pcap}(\mathcal{T}) \leq \mathrm{wcap}(\mathcal{T}) \leq \min_{i \in K} \left|\mathrm{feas}\left(U_i^{\mathcal{T}}\right)\right|$ and $\mathrm{pcap}(\mathcal{T}) \leq \left\lfloor \sqrt[k-1]{\psi_{\mathcal{T}}} \right\rfloor$. These bounds hold for all view functions for Eve. We do not know how to use the additional information given to Eve to improve these results for any particular view function.

We explore the use of a random deal of cards for secret key exchange and show that it is possible to use a random deal of cards to achieve perfect secrecy. We exhibit several card game systems for perfect secret key exchange against the extra cards view. These systems also provide lower bounds on the capacity of card games (against the empty view). In Chapter 6, we exhibit the class of key set protocols. We use a game-theoretic argument to exhibit an optimal key set protocol. Using key set protocols, it follows that $\mathrm{pcap}(s_1, \ldots, s_k; d) \geq 2$ if $\max s_i + \min s_i \geq k + d - \sum s_i$. In Chapter 7, we exhibit the transformation protocol and use a potential argument to analyze the protocol. Using the transformation protocol and protocol composition, we show that if $d \geq (42.8) \cdot (1/\beta)^{2.71}(n + 0.07)$ and $s_i \geq \lfloor \beta d \rfloor$ for $i \in K$, then $\mathrm{pcap}(s_1, \ldots, s_k; d) \geq 2^n$. Finally, we show that $\mathrm{pcap}(1^{(k)}; k) = \mathrm{wcap}(1^{(k)}; k) = 1$ if $k \geq 3$.

Example 1 (§5.3) exhibits a signature for which deterministic secret bit exchange against the empty view is possible but deterministic secret bit exchange against the extra cards view is not possible. This shows that allowing Eve to see the remaining cards can sometimes affect the possibility of secret key exchange.

Example 2 (§4.2) exhibits a source $\mathcal{T}$ such that the perfect capacity of $\mathcal{T}$ is strictly less than the weak capacity of $\mathcal{T}$. It is open whether there exist *signatures* ξ such that the perfect capacity of ξ is strictly less than the weak capacity of ξ. We conjecture that such a signature does not exist because of the symmetry inherent in the structure of a deal.

Except in some simple cases of ξ, the upper and lower bounds on the capacity of ξ given by results in this dissertation do not coincide. It remains open to determine exact characterizations of the weak and perfect capacities of an arbitrary signature ξ.

Our secret key exchange requirements are quite strict in that they require every run of a system to succeed in obtaining a secret key. In the case of perfect secret key exchange, this key is required to be perfectly uniform and perfectly secret. An obvious direction for future work is to relax the strictness of these requirements. One such relaxation is to allow a small fraction of the runs to fail in some way, such as the players' outputs not agreeing or Eve learning the outputs. Orthogonally, a small margin of error could be allowed everywhere. Our weak secrecy is an example of such a relaxation. A less drastic example is to require Eve's probability of guessing the output correctly to be at most $1/N + \epsilon$ for some constant ϵ. It is our hope that the work in this dissertation will provide a solid foundation from which to further investigate these issues.

BIBLIOGRAPHY

[Abe78] H. Abelson. Lower bounds on information transfer in distributed computations. In *Proc. 19th IEEE Symposium on the Foundations of Computer Science*, pages 151–158, October 1978.

[BB85] C. H. Bennett and G. Brassard. An update on quantum cryptography. In *Advances in Cryptology: Proceedings of CRYPTO '84*, volume 196 of *LNCS*, pages 475–480. Springer-Verlag, 1985.

[BBE92] C. H. Bennett, G. Brassard, and A. K. Ekert. Quantum cryptography. *Scientific American*, pages 50–57, October 1992.

[BBR88] C. H. Bennett, G. Brassard, and J. Robert. Privacy amplification by public discussion. *SIAM J. Computing*, 17(2):210–229, April 1988.

[BCG82] E. R. Berlekamp, J. H. Conway, and R. K. Guy. *Winning Ways*, volume I. Academic Press, Inc., London, 1982.

[BHW93] D. Beaver, S. Haber, and P. Winkler. On the isolation of a common secret. Preprint, Bellcore, 1993.

[BOGKW88] M. Ben-Or, S. Goldwasser, J. Kilian, and A. Wigderson. Multi-prover interactive proofs: How to remove intractability assumptions. In *Proc. 20th ACM Symposium on Theory of Computing*, pages 113–131, May 1988.

[CCD88] D. Chaum, C. Crépeau, and I. Damgård. Multiparty unconditionally secure protocols. In *Proc. 20th ACM Symposium on Theory of Computing*, pages 11–19, May 1988.

[DH76] W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Trans. Information Theory*, IT–22(6):644–654, November 1976.

[ElG85] T. ElGamal. A public key cryptosystem and a signature scheme based on discrete logarithms. *IEEE Trans. Information Theory*, IT–31(4):469–472, July 1985.

[Fli81] J. Flint. Cheating by degrees. *The Times Saturday Review*, May 9, 1981.

[FLP85] M. J. Fischer, N. A. Lynch, and M. S. Paterson. Impossibility of distributed consensus with one faulty process. *J. ACM*, 32(2):374–382, April 1985.

[FPR91] M. J. Fischer, M. S. Paterson, and C. Rackoff. Secret bit transmission using a random deal of cards. In *Distributed Computing and Cryptography*, volume 2 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 173–181. American Mathematical Society, 1991.

[FW] M. J. Fischer and R. N. Wright. Bounds on secret key exchange using a random deal of cards. *Journal of Cryptology*. To appear.

[FW92] M. J. Fischer and R. N. Wright. Multiparty secret key exchange using a random deal of cards. In *Advances in Cryptology – CRYPTO '91*, volume 576 of *LNCS*, pages 141–155. Springer-Verlag, 1992.

[FW93a] M. J. Fischer and R. N. Wright. An application of game theoretic techniques to cryptography. In *Complexity Theory*, volume 13 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, pages 99–118. American Mathematical Society, 1993.

[FW93b] M. J. Fischer and R. N. Wright. An efficient protocol for unconditionally secure secret key exchange. In *Proc. 4th Annual Symposium on Discrete Algorithms*, pages 475–483, January 1993.

[Gal68] R. Gallager. *Information Theory and Reliable Communication*. John Wiley & Sons, Inc., New York, 1968.

[GMR89] S. Goldwasser, S. Micali, and C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Computing*, 18(1):186–208, February 1989.

[GMW86] O. Goldreich, S. Micali, and A. Wigderson. Proofs that yield nothing but their validity and a methodology of cryptographic protocol design. In *Proc. 27th IEEE Symposium on Foundations of Computer Science*, pages 174–187, October 1986.

[IR89] R. Impagliazzo and S. Rudich. Limits on the provable consequences of one-way permutations. In *Proc. 21st ACM Symposium on Theory of Computing*, pages 44–61, May 1989.

[Lov90] L. Lovász. Communication complexity: A survey. In *Paths, Flows, and VLSI-Layout*, volume 9 of *Algorithms and Combinatorics*, pages 235–265. Springer-Verlag, 1990.

[LSP82] L. Lamport, R. Shostak, and M. Pease. The Byzantine generals problem. *ACM Trans. Programming Languages and Systems*, 4(3):382–401, July 1982.

[Mau93] U. M. Maurer. Secret key agreement by public discussion from common information. *IEEE Trans. Information Theory*, 39(3):733–742, May 1993.

[Mer78] R. C. Merkle. Secure communication over insecure channels. *Comm. ACM*, 21(4):294–299, April 1978.

[NOS93] M. Naor, A. Orlitsky, and P. Shor. Three results on interactive communication. *IEEE Trans. Information Theory*, 39(5):1608–1615, September 1993.

[Orl90] A. Orlitsky. Worst-case interactive communication I: Two messages are almost optimal. *IEEE Trans. Information Theory*, 36(5):1111–1126, September 1990.

[Orl91] A. Orlitsky. Worst-case interactive communication II: Two messages are not optimal. *IEEE Trans. Information Theory*, 37(4):995–1005, July 1991.

[RSA78] R. L. Rivest, A. Shamir, and L. Adleman. A method for obtaining digital signatures and public-key cryptosystems. *Comm. ACM*, 21(2):120–126, February 1978.

[Rud88] S. Rudich. *Limits on the Provable Consequences of One-way Functions.* PhD thesis, University of California at Berkeley, 1988.

[Sha48] C. E. Shannon. A mathematical theory of communication. *Bell System Technical Journal*, pages 379–423, 623–656, 1948.

[Sha49] C. E. Shannon. Communication theory of secrecy systems. *Bell System Technical Journal*, pages 656–715, 1949.

[Win81a] P. Winkler. Cryptologic techniques in bidding and defense: Parts I, II, III, and IV. *Bridge Magazine*, April–July 1981.

[Win81b] P. Winkler. My night at the Cryppie club. *Bridge Magazine*, pages 60–63, August 1981.

[Win83] P. Winkler. The advent of cryptology in the game of bridge. *Cryptologia*, 7(4):327–332, October 1983.

[Yao78] A. C.-C. Yao. Some complexity questions related to distributed computing. In *Proc. 11th ACM Symposium on Theory of Computing*, pages 209–213, May 1978.

[Yao82] A. C.-C. Yao. Protocols for secure computation. In *Proc. 23rd IEEE Symposium on the Foundations of Computer Science*, pages 160–164, October 1982.