

Secure Group Communication Using Robust Contributory Key Agreement

Yair Amir, *Member, IEEE Computer Society*, Yongdae Kim, Cristina Nita-Rotaru, *Member, IEEE*, John L. Schultz, Jonathan Stanton, *Member, IEEE Computer Society*, and Gene Tsudik, *Member, IEEE*

Abstract—Contributory group key agreement protocols generate group keys based on contributions of all group members. Particularly appropriate for relatively small collaborative peer groups, these protocols are resilient to many types of attacks. Unlike most group key distribution protocols, contributory group key agreement protocols offer strong security properties such as key independence and perfect forward secrecy. This paper presents the first robust contributory key agreement protocol resilient to any sequence of group changes. The protocol, based on the Group Diffie-Hellman contributory key agreement, uses the services of a group communication system supporting Virtual Synchrony semantics. We prove that it provides both Virtual Synchrony and the security properties of Group Diffie-Hellman, in the presence of any sequence of (potentially cascading) node failures, recoveries, network partitions, and heals. We implemented a secure group communication service, Secure Spread, based on our robust key agreement protocol and Spread group communication system. To illustrate its practicality, we compare the costs of establishing a secure group with the proposed protocol and a protocol based on centralized group key management, adapted to offer equivalent security properties.

Index Terms—Security and protection, fault tolerance, network protocols, distributed systems, group communication, contributory group key agreement, cryptographic protocols.

1 INTRODUCTION

MANY collaborative settings such as audio and videoconferencing, white-boards, clustering, and replication applications, require services which are not provided by the current network infrastructure. A typical collaborative application operates as a peer group where members communicate via reliable many-to-many multicast, sometimes requiring reliable ordered message delivery. In some settings, group members must be aware of the exact (agreed upon) group membership. Since group communication systems provide these services, many collaborative applications use group communication systems (GCS) as the underlying messaging infrastructure.

Security is crucial for distributed and collaborative applications that operate in a dynamic network environment and communicate over insecure networks such as the Internet. Basic security services needed in such a group setting are largely the same as in point-to-point communication: data secrecy and integrity, and entity authentication. These services cannot be attained without secure, efficient,

and robust group key management. Many critical applications (e.g., military and financial) require that all intragroup communication remain confidential. Consequently, not only sufficiently strong encryption must be used to protect intragroup messages, but the underlying group key management must also provide strong security guarantees.

Group keys can be viewed as a sequence of values sorted by time of use, with each key corresponding to a different “snapshot” of a group. A group key is changed whenever the group changes or a periodic rekey is needed. The strongest known security guarantees are *key independence* and *perfect forward secrecy* (PFS). Key independence states that a passive adversary—who, in the worst case, might know all group keys except one—cannot use its knowledge to discover the one key that is missing. PFS demands that the compromise of group members’ long-term keys should not lead to the compromise of any previously used group keys (see [3] for formal definitions).

Contributory group key agreement protocols that compute a group key as a (usually, one-way) function of individual contributions from all members can provide both key independence and PFS properties. At the same time, contributory group key agreement presents a tough practical challenge: Its multiround nature must be reconciled with the possibility of crashes, partitions, and other events affecting group membership that can occur *during* the execution of the group key agreement. Therefore, this paper focuses on **robust** contributory group key agreement.

1.1 Group Key Management

Traditional centralized key management relies on a single fixed key server to generate and distribute keys to the group. This approach is not well-suited for group communication systems that guarantee continuous operation in any possible group subset and any arbitrary number of partitions in the event of network partitions or faults. Although a key server can be made constantly available and

- Y. Amir is with the Department of Computer Science, 224 NEB, Johns Hopkins University, Baltimore, MD 21218. E-mail: yairamir@cs.jhu.edu.
- Y. Kim is with the Computer Science and Engineering Department, University of Minnesota, Twin Cities, 200 Union Street S.E., Minneapolis, MN 55455. E-mail: kyd@cs.umn.edu.
- C. Nita-Rotaru is with the Department of Computer Science, Purdue University, 250 N. University Street, West Lafayette, IN 47907-2066. E-mail: crism@cs.purdue.edu.
- J.L. Schultz, 335 S. Ann St., Baltimore, MD 21231. E-mail: jschultz@d-fusion.net.
- J. Stanton is with the Department of Computer Science, 801 22nd Street NW, Suite 704, George Washington University, Washington, DC 20052. E-mail: jstanton@gwu.edu.
- G. Tsudik is with the ICS Department, University of California, Irvine, 458 CS Building, Irvine, CA 92697-3425. E-mail: gts@ics.uci.edu.

Manuscript received 7 May 2002; revised 3 Mar. 2003; accepted 26 Sept. 2003.

For information on obtaining reprints of this article, please send e-mail to: tps@computer.org, and reference IEEECS Log Number 116498.

attack-resistant with the aid of various fault tolerance and replication techniques, it is very difficult (in a scalable and efficient manner) to make a centralized server present in every possible group subset. We note that centralized approaches work well in one-to-many multicast scenarios since a key server (or a set thereof) can support continued operation within an arbitrary partition as long as it includes the source.

The requirement to provide continued operation in an arbitrary partition can be overcome by dynamically selecting a group member to act as a group key server. However, most centralized key distribution protocols do not provide strong security properties such as key independence and PFS. These properties can only be provided if the key server maintains pairwise secure channels with each group member in order to distribute group keys. Although this approach seems appealing, each time a new key server comes into play, significant costs must be incurred to set up pairwise secure channels. In addition, this method has a disadvantage (common to all centralized fixed-server methods) in that it relies on a single entity to generate good (i.e., cryptographically strong) random keys.

Our approach is to use a fully distributed, contributory group key management algorithm where a group key is not selected by one entity, but instead, is a function of each group member's contribution. This avoids the issues with centralized trust, single point of failure (and attack), and the requirement to establish pairwise secret channels, and provides strong security properties such as forward and backward secrecy, key independence, and PFS [3].

1.2 Goal and Contribution

Secure, robust, and efficient key management is critical for secure group communication. However, designing key management protocols that are robust and efficient in the presence of network and process faults is a big challenge. The goal of this work is to provide a robust and secure group communication that offers Virtual Synchrony (VS) [4] semantics. Our contribution is three-fold:

1. We present the first robust contributory key agreement protocols that are resilient to any *finite* (even cascading) sequence of events. Our protocols (basic and optimized) are based on Group Diffie-Hellman (GDH) [5] key agreement.
2. We design a robust and secure group communication service by combining our robust key agreement with a reliable group communication service. We prove that the resulting system preserves the Virtual Synchrony properties as well as the security properties of GDH.
3. We provide an insight into the cost of adding security services to GCS, focusing on group key management costs. We describe the implementation of a secure group communication service—Secure Spread—based on our optimized robust key agreement protocol and the Spread [6] group communication system. We present experimental results measuring the delay incurred by a group installing a secure membership following group membership changes. The cost of establishing a secure group when our protocol is used is compared with the cost of establishing a secure group when a centralized key management protocol, modified such that it

provides the **same strong security properties** as our group key agreement, is used.

The rest of the paper is organized as follows: We present our failure and security models in Section 2. Section 3 presents both the group communication service and the key agreement protocol used in designing the robust secure group communication service. We then describe our protocols in Sections 4 and 5 and provide implementation details and performance results in Sections 6 and 7, respectively. Related work is overviewed in Section 8, and the paper concludes with a brief summary in Section 9.

2 FAILURE MODEL AND SECURITY ASSUMPTIONS

We consider a *distributed system* composed of a group of processes executing on one or more CPUs and coordinating their actions by exchanging messages. Message exchange is achieved via asynchronous multicast and unicast. While messages can be lost, we assume that message corruption is masked by a lower layer. Any process can crash and recover. A crash of any component of a process (i.e., key agreement layer or the group communication system) is considered a process crash. We assume that the crash of one of any component is detected by all the other components and is treated as a process crash.

Due to congestion or outright failures, the network can be split into disconnected fragments. At the group communication layer, this is referred to as a *partition*. When a partition is repaired, disconnected components merge into a larger connected component; this is referred to at the group communication layer as a *merge*. While processes are in separate disconnected components, they cannot exchange messages. Since we are interested in a practical and reasonably efficient solution, we do not consider Byzantine failures in this work.

We do not assume authenticity of membership events. Authentication of new members is obtained as part of group key management. When members leave, no explicit authentication of their departure is obtained. Furthermore, we do not assume any access control mechanisms to enforce membership policies, if any. We recognize that such mechanisms are necessary in real group applications; their development is the subject of recent and ongoing work [7], [8], [9].

Our adversary model takes into account only outside adversaries, both passive and active. An outsider is anyone who is not a current group member. Any former or future member is an outsider according to this definition. We do not consider insider attacks, as our focus is on the secrecy of group keys and the integrity of group membership. The latter means the inability to spoof authenticated membership. Consequently, insider attacks are not relevant in this context since a malicious insider can always reveal the group key or its own private key, thus allowing for fraudulent membership.

Passive outsider attacks involve eavesdropping with the aim of discovering the group key(s). Active outsider attacks involve injecting, deleting, delaying, and modifying protocol messages. Some of these attacks aim to cause denial of service and we do not address them. Attacks that aim to impersonate a group member are prevented by the use of public key signatures. Every protocol message is signed by its sender and verified by all receivers. Other more subtle active attacks aim to introduce a known (to the attacker) or old key. These attacks are prevented by the combined use of timestamps, unique protocol message identifiers, and

sequence numbers which identify the particular protocol run. This modification of GDH was formally proven secure against active adversaries in [10], [11].

3 PROBLEM DEFINITION

Our goal is to design a secure group communication service by combining a robust key agreement algorithm with a reliable group communication system (GCS). We define the semantics provided by the GCS and overview the GDH key agreement protocol suite, both of which are used later in the paper.

3.1 Group Communication Service

A GCS provides two important services: group membership and dissemination, reliability, and ordering of messages. The membership service notifies the application of the current list of group members every time the group changes. The output of this notification is called a *view*.

Several different group communication models [12], [13] have been defined in the literature, each providing a different set of semantics to the application. Many communication models claim to offer Virtual Synchrony or some variant thereof. Such claims are often based on a loose definition of *Virtual Synchrony* stating that: Processes moving together from one view to another deliver the same set of messages in the former view. However, not all the models offer the same set of properties and, to the best of our knowledge, a canonical “Virtual Synchrony (VS) Model” has not been defined in the literature. A good survey of many flavors of virtual synchrony semantics can be found in [14].

The ordering and reliability guarantees are provided within a view. In order to specify when the ordering and delivery properties are met, GCS delivers to the application an additional notification referred to as a *transitional signal*. Additional information provided with the view by a GCS is what is referred to as the *transitional set*. This set represents the set of processes that continue together with the process to which the membership notification was delivered and allows processes to locally determine if a state transfer is required. Different transitional sets may be delivered with the same view at different processes.

One property of the VS model that also has relevance for security is the *Sending View Delivery* [14] property, which requires messages to be delivered in the same view they were sent in. This enables the use of a shared view-specific key to encrypt data, since the receiver is guaranteed to have the same view as the sender and, therefore, the same key. To satisfy *Sending View Delivery* without discarding messages from group members, a GCS must block the sending of messages before the new view is installed [15]. This is achieved as follows: When a group membership change occurs, the GCS sends a message, *flush request*, to the application asking for permission to install a new view. The application must respond with a *flush acknowledgment* message which follows all the messages sent by the application in the old view. After sending the acknowledgment, the application is not allowed to send any message until the new view is delivered.

3.1.1 Virtual Synchrony Semantics

The GCS is assumed to support VS semantics as defined below. This set of properties is largely based on the survey in [14] and the definition of related semantics in [12] and [16].

We define that some event occurred in view v at process p if the most recent view installed by process p before the event was v .

1. *Self-Inclusion*: If process p installs a view v , then p is a member of v .
2. *Local Monotonicity*: If process p installs a view v after installing a view v' , then v 's identifier id_v is greater than v' 's identifier $id_{v'}$.
3. *Sending View Delivery*: A message is delivered in the view that it was sent in.
4. *Delivery Integrity*: If process p delivers a message m in a view v , then there exists a process q that sent m in v causally before p delivered m .
5. *No Duplication*: A message is sent only once. A message is delivered only once to the same process.
6. *Self Delivery*: If process p sends a message m , then p delivers m unless it crashes.
7. *Transitional Set*:
 - a. Every process is part of its transitional set for a view v .
 - b. If two processes p and q install the same view and q is included in p 's transitional set for this view, then p 's previous view was identical to q 's previous view.
 - c. If two processes p and q install the same view v and q is included in p 's transitional set for v , then p and q have the same transitional set for v .
8. *Virtual Synchrony*: Two processes p and q that move together¹ through two consecutive views v and v' deliver the same set of messages in v .
9. *FIFO Delivery*: If message m is sent before message m' by the same process in the same view, then any process that delivers m' delivers m before m' .
10. *Causal Delivery*: If message m causally precedes message m' and both are sent in the same view, then any process that delivers m' delivers m before m' .
11. *Agreed Delivery*:
 - a. Agreed delivery maintains all causal delivery guarantees.
 - b. If agreed messages m and, later, m' are delivered by process p , and m and m' are also delivered by process q , then q delivered m before m' .
 - c. If agreed messages m and, later, m' are delivered by process p in view v , and m' is delivered by process q in v before a transitional signal, then q delivers m . If messages m and, later, m' are delivered by process p in view v , and m' is delivered by process q in v after a transitional signal, then q delivers m if r , the sender of m , belongs to q 's transitional set.
12. *Safe Delivery*:
 - a. Safe delivery maintains all agreed delivery guarantees.
 - b. If process p delivers a safe message m in view v before the transitional signal, then every process q

1. If process p installs a view v with process q in its transitional set and process q installs v as well, then p and q are said to move together.

of view v delivers m unless it crashes. If process p delivers a safe message m in view v after the transitional signal, then every process q that belongs to p 's transitional set delivers m after the transitional signal unless it crashes.

13. *Transitional Signal*: Each process delivers exactly one transitional signal per view.

3.2 GDH Contributory Key Agreement Protocol

GDH IKA.2 [5] is an extension of the 2-party Diffie-Hellman key exchange protocol [17] to groups. The shared key is never transmitted over the network, even in encrypted form. Instead, a set of partial keys (that are used by individual members to compute the group secret) is sent. One particular member, group controller, is charged with the task of building and distributing the set. This is done by passing a token between the members of the group to collect contributions of the new members. The group controller is not fixed and has no special security privileges.

The protocol works as follows: When a merge event occurs, the current controller refreshes its own contribution to the group key (to prevent any incoming members from discovering the old group key), generates a new token, and passes it to one of the new members. When the chosen new member receives the token, it adds its own contribution and then passes the token to the next new member.² Eventually, the token reaches the last new member. This new member, who is slated to become the new controller, broadcasts the token to the group without adding its contribution. Upon receiving the broadcast token, each group member (old and new) factors out its contribution and unicasts the result (called a factor-out token) to the new controller. The new controller collects all the factor-out tokens, adds its own contribution to each of them, builds the set of partial keys, and broadcasts it to the group. Every member can then obtain the group key by factoring in its contribution.

When some of the members leave the group, the controller (who, at all times, is the most recent remaining group member) removes their corresponding partial keys from the set of partial keys, refreshes each partial key in the set, and broadcasts the set to the group. Every remaining member can then compute the shared key. Note that, if the current controller leaves the group, the newest remaining member becomes the group controller.

4 BASIC ROBUST ALGORITHM

This section discusses the details of a basic robust group key agreement algorithm (GKA). We describe the algorithm and prove its correctness, i.e., we show that it preserves virtual synchrony semantics presented in Section 3.1. Throughout the remainder of the paper, we use the term GCS to mean a group communication system providing virtual synchrony semantics.

4.1 Algorithm Description

The GDH IKA.2 protocol, briefly presented in Section 3.2, is secure and correct. Security is preserved independently of any sequence of membership events, while correctness holds only as long as no additional group view change

takes place before the protocol terminates. To elaborate further, consider what happens if a leave or partition event occurs while the protocol is in progress, e.g., while the group controller is waiting for individual unicasts from all group members. Since the GDH protocol does not incorporate a membership protocol (including a fault-detection mechanism), it is not aware of the membership change and the group controller does not proceed until all factor-out tokens (including those from departed members) are collected. Therefore, the system simply blocks. Similar scenarios are also possible if one of the new members crashes while adding its contribution to a group key. In this case, the token never reaches the new group controller and the GDH protocol, once again, blocks.

If the nested event is additive (join or merge), the protocol operates correctly. In other words, it runs to completion and the nested event is handled serially. However, this is not optimal since, ideally, multiple additive events ought to be "chained," effectively reducing broadcasts and factor-out token implosions. As the above examples illustrate, the GDH protocol does not operate correctly in the face of certain cascaded membership events (specifically, when the interrupting events are subtractive events). This behavior basically violates the high degree of robustness and fault tolerance of the GCS.

We propose a basic solution as follows: Each time a membership change occurs, the group deterministically selects a member (say, the oldest) to initiate the GDH merge protocol. The algorithm uses the membership service to consistently choose that member and the FIFO and Agreed ordering services to ensure that, if one member installs a secure view, all other members eventually install the same view. The approach we propose is twice as expensive in computation and requires $O(n)$ more messages for the common case with no cascading membership events, n being the group size. We discuss it because it is simpler and it allows us to show algorithm correctness with respect to the group communication semantics and stated security goals. In Section 5, we present an optimized algorithm that offers better performance and uses the basic algorithm as a "subroutine" in exceptional cases.

Since the output of the algorithm is a secure GCS, VS semantics as defined in Section 3.1 must be preserved. To achieve this, our algorithm takes extra care to provide delivery of the correct views, transitional signal, and transitional sets to the applications, as well as the list of connected group members.

We model the algorithm as a state machine (see Fig. 2), where transitions from one state to another take place based on the event that occurred. An event is defined as receiving a particular type of message. In Fig. 2, all transitions numbered with the same number denote the same set of events and actions for that particular state. The following types of messages are used: GDH messages (see [18]) (partial_to ken_msg, final_token_msg, key_list_msg, fact_out_msg); membership notification messages (memb_msg); transitional signal messages (trans_signal_msg); data messages (data_msg); flush mechanism messages (flush_request_msg, flush_ok_msg).

Fig. 1 presents the secure GCS protocol stack. Our group key agreement (GKA) protocol interacts with both the application and GCS and implements the blocking mechanism as follows: When a flush_request_msg message is received from GCS, it is delivered to the application. When

2. The set of new members and its ordering is decided by the underlying group communication system. The actual order is irrelevant to GDH.

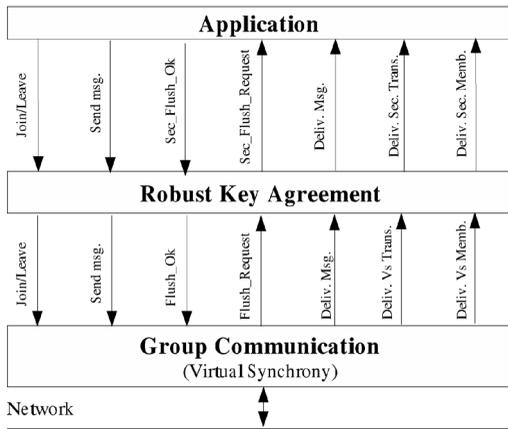


Fig. 1. Secure GCS protocol stack.

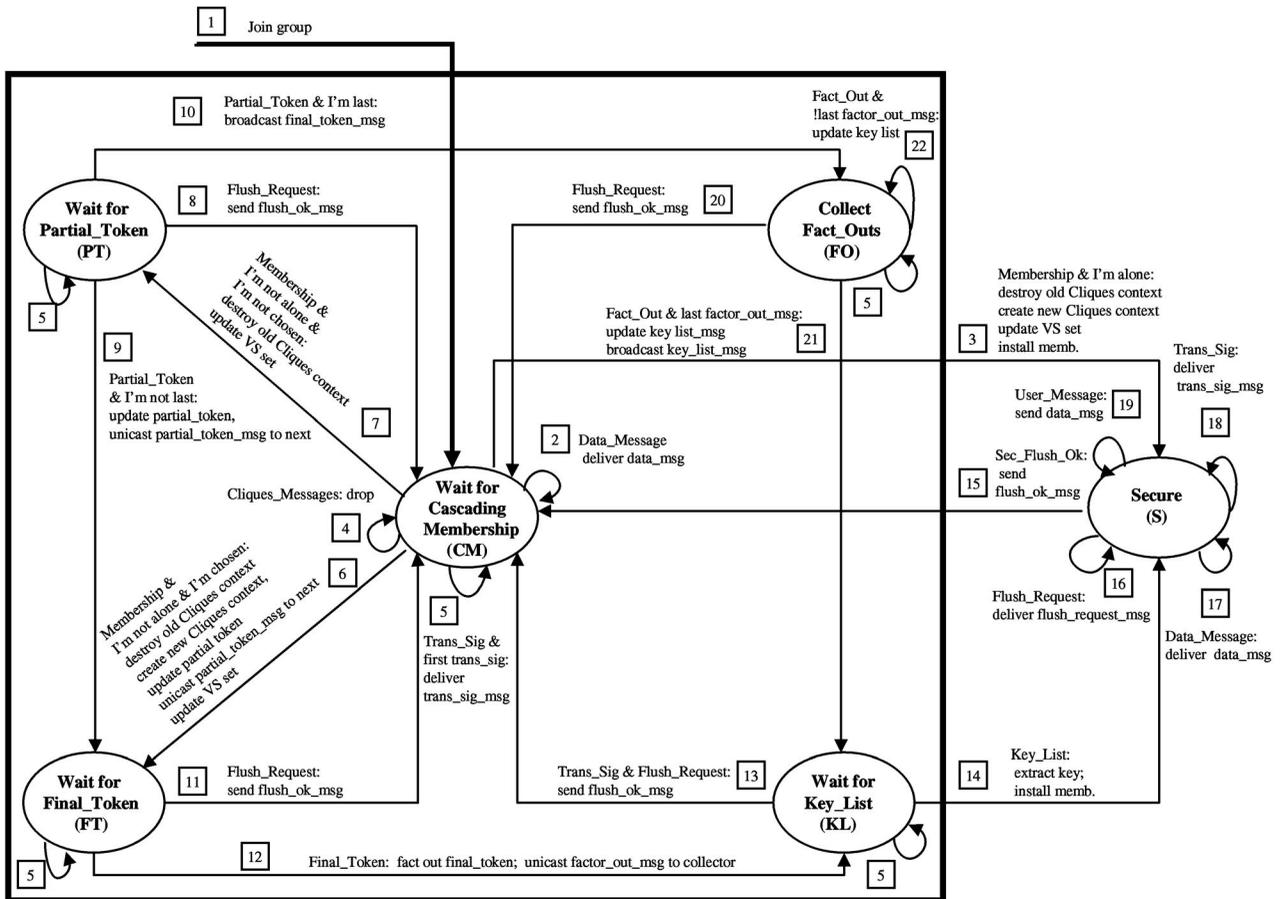
the application acknowledgment message is received, it is sent down to the GCS.

A process starts executing the algorithm by invoking the *join* primitive of the key agreement module, which translates into a GCS join call. In any state of the algorithm, a process can voluntarily leave by invoking the *leave* primitive of the key agreement module, which translates into a GCS leave call.

The set of events that can trigger transitions from one state to another are presented in Table 1. The events are associated with a specific group and are received by the GKA. Note that both *User_Message* and *Data_Message* events are associated with a *data_msg* message received by the GKA, but, in the first case, the source of the message is an application, while, in the second case, the source is the GCS. Also note that all messages specific to the GDH protocol are particular cases of *data_msg*. We define specific messages and associate events with them to simplify the description of the protocol.

Each state of the algorithm (in Fig. 2) is described by three types of events: events that trigger transitions, events that are considered illegal (in which case, an error is returned to the application), and events that should not occur when the algorithm operates correctly. The following states are used:

- SECURE (S): group is functional, all members have the group key and can communicate securely; possible events are: *Data_Message*, *User_Message*, *Secure_Flush_Ok*, *Flush_Request*, and *Transitional_Signal*; receiving a *Secure_Flush_Ok* without receiving a corresponding *Flush_Request* is illegal;
- WAIT_FOR_PARTIAL_TOKEN (PT): process is waiting for the token accumulating contributions; possible



Notes: VS_set is delivered as part of the membership
 : All Cliques messages but key_list_msg are sent FIFO; key_list_msg is sent as an AGREED message
 : A process can leave the group in any state

Fig. 2. Basic GKA.

TABLE 1
Events Received by GKA

Event	Message	Source
Partial-Token	partial_token_msg	GCS
Final-Token	final_token_msg	GCS
Fact_Out	factor_out_msg)	GCS
Key_List	key_list_msg	GCS
User_Message	data_msg	App.
Data_Message	data_msg	GCS
Transitional_Signal	trans_signal_msg	GCS
Membership	memb_msg	GCS
Flush_Request	flushrequest_msg	GCS
Secure_Flush_Ok	flushok_msg	App.

events are: Partial-Token, Flush_Request, and Transitional_Signal; User_Message and Secure_Flush_Ok are illegal;

- WAIT_FOR_FINAL_TOKEN (FT): process is waiting for the token containing all contributions; possible events are: Final-Token, Flush_Request, and Transitional_Signal; User_Message and Secure_Flush_Ok are illegal;
- COLLECT_FACT_OUTS (FO): process is waiting for $n - 1$ factor out messages; possible events are: Fact_Out, Flush_Request, and Transitional_Signal; User_Message and Secure_Flush_Ok are illegal;
- WAIT_FOR_KEY_LIST (KL): process is waiting for the set of partial keys; possible events are: Key_List, Flush_Request, Transitional_Signal; User_Message, and Secure_Flush_Ok are illegal;
- WAIT_FOR_CASCADING_MEMBERSHIP (CM): process is waiting for membership and transitional signal messages; possible events are: Membership, Transitional_Signal, Data_Message (possible only the first time the process gets in this state), Partial-Token, Final-Token, Fact_Out, and Key_List (they correspond to GDH messages from a previous instance of the GKA when cascaded events happen); User_Message and Secure_Flush_Ok are illegal.

The state machine is built around the CM state, which is used to restart the protocol. Four other states are just a map of the GDH merge protocol and are used to pass the token accumulating contributions and used to build the set of partial keys. The S state is the operational state. The pseudocode corresponding to the state machine from Fig. 2 and the correctness proofs are presented in Appendix 1 which can be found on the Computer Society Digital Library at <http://computer.org/tpds/archives.htm>.

4.2 Security Considerations

The GDH protocol was proven secure against passive adversaries in [5]. As evident from the state machine in

Fig. 2, the protocol remains intact, i.e., all protocol messages are sent and delivered in the same order as specified in [5]. More precisely, with no cascaded events, our protocol is exactly the same as the original GDH join protocol [5]. In the case of a cascaded event, the protocol is the same as the IKA.2 [5] group key agreement protocol. Since both of these protocols are proven secure, our robust protocol is, therefore, also provably secure. In this context, security means that it is computationally infeasible to compute a group key by passively observing any number of protocol messages. As discussed in Section 2, stronger, active attacks are averted by the combined use of timestamps, protocol message type, and protocol run identifiers, explicit inclusion of message source and destination and, most importantly, digital signatures by the source of the message. These measures make it impossible for the active adversary to impersonate a group member or to interfere with the key agreement protocol and thereby influence or compute the eventual group key [10], [11].

5 AN OPTIMIZED ROBUST ALGORITHM

In this section, we show how the algorithm presented in the previous section can be optimized, resulting in lower-cost handling of common, noncascaded events, while preserving the same set of group communication and security guarantees.

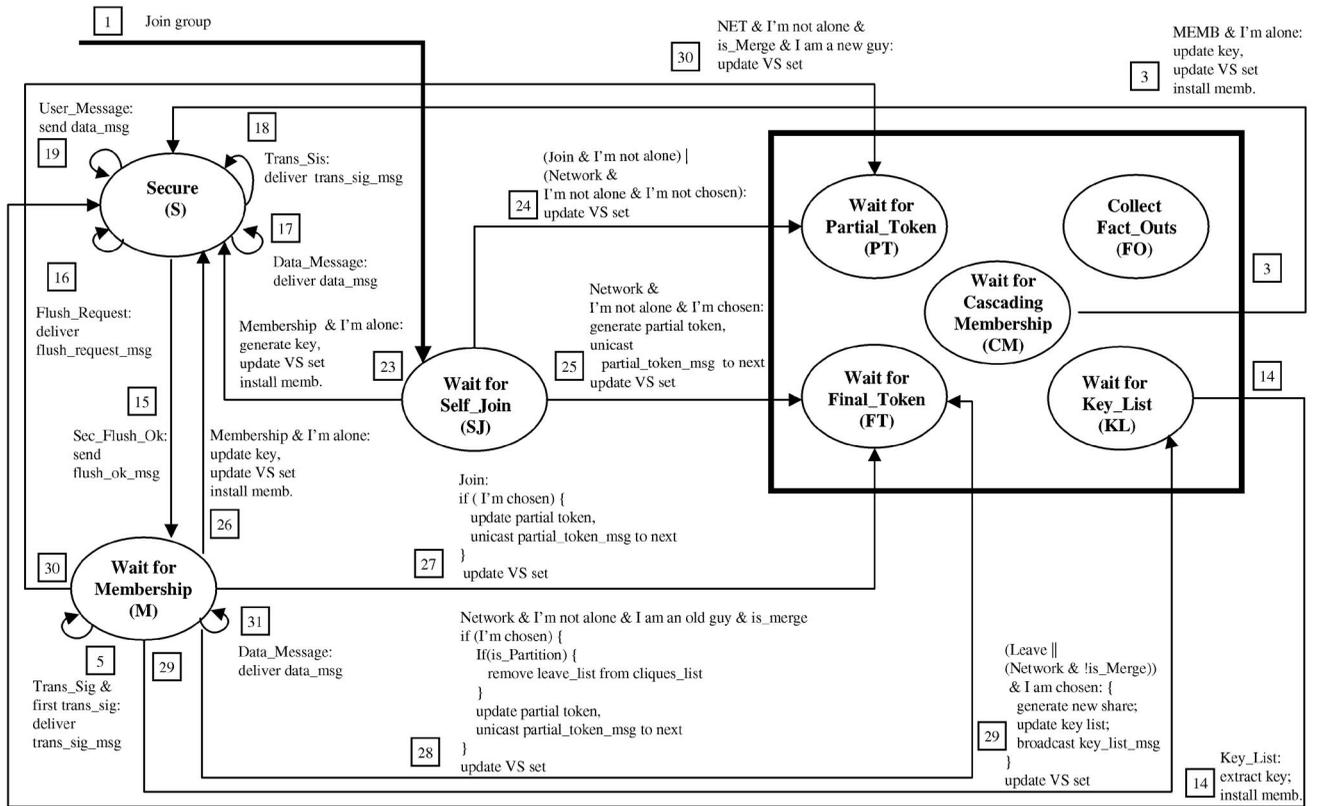
5.1 Algorithm Description

The basic algorithm presented in Section 4 is robust even when cascaded group events occur. Every time a membership notification is delivered by the GCS, the algorithm “forgets” all previous key agreement information (i.e., the set of partial keys) and restarts the merge protocol, selecting a member from the new group to initialize it. Thus, this algorithm costs more than necessary since it does not attempt to use the existing accumulated information (partial keys) and avoid unnecessary computation.

We propose to improve the basic protocol by using optimized protocols for each type of group change (join, leave, partition, merge, or a combination of partition and merge) and by taking advantage of the already existing set of partial keys. We also utilize the basic algorithm to handle more complex cascaded membership events. For example, in case of a leave, the leave protocol is invoked, which requires the group controller to remove the leaving member(s) from the set, refresh the set of partial keys, and broadcast it. Thus, leave events can be handled immediately, with lower communication and computation costs than those of in the basic algorithm. In Section 5.2, we discuss how a combined event—including both joins and leaves—can be handled by a modified version of the GDH merge protocol.

The optimized algorithm is modeled by a state machine that, in addition to the states in the basic algorithm, uses two more states, as shown in Fig. 3. Each state is described by three types of events: events that trigger transitions, events that should never occur when the algorithm operates correctly, and events that are considered illegal:

- WAIT_FOR_SELF_JOIN (SJ): initial state wherein a process that joins a group enters the state machine; the process is waiting for the membership message



Notes: VS_set is delivered as part of the membership
 : All Cliques messages but key_list_msg are sent FIFO; key_list_msg is sent as an AGREED message.
 : A process can leave the group in any state

Fig. 3. Optimized GKA.

that notifies the group about its joining. In case a network event happens between the join request and the membership notification delivery, the GCS will report a network event and the transitional set will contain only the joining member; possible event is Membership; User_Message and Secure_Flush_Ok are illegal.

- WAIT_FOR_MEMBERSHIP (M): process is waiting for a membership notification; possible events are: Transitional_Signal, Data_Message, and Membership; User_Message and Secure_Flush_Ok events are illegal.

While a process starts the basic algorithm in the CM state, in the optimized algorithm, a process starts the algorithm in state SJ, by invoking the *Join* primitive. At any time, a process can voluntarily leave the algorithm by invoking the *Leave* primitive. The main difference between the robust and the optimized algorithm is that, in case of a membership change, the process moves to the M state and tries to handle the event depending on its nature (subtractive, additive, or both). In case of cascading memberships, everything is abandoned and the basic algorithm is invoked by moving to the CM state.

The pseudocode corresponding to the state machine from Fig. 3 and the correctness proofs are presented in Appendix 2 which can be found on the Computer Society Digital Library at <http://computer.org/tpds/archives.htm>.

5.2 Handling Bundled Events

Most group events are homogeneous in nature: leave (partition) or join (merge) of one or more members. However, a GCS can decide to bundle several such events if they occur within a very short time interval. The main incentive for doing so is to reduce the impact and overhead on the application.

Recall that GDH defines two separate protocols for leave and merge. Each of them can trivially handle bundled events of the same type: The GDH merge protocol can accommodate any combination of bundled merges, while the GDH leave protocol can do the same for any combination of partitions. A more interesting scenario is when a single membership event bundles merges/joins with leaves/partitions. One way to handle such an event is to first invoke GDH leave to process all leaves/partitions and then invoke GDH merge to process joins/merges. However, this is inefficient since the group would perform two separate key agreement protocols where only one is truly needed. Since both GDH protocols are initiated by the group controller, we propose the following optimized solution. After processing all leaves/partitions, the group controller can suppress the usual broadcast of new partial keys and, instead, forward the resulting set to the first merging/joining member thereby initiating a merge protocol. This saves an extra round of broadcast and at least one cryptographic operation for each member.

5.2.1 Security Considerations

Recall that, in the merge protocol, the current controller begins by refreshing its contribution (to the group key) and forwarding the result to the first merging member. This message actually contains a set of partial keys, one for each "old" member and an extra partial key for the first new member. This message is also signed by the controller and includes the list of all members believed by the controller to be in the group at that instant. In the optimized protocol, the controller effectively suppresses all partial keys corresponding to members who are leaving the group. This modification changes nothing as far as any outside attacks or threats are concerned. The only issue of interest is whether any members leaving the group can obtain the new key. We claim that this is impossible since the set of partial keys forwarded by the controller is essentially the same as the partial key set broadcast in the normal leave protocol. Therefore, former members are no better off in the optimized than in the leave protocol. Also, the new (merging) members are still unable to compute any prior group keys just as in the plain merge protocol. This is because the information available to the new members in the optimized protocol is identical to that in the plain merge.

6 IMPLEMENTATION

We implemented the optimized algorithm described above using the Spread [6] GCS and the Cliques key agreement library. In this section, we overview the Spread and Cliques toolkits and present the concrete outcome of this work, the Secure Spread library.

6.1 The Spread Toolkit

Spread [6] is a general-purpose GCS for wide and local-area networks, where any group member can be both a sender and a receiver. Although designed to support small to medium-size groups, it can accommodate a large number of collaboration sessions, each spanning the Internet.

The main services provided by the system are reliable and ordered delivery of messages (FIFO, causal, total/Agreed order, safe) and a membership service in a model that considers benign network and computer faults (crashes, recoveries, partitions and merges). Spread supports two well-known semantics, Virtual Synchrony (VS) [13], [16] and Extended Virtual Synchrony (EVS) [12], [19]. In this work, we use only the latter.

The Spread toolkit is publicly available and is being used by several organizations in research, educational, and production settings. It supports cross-platform applications and has been ported to several Unix platforms as well as to Windows and Java environments.

6.2 The Cliques Toolkit

Cliques is a cryptographic toolkit providing key management services for dynamic peer groups. The toolkit assumes the existence of a communication platform for transporting protocol messages and maintaining group membership. It includes several protocol suites: GDH, CKD, TGDH [20], STR [21], [22], and BD [23]. GDH is based on group extensions of the 2-party Diffie-Hellman key exchange [5] and provides fully contributory group key agreement.

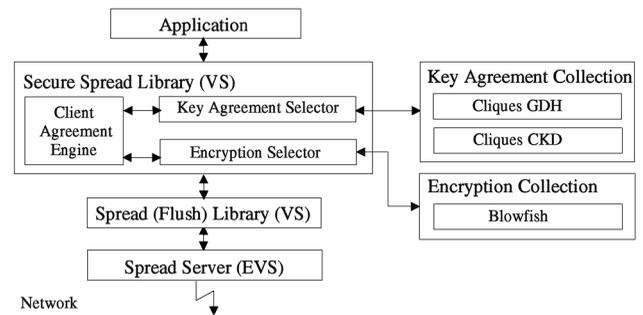


Fig. 4. Secure spread architecture.

TGDH, STR, and BD are Diffie-Hellman-based protocols that make different trade offs of communication and computation.

All Cliques protocol suites offer key independence, perfect forward secrecy, and resistance to known key attacks. (See [3] for precise definitions of these properties.) In this paper, we focus only on the GDH protocol suite within the Cliques toolkit.

6.3 The Secure Spread Library

The Secure Spread library provides client data confidentiality and integrity. It is built on top of the VS Spread client library; it uses Spread as its communication infrastructure and Cliques [18] library for key management.

Fig. 4 shows the architecture of our secure GCS. The Flush library is the component of the Spread Toolkit providing the Virtual Synchrony model as described in Section 3.1.³

The core of the library is the Client Agreement Module, which is the connection between the library and the GCS. When it receives a notification from GCS about a group membership change, the module starts the key agreement protocol. When the key agreement protocol completes and a new key is available, the module delivers a secure group membership change notification to the application.

The library has two components: Key Agreement Selector and Encryption Selector that allow, respectively, the selection of a specific key agreement module and a specific encryption module.

Secure Spread currently has two different modules for key agreement, both using primitives provided by the Cliques library: the robust optimized GDH protocol (presented in this paper) and a centralized key management protocol (described below), both having the same security properties. The architecture allows each group to run its own key agreement protocol. The library uses Blowfish for encryption.

6.4 Centralized Key Distribution Protocol

In general, centralized key distribution protocols do not provide key independence since, for efficiency reasons, they rely on previous group or subgroup keys to distribute new

3. Actually, the Flush library provides all the properties we described in Section 3.1 but one. It does not deliver exactly one transitional signal per view. However, in the Flush library the Flush_Request and Transitional_Signal events are delivered in AGREED order. Using this property, with a minor modification, our GKA can avoid generating unnecessary transitional signals for the application.

keys. When using such a method, the compromise of some group keys can lead to the compromise of other group keys. To compare protocols having the same security properties, we designed a Centralized Key Distribution (CKD) scheme that provides the same level of security as GDH, as far as key independence and PFS [3].

The group secret is always generated by one member, the current group controller.⁴ Following each membership change, the controller generates a new secret and distributes it securely to the group. Of course, an efficient symmetric cipher can be used to securely distribute the group key. However, the resulting security properties would differ from those of our key agreement protocol, which relies solely on the Decision Diffie-Hellman assumption [24] and the Discrete Logarithm Problem [3]. Therefore, to provide an equivalent level of security, we encrypt the group key via modular exponentiation.

The controller in CKD is always the oldest member. Regardless of the group operation, the CKD protocol consists of two phases (see also Algorithm 1):

Algorithm 1 CKD protocol

Let (x_1, α^{x_1}) and $(x_{n+1}, \alpha^{x_{n+1}})$ be the secret and public keys of M_1 , the group controller, and M_{n+1} respectively.

Let $K_{1n+1} = \alpha^{x_1 x_{n+1}} \bmod p$. Assume that the group has n members, and that M_{n+1} wants to join the group.

Round 1:

M_1 selects random $r_1 \bmod q$ (this selection is performed only once),

$M_1 \rightarrow M_{n+1} : \alpha^{r_1} \bmod p$

Round 2:

M_{n+1} selects random $r_{n+1} \bmod q$,

$M_1 \leftarrow M_{n+1} : \alpha^{r_{n+1}} \bmod p$

Round 3:

M_1 selects a random group secret K_s and computes

$M_1 \rightarrow M_i : K_s^{\alpha^{r_1 r_i}} \bmod p \quad \forall i \in [2, n+1]$

1. Each group member and the controller agree on a unique pairwise key using authenticated two-party Diffie-Hellman. This key does not need to change as long as both users remain in the group. If the controller leaves the group, the new controller has to perform this operation with every member. If a regular member leaves, the controller simply discards this pairwise key.
2. The group controller unilaterally generates and distributes the group secret.

7 PERFORMANCE EVALUATION

In this section, we compare the GDH key agreement protocol with the CKD protocol presented in Section 6.4, in a LAN environment. We evaluate the time it takes the system to establish secure membership for the most common group events: join and leave.

Cryptography relies on expensive exponentiations, so it seems that measuring CPU time will be a good approach

4. We use the term *current* to mean that a controller can fail or be partitioned out, thus causing the controller role to be reassigned to the oldest surviving member.

demonstrating the cryptographic overhead. Table 2 presents the number of serial exponentiations for a join or leave event, where n is the group size before the operation, while m and p represent the number of new and partitioned members, respectively. A more relevant measure for a GCS is the latency that a user experiences from the moment the group change was detected until the new secure group is established. This time is greater than just the analytical cryptographic cost since it includes network latency. It can also exhibit increased computation cost if several processes compete for the same CPU.

Our experimental testbed is a cluster of 13 667 MHz Pentium III dual-processor PCs running Linux. Each machine runs a Spread server, while group members are uniformly distributed on the machines. Therefore, more than one process can be running on a single machine (which is frequent in many collaborative applications).

Each member measures the time it took to complete the key agreement and establish a secure view. We compute the average cost of the membership service and secure membership with GDH and CKD, respectively. This time was averaged over 20 distinct runs of the experiment.

Since CKD is particularly expensive if the current controller leaves the group, we take this into account by considering that, with $1/n$ probability, the member leaving the group is the group controller.

Experiments performed on our testbed for the insecure GCS show that the average cost of sending and delivering one agreed multicast message is almost constant, ranging between 0.75 and 0.92 milliseconds for a group size ranging between 2 to 50 members. The cost of the membership service (see Figs. 5 and 6) is negligible with respect to key agreement overhead, varying between 2 and 8 milliseconds for a group size between 2 and 50 members. We use 1,024-bit RSA signatures for message origin and data authentication, with the public exponent of 3 to reduce the verification overhead. On our hardware platform, the RSA sign and verify operations take 9.6 and 0.2 milliseconds, respectively. For the short-term group key, we use both 512 and 1,024-bit Diffie-Hellman parameter p and 160-bit q . The cost of a single exponentiation is 1.7 and 5.3 milliseconds for a 512 and a 1,024-bit modulus, respectively.

Figs. 5 and 6 present the respective costs of:

1. our robust GDH key agreement protocol,
2. the centralized protocol (CKD), and
3. the insecure group communication membership service.

Note that the cost of the membership service is insignificant when compared to key agreement overhead.

For join (Fig. 5), the contributory protocol is more expensive than CKD. For example, for a group of 20 members, the time to install a secure membership is about 142 milliseconds for the contributory protocol, while, for the centralized protocol, it is about 91 milliseconds, when the 512-bit modulus is used. The difference between CKD and GDH comes from exponentiation and signature verifications: extra operations in GDH include n verifications, one RSA signature, and one modular exponentiation.

For leave events (see Fig. 6), the centralized protocol is more expensive. For the 512-bit modulus, for a group of 20 members, it takes about 49 milliseconds to establish secure membership with the contributory protocol, while

TABLE 2
Communication and Computation Costs

		Communication				Computation		
		Rounds	Messages	Unicast	Multicast	Exp.	Sign.	Verif.
GDH	Join	4	$n + 3$	$n + 1$	2	$n + 3$	4	$n + 3$
	Leave	1	1	0	1	$n - 1$	1	1
	Merge	$m + 3$	$n + 2m + 1$	$n + 2m - 1$	2	$n + 2m + 1$	$m + 3$	$n + 2m + 1$
	Partition	1	1	0	1	$n - p$	1	1
CKD	Join	3	3	2	1	$n + 2$	3	3
	Leave	1	1	0	1	$n - 2$	1	1
	Merge	3	$m + 2$	m	2	$n + 2m$	3	$m + 2$
	Partition	1	1	0	1	$n - p - 1$	1	1
	Controller Leave	3	$3n - 6$	$2n - 4$	2	$2n - 3$	$2n - 2$	n

the centralized protocol takes about 58 milliseconds. The overhead of the centralized protocol over the contributory protocol comes from one exponentiation, plus the cost of establishing $n - 1$ secure channels if the leaving member is the group controller.

Both protocols scale linearly with group size in the number of exponentiations. Performance deteriorates when 1,024-bit modulus is used for shared key generation, as shown in Figs. 5 and 6.

8 RELATED WORK

In this section, we consider related work in group key management and reliable group communication.

8.1 Group Key Management

Cryptographic techniques for securing all types of multicast or group-based protocols require all parties to share a common key. This requires a group key management protocol to generate new group keys and update existing keys. Group key management protocols generally fall into two classes:

- Protocols designed for large-scale (e.g., IP Multicast) groups, with a one-to-many communication paradigm and relatively weak security requirements [25],

[26]. Most of such protocols are centralized key distribution schemes.

- Protocols designed to support medium size tightly coupled dynamic peer groups, with a many-to-many communication paradigm and strong security requirements [23], [5]. Both distributed group key distribution and group key agreement methods are applicable to such settings.

Many protocols of the first type are being developed in the context of IETF/IRTF: Group Key Management Protocol (GKMP) [25], Multicast Key Management Protocol (MKMP) [27], Scalable Multicast Key Distribution [28], the Intradomain Group Key Management work of [26], One-way Function Trees [29], Group Secure Association Key Management Protocol (GSAKMP) [30], GSAKMP-light [31], Group Domain of Interpretation (GDOI) [32], while [33] defines an architecture for large scale group key management. Since the focus of our work is on dynamic peer groups key management, we discuss only distributed group key distribution and contributory key agreement protocols.

Most group key agreement schemes [5], [22], [23], [20], [34], [21] extend the well-known Diffie-Hellman key exchange [17] method to groups of n parties. Steer et al. proposed a group key agreement protocol [22] for static conferencing. While the protocol is well-suited for adding new group members as it takes only two rounds and two

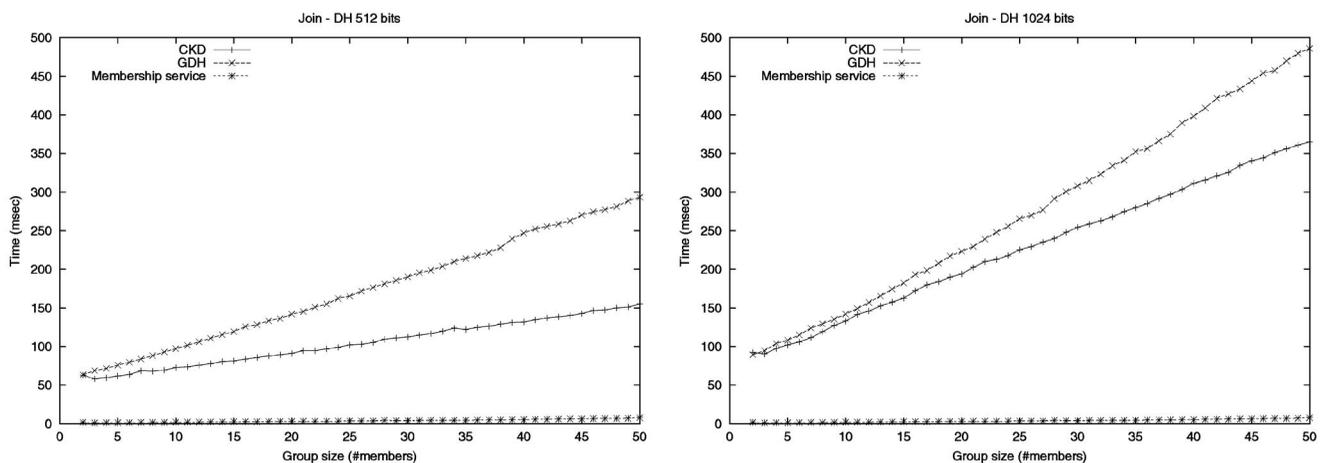


Fig. 5. Join—average time.

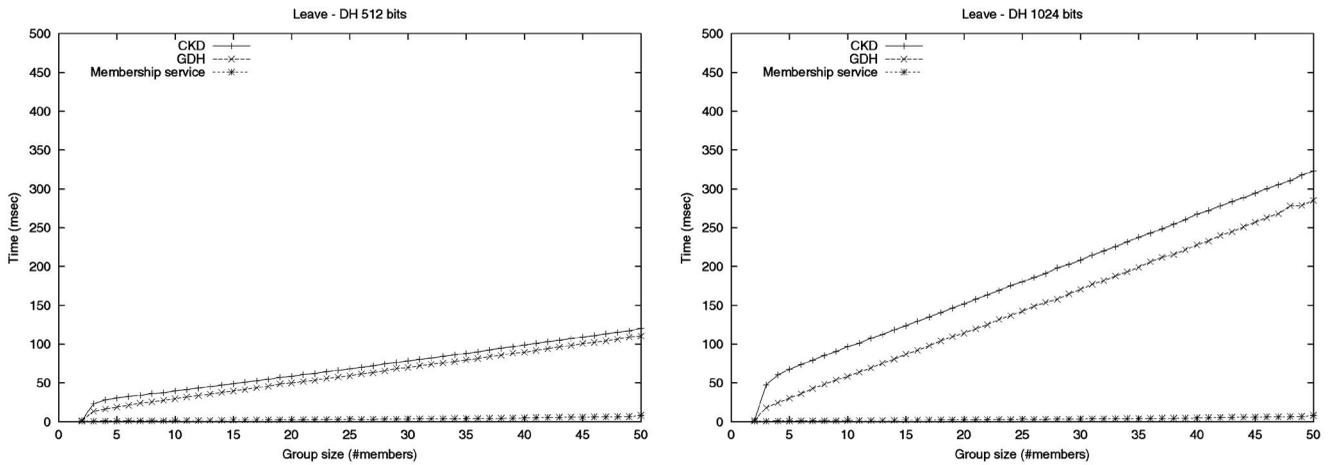


Fig. 6. Leave—average time.

modular exponentiations, it is relatively expensive when excluding members. In 1994, Burmester and Desmedt [23] proposed an efficient protocol that takes only three rounds and three modular exponentiations per member to generate a group key. This protocol allows all members to recompute the group key for any membership change with a constant small CPU cost. However, it requires $2n$ broadcast messages, which can be expensive on a wide area network. Tzeng and Tzeng also proposed an elegant authenticated key agreement scheme based on secure multiparty computation [34]. Their protocol is optimized in terms of communication rounds, but also uses $2n$ simultaneous broadcast messages. The resulting group key does not provide PFS, which represents a major drawback.

Steiner et al. address dynamic membership issues [5] in group key agreement and propose a family of protocols based on straight forward extensions of the two-party Diffie-Hellman protocol. Their protocol suite is fairly efficient in leave and partition operation, but the merge protocol requires as many rounds as the number of new members to complete key agreement. The entire protocol suite has been proven secure with respect to both passive and active attacks. Follow-on work yielded more efficient protocols in either communication or computation [20], [21].

Dynamic group key distribution methods are also amenable to dynamic peer groups. Centralized Key Distribution (CKD) is a simple example of distributed key distribution (see Section 6.4), where the oldest group member acts as a key distribution center and, in the event of a partition or a leave of the center, the role shifts to the oldest remaining member. Rodeh et al. proposed more advanced key distribution protocols, combining a key tree structure with dynamic key server election [35] or taking advantage of efficient data structures such as AVL trees [36]. Although they have some of the disadvantages of key distribution schemes, the communication and computation costs are appreciably lower than those in CKD.

8.2 Reliable Group Communication

Reliable group communication in LAN environments has a long history beginning with ISIS [37], and more recent systems such as Transis [38], Horus [39], Totem [40], and RMP [41]. These systems explored several different models of group communication such as Virtual Synchrony [4] and Extended Virtual Synchrony [12]. More recent work in this

area focuses on scaling group membership to wide-area networks [42], [43]. Research on securing group communication is fairly new. The only implementations of GCS that focus on security (in addition to ours) are the SecureRing [44] project at UCSB, the Horus/Ensemble work at Cornell [35], and the Rampart system at AT&T [45].

Some GCSs (Rampart and SecureRing) address Byzantine failures. They suffer from limited performance since they use costly protocols and rely intensively on public key cryptography. Rampart builds the group multicast protocols over a secure group membership protocol, while SecureRing system protects a low-level ring by authenticating each transmission of the token and data message received.

The Ensemble work is state-of-the-art in secure reliable group communication. It allows application-dependent trust models and optimizes certain aspects of the group key generation and distribution protocols. Ensemble achieves data confidentiality by using a shared group key obtained by means of group key distribution protocols. In comparison with our approach, although efficient, the scheme does not provide forward secrecy, key independence, and PFS.

Some other approaches focus on building highly configurable dynamic distributed protocols. Cactus [46] is a framework that allows the implementation of configurable protocols as composition of microprotocols. Survivability of the security services is enhanced by using redundancy [47].

Another toolkit that can be used to build secure group oriented applications is Enclaves [48]. It provides group control and communication (both unicast and multicast) and data confidentiality. The system uses a centralized key distribution scheme where a member of the group (group leader) selects a new key every time the group changes and securely distributes it to all group members. The main drawback of the system is that it does not address failure recovery when the leader of the group fails.

Antigone [49] is a framework that provides mechanisms which allow flexible application security policies. The system implements group rekeying mechanisms in two flavors: session rekeying—all group members receive a new key, and session key distribution—the session leader transmits an existing session key. Both schemes present problems, distributing the same key when the group

changes breaks PFS, while the session rekeying mechanism does not recover from the leader's failure.

9 CONCLUSIONS

In this paper, we showed that, although difficult, it is possible to harden security protocols to make them robust to asynchronous network events. In particular, we demonstrated how robust contributory key agreement protocols can be designed by taking advantage of group communication services. We presented two such robust protocols based on the GDH key protocol suite and the Virtual Synchrony group communication semantics. We also showed how such protocols can be used to design secure group communication services and argued that, by integrating them with a GCS supporting Virtual Synchrony, group communication membership and ordering guarantees are preserved. We exemplified by presenting Secure Spread, a client library that uses Spread as its GCS and relies on a group key management protocol that is robust to process crashes and network partitions and merges and protects confidentiality of the data even when long-term keys of the participants are compromised.

ACKNOWLEDGMENTS

This work was supported in part by a grant from the National Security Agency under the LUCITE program and by grant F30602-00-2-0526 from the US Defense Advanced Research Projects Agency. Parts of this work have appeared as conference publications in ICDCS 2000 [1] and ICDCS 2001 [2].

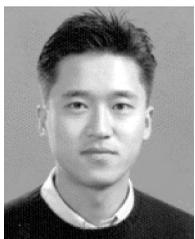
REFERENCES

- [1] Y. Amir, G. Ateniese, D. Hasse, Y. Kim, C. Nita-Rotaru, T. Schlossnagle, J. Schultz, J. Stanton, and G. Tsudik, "Secure Group Communication in Asynchronous Networks with Failures: Integration and Experiments," *Proc. 20th IEEE Int'l Conf. Distributed Computing Systems*, pp. 330-343, Apr. 2000.
- [2] Y. Amir, Y. Kim, C. Nita-Rotaru, J. Schultz, J. Stanton, and G. Tsudik, "Exploring Robustness in Group Key Agreement," *Proc. 21st IEEE Int'l Conf. Distributed Computing Systems*, pp. 399-408, Apr. 2001.
- [3] A. Menezes, P. van Oorschot, and S. Vanstone, *Handbook of Applied Cryptography*. year?
- [4] K.P. Birman and T. Joseph, "Exploiting Virtual Synchrony in Distributed Systems," *Proc. 11th Ann. Symp. Operating Systems Principles*, pp. 123-138, Nov. 1987.
- [5] M. Steiner, G. Tsudik, and M. Waidner, "Key Agreement in Dynamic Peer Groups," *IEEE Trans. Parallel and Distributed Systems*, vol. 11, no. 8, Aug. 2000.
- [6] Y. Amir and J. Stanton, "The Spread Wide Area Group Communication System," Johns Hopkins Univ., Center of Networking and Distributed Systems, Technical Report 98-4, 1998.
- [7] Y. Amir, C. Nita-Rotaru, and J. Stanton, "Framework for Authentication and Access Control of Client-Server Group Communication Systems," *Proc. Third Int'l Workshop Networked Group Comm.*, Nov. 2001.
- [8] Y. Kim and G. Tsudik, "Membership Control in Peer Groups," *Proc. Workshop New Directions on Scalable Cyber-Security*, Mar. 2003.
- [9] Y. Kim, D. Mazzocchi, and G. Tsudik, "Admission Control in Collaborative Groups," *Proc. Second IEEE Int'l Symp. Network Computing and Applications*, Apr. 2003.
- [10] E. Bresson, O. Chevassut, and D. Pointcheval, "Provably Authenticated Group Diffie-Hellman Key Exchange—The Dynamic Case," *Proc. Asiacrypt*, 2001.
- [11] E. Bresson, O. Chevassut, D. Pointcheval, and J.-J. Quisquater, "Provably Authenticated Group Diffie-Hellman Key Exchange," *Proc. Eighth ACM Conf. Computer and Comm. Security*, Nov. 2001.
- [12] L.E. Moser, Y. Amir, P.M. Melliar-Smith, and D.A. Agarwal, "Extended Virtual Synchrony," *Proc. IEEE 14th Int'l Conf. Distributed Computing Systems*, pp. 56-65, June 1994.
- [13] A. Fekete, N. Lynch, and A. Shvartsman, "Specifying and Using a Partitionable Group Communication Service," *ACM Trans. Computer Systems*, vol. 19, no. 2, pp. 171-216, May 2001.
- [14] G.V. Chockler, I. Keidar, and R. Vitenberg, "Group Communication Specifications: A Comprehensive Study," *ACM Computing Surveys*, no. 4, pp. 427-469, Dec. 2001.
- [15] R. Friedman and R. van Renesse, "Strong and Weak Virtual Synchrony in Horus," Cornell Univ., Computer Science, Technical Report 95-1537, Aug. 1995.
- [16] J. Schultz, "Partitionable Virtual Synchrony Using Extended Virtual Synchrony," master's thesis, Dept. of Computer Science, Johns Hopkins Univ., Jan. 2001.
- [17] W. Diffie and M.E. Hellman, "New Directions in Cryptography," *IEEE Trans. Information Theory*, vol. 22, pp. 644-654, Nov. 1976.
- [18] G. Ateniese, O. Chevassut, D. Hasse, Y. Kim, and G. Tsudik, "Design of a Group Key Agreement API," *Proc. DARPA Information Security Conf. and Exposition*, Jan. 2000.
- [19] Y. Amir, "Replication Using Group Communication over a Partitioned Network," PhD dissertation, Inst. of Computer Science, The Hebrew Univ. of Jerusalem, Israel, 1995.
- [20] Y. Kim, A. Perrig, and G. Tsudik, "Simple and Fault-Tolerant Key Agreement for Dynamic Collaborative Groups," *Proc. Seventh ACM Conf. Computer and Comm. Security*, pp. 235-244, Nov. 2000.
- [21] Y. Kim, A. Perrig, and G. Tsudik, "Communication-Efficient Group Key Agreement," *Proc. Int'l Conf. Information Security IFIP SEC*, June 2001.
- [22] D. Steer, L. Strawczynski, W. Diffie, and M. Wiener, "A Secure Audio Teleconference System," *Proc. Conf. Advances in Cryptology*, Aug. 1990.
- [23] M. Burmester and Y. Desmedt, "A Secure and Efficient Conference Key Distribution System," *Proc. Conf. Advances in Cryptology*, May 1994.
- [24] D. Boneh, "The Decision Diffie-Hellman Problem," *Lecture Notes in Computer Science*, vol. 1423, pp. 48-63, 1998.
- [25] H. Harney and C. Muckenhirn, "Group Key Management Protocol (GKMP) Specification," RFC 2093, July 1997.
- [26] T. Hardjono, B. Cain, and I. Monga, "Intradomain Group Key Management Protocol," Sept. 2000.
- [27] D. Harkins and N. Doraswamy, "A Secure Scalable Multicast Key Management Protocol (MKMP)," Nov. 1997.
- [28] T. Ballardie, "Scalable Multicast Key Distribution," RFC 1949, 1996.
- [29] D. Balenson, D. McGrew, and A. Sherman, "Key Management for Large Dynamic Groups: One-Way Function Trees and Amortized Initialization," 2000.
- [30] H. Harney, A. Schuett, U. Meth, and A. Colegrove, "GSAKMP," Feb. 2003.
- [31] H. Harney, A. Schuett, and A. Colegrove, "GSAKMP Light," July 2002.
- [32] M. Baugher, T. Hardjono, H. Harney, and B. Weis, "The Group Domain of Interpretation," Dec. 2002.
- [33] M. Baugher, R. Canetti, L. Dondeti, and F. Lindholm, "Group Key Management Architecture," 2002.
- [34] W.-G. Tzeng and Z.-J. Tzeng, "Round-Efficient Conference-Key Agreement Protocols with Provable Security," *Proc. Conf. Advances in Cryptology*, Dec. 2000.
- [35] O. Rodeh, K. Birman, and D. Dolev, "The Architecture and Performance of Security Protocols in the Ensemble Group Communication System," *ACM Trans. Information and System Security*, vol. 4, no. 3, pp. 289-319, Aug. 2001.
- [36] O. Rodeh, K. Birman, and D. Dolev, "Using AVL Trees for Fault Tolerant Group Key Management," Cornell Univ., Computer Science, Technical Report 2000-45; Hebrew Univ., Computer Science, Technical Report 2000-1823, 2000.
- [37] K.P. Birman and R.V. Renesse, *Reliable Distributed Computing with Thelsis Toolkit*. IEEE Press, Mar. 1994.
- [38] Y. Amir, D. Dolev, S. Kramer, and D. Malki, "Transis: A Communication Sub-System for High Availability," *Proc. 22nd Int'l Symp. Fault-Tolerant Computing Systems*, pp. 76-84, 1992.

- [39] R.V. Renesse, K. Birman, and S. Maffei, "Horus: A Flexible Group Communication System," *Comm. ACM*, vol. 39, pp. 76-83, Apr. 1996.
- [40] Y. Amir, L.E. Moser, P.M. Melliar-Smith, D. Agarwal, and P. Ciarfella, "The Totem Single-Ring Ordering and Membership Protocol," *ACM Trans. Computer Systems*, vol. 13, no. 4, pp. 311-342, Nov. 1995.
- [41] B. Whetten, T. Montgomery, and S. Kaplan, "A High Performance Totally Ordered Multicast Protocol," *Proc. Theory and Practice in Distributed Systems, Int'l Workshop*, p. 938, Sept. 1994.
- [42] T. Anker, G.V. Chockler, D. Dolev, and I. Keidar, "Scalable Group Membership Services for Novel Applications," *Proc. Workshop Networks in Distributed Computing*, 1998.
- [43] I. Keidar, K. Marzullo, J. Sussman, and D. Dolev, "A Client-Server Oriented Algorithm for Virtually Synchronous Group Membership in WANs," *Proc. 20th Int'l Conf. Distributed Computing Systems*, pp. 356-365, Apr. 2000.
- [44] K.P. Kihlstrom, L.E. Moser, and P.M. Melliar-Smith, "The SecureRing Protocols for Securing Group Communication," *Proc. IEEE 31st Hawaii Int'l Conf. System Sciences*, pp. 317-326, Jan. 1998.
- [45] M.K. Reiter, "Secure Agreement Protocols: Reliable and Atomic Group Multicast in Rampart," *Proc. Second ACM Conf. Computer and Comm. Security*, pp. 68-80, Nov. 1994.
- [46] M.A. Hiltunen and R.D. Schlichting, "Adaptive Distributed and Fault-Tolerant Systems," *Int'l J. Computer Systems Science and Eng.*, vol. 11, no. 5, pp. 125-133, Sept. 1996.
- [47] M.A. Hiltunen, R.D. Schlichting, and C. Ugarte, "Enhancing Survivability of Security Services Using Redundancy," *Proc. Int'l Conf. Dependable Systems and Networks*, June 2001.
- [48] L. Gong, "Enclaves: Enabling Secure Collaboration over the Internet," *IEEE J. Selected Areas in Comm.*, vol. 15, no. 3, pp. 567-575, Apr. 1997.
- [49] P. McDaniel, A. Prakash, and P. Honeyman, "Antigone: A Flexible Framework for Secure Group Communication," *Proc. Eighth USENIX Security Symp.*, pp. 99-114, Aug. 1999.



Yair Amir received the BS (1985) and MS (1990) degrees from the Technion and the PhD degree (1995) from the Hebrew University of Jerusalem. Prior to receiving the PhD, he gained extensive experience building C3I systems. He is currently with the Department of Computer Science, The Johns Hopkins University, where he has been an assistant professor since 1995 and an associate professor since 2000. He has been a member of the program committees of the IEEE International Conference on Distributed Computing Systems (ICDCS) in 1999 and 2002, the ACM Conference on Principles of Distributed Computing (PODC) in 2001, and the IEEE International Conference on Dependable Systems and Networks (DSN) in 2001 and 2003. He is a member of the IEEE Computer Society.



Yongdae Kim received the PhD degree from the Computer Science Department at the University of Southern California (USC) in May 2002. He is an assistant professor in the Department of Computer Science and Engineering and a member of the Digital Technology Center (DTC) at the University of Minnesota, Twin Cities, Minneapolis. From 1993 to 1998, he was a research staff member at the Electronics and Telecommunication Research Institute (ETRI), Korea. From January 2001 until July 2002, he was working as a research scientist at the University of California at Irvine. His academic interests include network security and cryptography. More information about his research is available at <http://www.cs.umn.edu/~kyd>.



Cristina Nita-Rotaru received the BS and MS degrees in computer science from the Politehnica University of Bucharest, Romania, in 1995 and 1996 and the MSE and PhD degrees in computer science from The Johns Hopkins University in 2000 and 2003. She is currently an assistant professor in the Computer Science Department at Purdue University. Her research interests include secure distributed systems, network security protocols, and security aspects in wireless ad hoc networks. She is a member of the ACM and the IEEE.



John L. Schultz received the BS degree in computer science and electrical engineering in 1999, and the MSE degree in computer science in 2001 from The Johns Hopkins University. He is a cofounder and the lead engineer of an information services company. He is also a minor partner in a small software company. His research interests include distributed systems and algorithms, their security, and software infrastructure.



Jonathan Stanton received the BA degree in mathematics in 1995 from Cornell University and the MSE and PhD degrees in computer science from The Johns Hopkins University in 1998 and 2002. He is currently an assistant professor in the Computer Science Department at George Washington University. He also holds an appointment as an adjunct assistant professor in the Computer Science Department of The Johns Hopkins University. His research interests include distributed systems, secure distributed messaging, network protocols, and middleware support for clustered systems. He is a member of the ACM and the IEEE Computer Society.



Gene Tsudik received the PhD in computer science from the University of Southern California in 1991; his dissertation focused on access control in internetworks. He is a professor in the Computer Science Department at the University of California, Irvine. He has been active in the area of internetworking, network security, and applied cryptography since 1987. Before joining the University of California, Irvine in 2000, he was a project leader at IBM Research, Zurich Laboratory (1991-1996) and the USC Information Science Institute (1996-2000). Over the years, his research interests included: internetwork routing, firewalls, authentication, mobile network security, secure e-commerce, anonymity, secure group communication, digital signatures, key management, ad hoc network routing and, more recently, database privacy and secure storage. Some of his notable research contributions include: interdomain policy routing (IDPR), IBM Network Security Program (KryptoKnight), IBM Internet Keyed Payment (iKP) protocols, Peer Group Key Management (CLIQUE), and Mediated Cryptographic Services (SUCSES). Professor Tsudik has more than 80 refereed publications and seven patents. He is currently serving as associate dean of research and graduate studies in the School of Information and Computer Science at the University of California, Irvine. He is a member of the IEEE.

► **For more information on this or any other computing topic, please visit our Digital Library at <http://computer.org/publications/dlib>.**