Cipher-Newsletter of the IEEE Technical Committee on Security & Privacy, Winter 1991, pages 7-18.

REPORT ON THE COMPUTER SECURITY FOUNDATIONS WORKSHOP III

Ravi Sandhu

Department of Information Systems
and Systems Engineering
George Mason University,
Fairfax, Virginia 22030-4444
sandhu@gmuvax2.gmu.edu

INTRODUCTION

The third computer security foundations workshop was held at the Franconia Inn in Franconia, N.H. on June 12-14, 1990. This series of annual workshops was founded by Jonathan Millen in 1988. It has been held at the Franconia Inn since its inception and will again be held there in 1991. The proceedings are published by the IEEE Computer Society. A number of copies are still available at the discount rate of $15/copy from:

Dr. William Young
Computational Logic, Inc.
1717 W. 6th St, Suite 290
Austin, TX 78703

Checks should be made payable to Computer Security Foundations Workshop III.

This year's program consisted of 22 papers and 3 discussion/work sessions. A new tradition was started this year with the inauguration of a croquet tournament. Congratulations to Simon Foley, this year's champion! On the technical side the workshop broke new ground by having a session on database security for the first time. An informal evening session organized by Leonard LaPadula on Integrity also attracted considerable attention. I hope this interest will translate into papers for an integrity session next year.

The bulk of the report consists of a chronological description of the proceedings organized by sessions. I have attempted to be accurate to the best of my ability given my hastily scribbled notes, but there will inevitably be unintended omissions and inaccuracies.

1 Logic and Protocol Analysis

The workshop opened with introductory remarks from Tom Haigh of the Secure Computing Technology Corporation (SCTC), this year's General Chairman, and John McLean of the Naval Research Laboratories (NRL), this year's Program Committee Chairman. The session on Logic and Protocol Analysis was chaired by John McLean and consisted of four papers.

Glen McEwen of Queen's University presented the first paper of the workshop titled, "A Logic for Reasoning about Security," co-authored with *Janice Glasgow* and *Prakash Panangaden*. In this paper the authors bring together their previous work on knowledge, permission and obligations into a single framework by presenting a modal logic incorporating these three notions. According to the authors the advantage of using possible-worlds semantics is that abstract security properties can be specified without considering implementation details.

Pierre Bieber of ONERA-CERT in Toulouse, France presented a paper on "A Logic of Communication in a Hostile Environment." In this work Pierre describes a new logic called CKT5 which builds on the KT5 logic of knowledge and time due to Sato. CKT5 distinguishes between possessing a message and knowing the meaning of a message, as was done by Merritt and Wolper in their analysis of cryptographic protocols. Pierre contrasted his work with the belief and action based logic of Burrows, Abadi and Needham and asserted that security statements using belief cannot guarantee security.

Catherine Meadows of NRL presented a paper on "Representing Partial Knowledge in an Algebraic Security Model." In this paper Cathy extends her earlier model presented at the 1989 Oakland conference (i.e., the IEEE Symposium on Security and Privacy held annually in Oakland, California). Her original model cannot account for situations where a penetrator knows a word but not its significance or knows properties about the word but not the word itself. Cathy's proposal is to add notation to each word to describe what the penetrator knows about it. The practical motivation stems from a recent protocol published by Lomas, Gong, Saltzer and Needham in which easily guessable keys are used but only to encrypt random data.

Paul Syverson of NRL presented a paper on "Formal Semantics for Logics of Cryptographic Protocols." Paul began by noting the importance of formal semantics both for evaluating the completeness and soundness of the logic as well as for giving a basis for semantic reasoning. He said his goals in developing CPL (Cryptographic Protocol Logic) were (i) to provide distinct means for representing knowledge of an individual word and propositional knowledge about the word, and (ii) to integrate cryptographic axioms in the syntax and semantics rather than viewing them as additions. Paul noted that in CPL accessibility relations are not equivalence relations as is the case for logics based on S5. In particular in CPL the Barcan formula and its converse are both rejected whereas in S5 based logics these two formulas cannot be equivalent.

2 Models of Information Flow

The session on Models of Information Flow was chaired by Daryl McCullough of Odyssey Research Associates and consisted of three papers.

Jeremy Jacob of Oxford University in U.K. presented a paper on "Categorizing Non-interference." In this paper Jeremy gives a definition of non-interference in category theory and proves an unwinding theorem (i.e., properties of sequences of commands are implied by properties of individual commands). Jeremy went on to describe ongoing work of his regarding the definition of information-flow security in terms of partial orders on systems. He speculated about the benefit of extending this work to categories of systems.

Vijay Varadharajan of Hewlett-Packard Laboratories in U.K. presented a paper on "Petri Net Modeling of Information Flow Security Requirements." Vijay contrasted the true concurrency semantics possible with Petri Nets with the usual interleaving semantics of Hoare's CSP and Milner's CCS. He conjectured that this distinction would turn out to be critical for composition and refinement of distributed systems. He went on to describe an extended Petri net model for information flow security as a 13-tuple (compared to the standard 5-tuple Petri net). He also mentioned his ongoing work on the problem of synthesizing nets with desired safety and liveness properties.

Simon Foley of the Cranfield Institute in U.K. presented a paper on "Secure Information Flow Using Security Groups." In earlier work presented at the 1989 Oakland Conference, Simon had extended Denning's Secure Flow Model to allow for non-transitive flows. In the current paper Simon introduces a further extension which is to bind an entity to a set or group of security classes rather the traditional singleton security label. The resulting mathematical structure is called a reflexive lattice. Simon demonstrated how security policies such as confidentiality and aggregation could be expressed in this framework.

3 Information Flow in Abstract Machines

This session was chaired by Ravi Sandhu of George Mason University and consisted of four papers.

Ira Moskowitz of NRL presented a paper on "Quotient States and Probabilistic Channels." Ira presented two related ideas in his talk. Firstly he described an interpretation and restatement of McCullough's concept of restrictiveness in terms of quotient sets and fiber bundles. This provides a new perspective and intuitive basis for restrictiveness. Secondly he argued that restrictiveness does not eliminate probabilistic channels of communication. He gave a specific example of such a channel whose capacity he calculated as 919 bits/second. These channels demonstrate that restrictiveness is not sufficient and information theory techniques must also be used in security modeling.

Jonathan Millen of Mitre presented a paper titled “Hookup Security for Synchronous Machines.” Jon reviewed the concept of non-deducibility on strategies which was introduced by Wittbold and Johnson in their 1990 Oakland paper. He also reviewed the Wittbold-Johnson demonstration that a Trojan Horse can defeat a non-deducibility secure but hookup unsafe system. He then gave a generalization of the Wittbold-Johnson synchronous machine model and proved that for this model (i) non-deducibility by strategies implies Sutherland nondeducibility, and (ii) non-deducibility by strategies is hookup secure. In conclusion Jon recommended that the following three steps should be required of any useful definition of information flow. His comments on each step are given in parentheses.

1. Explain the state machine model. (Existing models are synchronous and attention should be given to asynchronous ones which are more realistic.)

2. Show that the definition implies Sutherland non-deducibility. (Sutherland's definition is the weakest one known so far. Jon was doubtful if a weaker one existed.)

3. Show that the definition is composable, i.e., it is hookup secure. (This is certainly required from an engineering viewpoint. Jon speculated that theoretically hookup security might be necessary and sufficient to eliminate Trojan Horses.)

David Rosenthal of Odyssey Research Associates presented a paper on “Security Models of Priority Buffering and Interrupt Handling.” He noted that this work was done in the context of McCullough's theory of restrictiveness but should extend easily to similar theories. He also noted that his work only deals with covert storage channels without consideration of probabilistic issues. The motivation for the work is to relax the input-total assumption of restrictive machines. David proposed the notion of input-limited restrictive for this purpose. He showed that a priority buffer can be modeled as a buffer restrictive machine which is a stronger requirement than restrictiveness. He proved that the composition of an input-limited machine and a buffer restrictive machine gives a restrictive system. In this manner the stringent requirement of input-total is required only of the buffering components in a system rather than all components. He concluded by speculating on the use of similar techniques to deal with interrupts.

Peter Ryan of the Communications Electronics Security Group in U.K. presented a paper on “A CSP Formulation of Non-Interference and Unwinding.” Peter gave a formulation of Gougen-Meseguer non-interference in CSP. This requires dispensing with the traditional distinction between inputs and outputs and focusing on events. He then gave an unwinding theorem in terms of two necessary and sufficient state transition rules. Peter noted that in his approach the quantification is over all valid traces of the system whereas other approaches tend to quantify over all input sequences, whether valid or not. Jonathan Millen questioned whether it was sensible to drop the classical input/output distinction? After all a system with all Low inputs and all High outputs is manifestly secure whereas the opposite is obviously insecure.

4 Panel on Covert Channels

Terry Vickers-Benzel of Trusted Information Systems (TIS), who chaired the panel, explained that the background was the First Workshop on Covert Channel Analysis, Sept. 19-21 1989, Los Angeles, CA. Terry gave an overview of the workshop which is omitted here since detailed minutes of the workshop have since been published in the July 1990 special issue of Cipher. The basic issue was the emergence of very high bandwidth covert channels (200,000 bits/sec) as opposed to the conventional evaluation standard of 100 bits/sec as an acceptable rate at the high end of the rating scale.

Sue Landauer of TIS gave a report on the deliberations of the Research Working Group at the workshop. This account is also contained in the above mentioned minutes of the Covert Channel workshop.

John Wray of Digital Equipment Corporation (DEC) gave a vendor's perspective on the problem. Since John was not at the Covert Channel workshop I have summarized his comments here. John's remarks came from DEC's experience with the VAX/VMM multiprocessor configuration. He said that Denning style formal analysis raised too many false alarms. On the other hand informal analysis looking at lower levels of abstraction found more real covert channels. He explained that Kemmerer's Shared Resource Matrix technique is good for storage channels but not so for timing channels. He called for research on formal analysis at the code/hardware level and on closure/audit techniques for reduction of bandwidth. He asked the community to consider whether an A1 system with a known large bandwidth covert channel was preferable to an unevaluated A1 target system? John said that once covert channels were found they are easy to exploit.

At this point discussion was opened to the audience and proceeded as follows.

John McLean. The distinction between overt and covert channels is a mistake. The point is that formal techniques are incomplete and therefore what we need to do is look for omissions.

Sue Landauer. Progress has been made at the FTLS (i.e., formal top level specification) level but that has only revealed additional problems.

Robert Morris. It does not matter what you call it! That is the key issue and John McLean has identified it.

William Young. How do we compare the security of a A1 system with a known large bandwidth covert channel versus a B1 with unknown covert channels? We really do not have enough information to decide unless we equate security to rating.

Terry Vickers-Benzel. The real question is how do we compare "trustworthiness" rather than "security?" That is the question cannot be resolved without considering the Yellow Book (National Computer Security Center, *Guidance for Applying the DoD Trusted Computer System Evaluation Criteria in Specific Environments, CSC-STD-003-85*).

John McLean. The question is which would you rather buy?

Daryl McCullough. Formal methods barely scratch the surface of the problem. Is it worthless to pursue them?

John Wray. We need to extend formal methods to lower levels of abstraction.

Daryl McCullough. Does top-level analysis help at all? Would not these channels be detected at lower levels anyway?

John Wray. Formal analysis has value beyond covert channel analysis.

Thomas Haigh. A lot depends on the level of detail of the FTLS. In our experience with LOCK designing with non-interference in mind eliminates some covert channels. Yet there always will be channels undetected at the specification level. To give a perspective on the magnitude of the problem note that the Shared Resource Matrix for LOCK developed by Hartman and Taylor of Computational Logic Inc. is huge—say 100 yards long.

John Wray. Architectural advances have led to high speed timing channels in the hardware. Perhaps we need to develop alternative new architectures which deliver performance without timing channels.

Todd Wittbold. Is the National Computer Security Center open to a covert channel user's guide (a concept which emerged at the Covert Channel workshop)?

Terry Ireland. Yes, but the process has not gone far enough yet. We might have a draft for the next workshop.

Todd Wittbold. The approach appears to be to decouple covert channels from the rest of the rating.

Robert Morris. They were never coupled.

John Wray. From a designer's perspective the key issue is how to trade off covert channel bandwidth versus performance.

The session concluded on this optimistic note.

5 Modeling New Properties

This session was chaired by John McLean and consisted of three papers.

Leonard LaPadula of Mitre presented a paper on "Formal Modeling in a Generalized Framework for Access Control." Len pointed out the name change from the earlier Unified Access Control to Generalized Framework for Access Control (GFAC).

GFAC is based on the concept of separating enforcement of access control from the decision on whether or not to grant access. The goal of the formal modeling effort is to make it possible to select security policies from some pre-evaluated set without having to re-evaluate the resulting configuration. Len described this as a very ambitious goal. It is also important to model policies other than the traditional MAC and DAC. For the moment Len has been focusing on Clark-Wilson as an example but other policies are readily available and would be considered in future work. Discussion of this example was deferred to an informal evening session. Len pointed out that one of the significant differences between formal modeling in GFAC versus traditional models is that in GFAC there is no 1-1 correspondence between operations and rules. To make an access control decision in principle the entire rule base may need to be consulted.

Ed Amoroso of AT&T Bell Laboratories presented a paper on "A Policy Model for Denial of Service." Ed noted the very sparse work which has been done in this area. He said that lack of a generally accepted policy model for denial of service has served to confuse the concepts of availability and denial of service. In his opinion availability includes many issues which are not part of security. He presented a policy model based on the notions of subject priority and object criticality. He sketched how compliance to a particular instance of this policy might be demonstrated for AT&T's B1 System V/MLS UNIX operating system.

Lee Badger of Trusted Information Systems presented a paper on "Providing a Flexible Security Override for Trusted Systems." Lee considers systems in which security controls must sometimes be overridden due to overall mission requirements. He outlined an approach based on relaxation lattices which attempts to minimize the security damage caused by an override. The idea is that a security override weakens the set of guarantees regarding information flow that a systems enforces. When the override is retracted the previous set of guarantees is restored, except for those violations which occurred during the override.

6 Modeling Work Session

The modeling work session was organized by *Sue Landauer* of Trusted Information Systems. Sue gave an overview of the Trusted Mach project. She explained the goal was to obtain a secure version of Mach but that the performance, portability, multiprocessing and other such features of Mach were to be given higher priority than security. In principle the target was for a B3 rating. Much of her presentation consisted of a description of Mach and the difficulties in identifying the subjects and objects. As a result TIS was considering a hybrid modeling approach based on information flow models as well as access control models. The audience was sympathetic to the difficulty of this effort, particularly since security was a low priority objective. Perhaps a more realistic target rating level would be appropriate in such cases.

7 Covert Channel Analysis

This section was chaired by Robert Morris of the National Computer Security Center. In his introductory comments Bob reiterated his definition of a covert channel as the passing of information in a system in a manner surprising to the authorities. The session consisted of two talks.

Jonathan Trostle of MITRE presented a talk on the "The Serial Product of Controlled Signaling Systems." Controlled signaling systems (css's) were introduced by Todd Wittbold at the 1989 Computer Security Foundations Workshop to model the environment faced by two conspiring Trojan Horse programs in a computer system. Todd had demonstrated that the maximum capacity of the serial product of two css's is greater than 1.5 times the maximum capacity of the individual css's. Jon presented a restricted version of css's called r-css's. He showed that every css can be converted to an equivalent r-css. He also established the product theorem for r-css's, i.e., $\mathrm{cap}(S_1 \times S_2) = \max\{\mathrm{cap}(S_1), \mathrm{cap}(S_2)\}$. He suggested that similar results should be feasible for other abstract machines.

Todd Wittbold of MITRE presented a talk on "Maximizing Covert Channel Capacity." Todd addressed two points during his talk. The first point was that a distributed collection of Trojan Horse pairs can leak information at high rates in a network. He said the good news was that the point-to-point (covert) channel capacity in a network can be computed given the (covert) channel capacities at the individual hosts. He said the bad news is that the capacities at the individual hosts add up in a network and that there are simple algorithms for achieving these capacities. The basic strategy is to use each host both as a covert channel and as a packet switch. He emphasized that in this work the individual covert channels are only in the hosts and not in the network links. His second point concerned the optimal exploitation of resource exhaustion channels. This work was done jointly with Doug Muder of MITRE. Todd presented a simple model of resource exhaustion channels as parameter setting channels. In this model the channel consists of a 4-tuple (N, t_+, t_-, r) where N is the number of resource units, t_+ is the time for a positive response, t_- is the time for a negative response and r is the time to reset the channel. Todd explained that simply increasing r is not as effective a strategy as might appear at first thought.

8 Covert Channel Detection

This session was chaired by Jonathan Millen. Two papers were presented and are described below. A third paper on "Information-Flow Analysis for Covert-Channel Identification in Multilevel Secure Operating Systems," authored by Virgil Gligor and Jingsha He of the University of Maryland appears in the proceedings but was not presented since the authors were unable to attend the workshop.

James Gray of the Naval Research Laboratory presented a paper on “Information Sharing in Secure Systems.” Jim said the goal of his work was to define a security property that precludes all covert channels and can be applied to source code. In doing so he explained there were three problem areas.

1. Nondeterminism which he dealt with in the sense of probabilistic behavior (rather than unknown but deterministic implementation).

2. Time which was modeled in the sense of local clock ticks (rather than real time).

3. Communication which required an accurate model of system behavior and a narrow gap between the model and source code. With these goals the only viable options were a polling mechanism or a bounded buffer without blocking. Jim had opted for the latter option in this paper.

Jim explored the secure readers-writers problem in this context also including serializability, failure atomicity and freedom from starvation as integrity and availability criteria.

John Wray of Digital Equipment Corporation presented a paper on “A Methodology for the Detection of Timing Channels.” John’s paper was based on DEC’s recent experience with the secure VAX/VMS multiprocessor system. He said there were three important points to note.

1. The fastest channels are timing channels.

2. Timing channels arise due to implementation quirks.

3. We only have informal techniques for detecting timing channels.

He noted that although the multiprocessor situation is by far the worse one, uniprocessors also have high bandwidth timing channels at the level of the cache. John’s basic thesis was that a storage channel requires only one clock (if only for synchronization) whereas a timing channel requires two clocks—a data clock and a reference clock. Information is conveyed in a timing channel by the frequency or phase modulation of one clock with respect to the other. He therefore proposed that timing channels can be identified by finding pairs of clocks in the system. He said that the major clock is the CPU itself. Other clocks are CPU asynchronous events which in practical terms means input/output. John claimed that almost all clocks, once they are identified, can be modulated. In the VAX/VMS the only unmodulatable clock was the timer quadword (i.e., the time of day).

9 Covert Channel Work Session

This work session was organized by *Jonathan Millen* of MITRE. He began by establishing the context for the specific problem that he was about to pose. The problem was motivated by John Wray's paper of the previous session. According to Millen the disk arm example of Wray's paper was a storage channel whereas Wray identified it as a timing channel with a storage exploitation. Jon stated the following motivations concerning why it is important to model timing channels.

- Distinguish between pure storage, pure timing and mixed channels.
- Settle Marv Schaefer's conjecture that any timing channel can be converted to a storage channel (and perhaps vice versa).
- Establish procedures to classify known channels.
- Assist in detection and bandwidth analysis.

Jon summarized John Wray's position as follows.

1. A clock is any series of events.
2. Storage channels require a clock for synchronization.
3. Timing channels need two clocks (whose phase difference can be modulated to transmit information).
4. It therefore follows that covert channels (storage or timing) can be detected by identifying all possible clocks.

Jon observed that a timing channel depends on the consistency of the execution times of certain activities. Therefore a timing channel is a channel that can be made to fail due to variations in the relative rates of processors. On the other hand such variations should not throw off a storage channel. Jon went on to describe two disk arm covert channels. He described the first one as a storage channel. He challenged the audience to classify the second one. Robert Morris suggested a variation of the second one thus giving us a total of three disk arm channels. The definitions were followed by somewhat heated but inconclusive discussion. The session was then formally closed and discussion continued in informal groups.

The three disk arm channels are described below. In all cases the disk is assumed to consist of tracks A, B, C, D, E listed sequentially from innermost to outermost. The scheduler discipline is assumed to be shortest seek time first. In the examples a single transmitter-receiver cycle for communicating one bit is defined. This cycle is repeated for every bit in the secret being transmitted. The transmitter algorithm is the same for all three examples.

9.1 DISK ARM I

The transmitter executes the following algorithm where x is the secret bit being transmitted.

```
if x=0 then read A else read E;
```

The receiver puts in two concurrent write requests as follows.

```
R0: write B, C;
        &
R1: write D, C;
```

If R0 terminates first the receiver records a 0 otherwise the receiver records a 1.

9.2 DISK ARM II

The transmitter executes the same algorithm as in Disk Arm I. The receiver clears two buffers BUF-B and BUF-D and initiates two concurrent read requests.

```
R0: read B into BUF-B;
          &
R1: read D into BUF-D;
```

The receiver then polls the two buffers until data is found in one of them. If data is first found in BUF-B the receiver records a 0 otherwise the receiver records a 1.

9.3 DISK ARM III

Robert Morris suggested the following variation of Disk Arm II. The receiver uses a single buffer BUF which he clears and then initiates two concurrent read requests.

```
R0: read B into BUF;
          &
R1: read D into BUF;
```

It is assumed that B and D store different data and that the receiver knows the data which is stored. The receiver waits until both opreations have terminated. At this point if the data in BUF is equal to B the receiver records a 0 otherwise the receiver records a 1.

10 Database Security

This session was chaired by Catherine Meadows and consisted of three papers. The session was notable in being the first one on Database Security in this series of workshops.

Ravi Sandhu of George Mason University (GMU) presented two papers. The first one was on "A New Polyinstantiation Integrity Constraint for Multilevel Relations" and was co-authored with *Sushil Jajodia* of GMU and *Theresa Lunt* of Stanford Research Institute (SRI). Ravi identified the common core polyinstantiation integrity (PI) properties required by both the SeaView model of SRI and the Jajodia-Sandhu 1990 Oakland model. He showed that the Oakland model accommodates a much richer class of multilevel relations than SeaView. He demonstrated that SeaView is unable to express some very simple multilevel relations of obvious practical value which are expressible in the Oakland model. This situation had led to a number of discussions among the three authors of this paper. The outcome was the Franconia model which shared the common core PI properties of SeaView and the Oakland model. The additional PI property of the Franconia model formalizes the notion of "a single tuple per access class." Ravi showed that a number of subtle issues arise in doing so, particularly when elements of a tuple are labeled with incomparable classifications.

Ravi Sandhu's second presentation was on "A Formal Framework for Single Level Decomposition of Multilevel Relations," co-authored with *Sushil Jajodia.* Trusted Database Management Systems which rely mostly on the Operating System to enforce MAC will decompose a multilevel relation into single level relations which are then stored in single level Operating System files. Ravi presented formal criteria for verifying that this decomposition is correct. He showed that the proofs given for the SeaView decomposition were incomplete since they only showed correctness at system high rather than at each element of the security lattice. Ravi noted that a complete proof for SeaView's decomposition would have failed because of an ambiguity in SeaViews's inter-instance integrity property (which has since been fixed). He showed that the decomposition given by Jajodia-Sandhu for their 1990 Oakland model satisfies the correctness criteria developed in this paper.

Bhavani Thuraisingham presented a paper on "Recursion Theoretic Properties of the Inference Problem in Database Security." Bhavani developed a recursion theory formulation of the inference problem and showed that the problem of deciding whether or not a database is insecure is undecidable. This of course is no surprise and could be demonstrated in any number of different ways. The recursion theory formulation did enable Bhavani to relate the inference problem to degrees of unsolvability in a natural way.

CONCLUSION

Planning for Next Year's Workshop

The audience was generally enthusiastic about the value of this series of workshops. Jonathan Millen noted that he had been a little apprehensive about handing over the workshop to Thomas Haigh and John McLean after running it the first two years. He said he was reassured by the success of this year's workshop.

There was general consensus about keeping the workshop at Franconia and maintaining a workshop flavor while at the same time having paper presentations. Thomas Haigh and William Young were reappointed respectively as next year's General Chairman and Publications Chairman. Ravi Sandhu takes over from John McLean as next year's Program Chairman. The Program Committee for next year has now been formed and consists of Thomas Haigh (SCTC), Dale Johnson (Mitre), Jonathan Millen (Mitre), Robert Morris (NCSC), John McLean (NRL), Cathy Meadows (NRL) and Todd Wittbold (Mitre). Next year's workshop is scheduled for June 18-20, 1991 at the Franconia Inn, Franconia, NH. Look out for the call for papers.

Covert Channel Work Session Revisited

In this final session there was a brief discussion regarding Jon Millen's Disk Arm Covert Channel. There was essentially no disagreement about how to classify these channels using Jon's definition. The discussion focussed on the issue of whether or not it was important to have such clear cut definitions and whether or not Jon's definition was a good one? No consensus was reached on these questions and discussion on this issue will surely continue.