

Towards Secure Information Sharing Models for Community Cyber Security

Ravi Sandhu

Dept. of Computer Science
Institute for Cyber Security
University of Texas at San Antonio
Email: ravi.sandhu@utsa.edu

Ram Krishnan

Dept. of Electrical and Computer Engineering
Institute for Cyber Security
University of Texas at San Antonio
Email: ram.krishnan@utsa.edu

Gregory B. White

Dept. of Computer Science
Institute for Cyber Security
University of Texas at San Antonio
Email: greg.white@utsa.edu

Abstract—In this paper, we motivate the need for new models for Secure Information Sharing (SIS) in the specific domain of community cyber security. We believe that similar models will be applicable in numerous other domains. The term community in this context refers to a county or larger city size unit with a clearly demarcated geographical boundary aligned more or less with a governance boundary. Our choice of the community domain is based on the decade long experience of the Center for Infrastructure Assurance and Security (CIAS), now part of the Institute for Cyber Security (ICS-CIAS) at the University of Texas at San Antonio. Over the past decade ICS-CIAS has conducted cyber security preparedness exercises and training at communities throughout the nation specifically dealing with communication, incident response, disaster recovery, business continuity, security awareness and similar issues. We discuss the insights gained from these frequent exercises to illustrate the limitations of prior models for SIS, such as discretionary access control, mandatory access control and role-based access control. Specifically, we argue that these traditional models, while effective in addressing the issues that they were developed for, lack the agility to dynamically configure a system to facilitate SIS scenarios such as monitoring and response during a community cyber security incident life cycle. We discuss how our current research efforts at the Institute for Cyber Security on group-centric SIS models directly address the limitations of existing models in such scenarios.

I. INTRODUCTION

The need to *share but protect* is one of the oldest and most challenging problems for trustworthy computing. Saltzer-Schroeder [21] identified the desirability and difficulty of maintaining “some control over the user of the information even after it has been released.” The ensuing three and half decades have further compounded the technical difficulties to the point where one may ask if it is even reasonable to seek solutions. The analog hole [29] wherein content is captured at the point it is rendered into human perceptible form and converted back into unprotected digital form underscores the intrinsic limits. At the same time our increasingly information-rich and information-dependent society needs to exploit *secure information sharing* (SIS) to fully benefit from the productivity, social and national security benefits of the ongoing cyber revolution.

SIS presents two major research challenges. The *containment challenge* is to ensure that protected information is accessible on the recipient’s computer only as permitted

by policy, including inability to make unprotected or less-protected copies. The latter has inherent limits such as the analog hole and covert channels. Containment requires a trusted computing base on the recipient’s machine and a mix of cryptography and access control, with the degree of assurance correlated with tamper-resistance. There is a rich literature on containment including the currently dominant TCG approach [2]. While high assurance is elusive and may remain so, there is consensus that low to medium assurance is within state-of-the-art. We assume that adequate assurance for containment is available commensurate with the application.

Our focus here is on the *policy challenge* of specifying, analyzing and enforcing SIS policies assuming adequate containment. A basic premise is that this requires new access control models that can integrate and go beyond earlier ones, have intuitive grounding and rigorous mathematical foundations, are usable by the ordinary citizen and enforceable in distributed systems. Another basic premise is that the policy challenge in specifying and analyzing the intrinsic application policy should be clearly separated from enforcement policy issues that arise due to the realities and practicalities of a distributed system. Following [13], [23], [26] we call these respectively P-layer (for application Policy) and E-layer (for Enforcement policy) concerns.

The premise of sharply separating P- and E-layers builds on the much practised policy/mechanism separation principle first articulated in HYDRA [15]. P-layer specifications express a policy that is ideal in the sense that it ignores issues such as distributed authorization state, network latency, caching, and requirements for off-line use. E-layer specifications define authorization decisions that approximate those given by the ideal policy in a manner that provides the desired application-dependent balance between resource availability and timely propagation of authorization-state changes. They also include additional entities such as trusted authorization/revocation servers which are abstracted out at the P-layer. In this paper, we focus exclusively on the P-layer.

Researchers at the Institute for Cyber Security have recently developed a novel approach called Group-centric Secure Information Sharing (g-SIS) [10], [11], [12]. In g-SIS, users and information are brought together in a group to facilitate sharing. Users gain access to group information by virtue of

membership. Likewise information is made available to members by adding it to the group. Constituting a group as the unit of SIS provides benefits similar to using roles versus individual users for permission distribution. Two useful metaphors for a g-SIS group are a subscription service and a secure meeting room. Subscription disseminates information to subscribers who participate in blogs and forums. A meeting room brings people together to share information available in the room. Within a group, various factors may affect authorization. For instance, the times at which users join and leave and at which objects are added and removed may affect user authorizations both during and after periods of group membership. For example, in the much studied secure multicast problem [20] new members joining the group cannot access content added prior to joining (backward secrecy) and members leaving the group cannot access new content thereafter (forward secrecy). The requirements of a committee meeting room could allow members access to older information once they join (no backward secrecy). These metaphors further indicate the need for multiple groups. In the simplest case we can have multiple groups that are *isolated* or *independent* in that membership in one group has no impact on what a user can do in another group, whereas with *coupled* or *connected* groups such impact can occur. A theory of g-SIS thus needs to model and enable specification of such temporal and coupling interactions.

In this paper, we discuss how g-SIS can conveniently handle dynamic information sharing scenarios arising in the domain of community cyber security. We believe that similar models will be applicable in numerous other domains. The term community in this context refers to a county or larger city size unit with a clearly demarcated geographical boundary aligned more or less with a governance boundary. Our choice of the community domain is based on the decade long experience of the Center for Infrastructure Assurance and Security (CIAS), now part of the Institute for Cyber Security (ICS-CIAS) at the University of Texas at San Antonio. Over the past decade ICS-CIAS has conducted cyber security preparedness exercises and training at communities throughout the nation specifically dealing with communication, incident response, disaster recovery, business continuity, security awareness and similar issues. We discuss the insights gained from these frequent exercises to illustrate the limitations of prior models for SIS, such as discretionary access control, mandatory access control and role-based access control. Specifically, we argue that these traditional models, while effective in addressing the issues that they were developed for, lack the agility to dynamically configure a system to facilitate SIS scenarios such as monitoring and response during a community cyber security incident life cycle.

The remainder of this paper is organized as follows. In section II, we discuss the requirements of SIS for incident monitoring and response based on scenarios grounded in community cyber security. We also discuss the limitations of prior SIS models based on these scenarios. In section III, we briefly review our roadmap for g-SIS and give our conclusions in section IV.

II. COMMUNITY CYBER SECURITY

In this section, we discuss the insights gained from the decade long experience of ICS-CIAS in cyber security exercises and training in communities all across the USA. We also develop scenarios and associated SIS requirements around the cyber incident life cycle grounded in this domain.

A. SIS Needs For Community Cyber Security

Currently most communities have no effective process established for information sharing with respect to cyber security. While most communities understand the procedure for responding to physical incidents (such as in case of a burglary, one calls 911), in case of a cyber incident (such as an attack on a critical computer network), they have little or no system for communication, incident response, business continuity and disaster recovery. More urgently, they have no methodology for sharing information for cooperative cyber incident management.

In terms of the different levels of sharing we see the needs of communities progressing roughly along the following lines:

- 1) The first step is to establish some initial contacts such as an exchange of business cards. In the event of a suspected incident, individuals could at least call others who might also be experiencing a similar issue (or in the early stages, might call to ask “are you seeing something like this?”). Advisory groups can be formed in communities that city leaders can turn to for advice on cyber security issues and incidents.
- 2) The next step is to establish some formal information sharing process. Individuals from critical sectors would be asked to provide specific information in the event of a suspected or actual incident. This information would go to a central point in the community where information from different organizations could be collected for analysis. In the early stages there might be little analysis conducted; the central point may simply make this information available to all the others.
- 3) As the process becomes more formalized, certain metrics may be established and organizations may opt to provide information related to these metrics. This might include information like the number of scans being seen or the number of unsuccessful logins being seen. It may also include items such as the amount of traffic being seen on web sites or other seemingly benign activities that might indicate an unusual level of interest in the community.
- 4) Each involved entity may push or pull information from the central point. At some level the amount of information that might be needed to paint an accurate picture of the current security status of a community could be unmanageable. This might mean that a system where certain information is always shared and other information is shared only if an impending incident is suspected.

B. Life Cycle of a Cyber Incident

Given the above-mentioned process, normal (steady state) cyber incident information sharing would consist of a large “open” group, unstructured in nature. Membership in such a group is typically self-motivated and open to any entity in the community, possibly conditioned by some level of credential such as an employee of a community government agency or a critical sector organization. In such a potentially large group, information overload is a major issue and some level of automatic analysis is required to manage and consolidate related information. Typically, information shared in such a group is not so sensitive. In the steady state, we would also have one “core” group. Membership in the core group is relatively tightly controlled. For instance, a set of local government agencies might form a core group and share domain specific information that is not available to the open group. This allows for ongoing sharing of more sensitive information.

When a cyber incident (or indications of a pending incident) is suspected based on information in the core and open groups, an incident specific group may be formed. As an example, organizations from which the core group is drawn monitor their systems/networks to determine if something “out of the ordinary” occurs. This might mean an inordinate number of scans or sweeps, or it may be an unusually high number of failed login attempts. It could also be scanning for an unusual port that normally is not scanned. An organization that sees this might convey to the core group that they are seeing this abnormal activity. Others may check and determine that they too are experiencing the same or similar activity. Those who experience it may then start exchanging additional information within the core group, essentially as an ad hoc subset of core group members interested in this specific issue. This would facilitate deeper investigation into questions such as the following, while keeping the shared information within the core group.

- Is the suspicious activity occurring on specific hosts running certain software?
- What software/hardware platforms do they use?
- Are there similarities between the environments in terms of software/hardware between the entities that are seeing the unusual traffic?

Some of this information may optionally be conveyed to the open group to cast a wider net for information gathering and analysis.

It may turn out that after some period of monitoring this specific issue that there is no indication of an attack/incident in which case monitoring and discussions may phase out. On the other hand, it may turn out that there has been an incident (e.g. an intrusion) or a potential incident, perhaps even multiple incidents, that merit continued investigation. In such cases an incident group may be formed of those who have actually been part of the incident and further, often sensitive, information may be shared as others are brought into the incident group (such as law enforcement). An alert might be sent to the

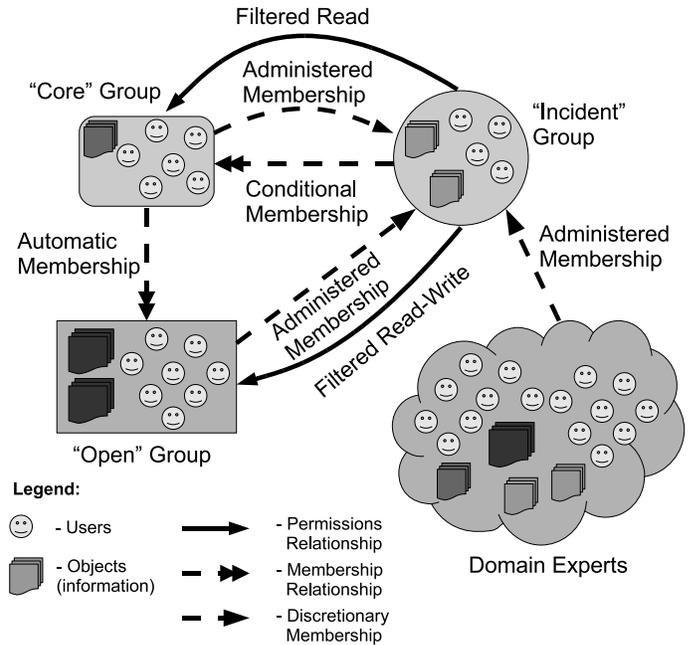


Fig. 1. Life cycle of a typical community cyber incident.

open group or even the community at large informing what to look for. As a result, others from outside the open and core groups (such as subject area experts) might be added to the small group handling the incident. Once the incident has been addressed and systems recovered (and law enforcement has gathered needed information), the incident group may be disbanded after appropriate debriefing.

C. Limitations of Classical Access Control Models

Figure 1 with “open”, “core” and “incident” groups illustrates the scenario discussed earlier. Each of these groups has different sets of users and objects (information). More likely, these groups have a few common users and objects. The dashed double-headed arrows indicate some form of membership relation between respective groups. The dashed single-headed arrows indicate administered discretionary membership from one group to the other in the direction of the arrow. The solid arrows indicate the permissions that users in one group may exercise on another in the direction of the arrow.

Users in the core group are automatically enrolled in the larger open community group by means of an automatic membership relationship as indicated. This relationship reduces administrative operations required when new members join the core since they most likely need to be abreast of information already available in the community. In response to an incident, an incident group is spun out by administering a select set of users and objects from the core group. A conditional membership relationship is established between the core and incident groups ensuring membership of such shared users in the incident group is contingent upon their membership in the core. In some cases, experts from external entities (such as local law enforcement personnel and researchers from

universities) may be called upon to participate in the incident group. This is indicated as the domain experts cloud. Since such users are not part of the core group, a filtered read relationship allows them to have restricted access to objects in the core. For instance, domain experts in the incident group may only access core group objects labeled as “port scan”. Although, such objects may be directly added to the incident group, the filtered read relationship reduces administrative burden since objects are typically created and modified at very high frequency.

Although many access control models have been published and analyzed, only three have received meaningful practical traction [27]. Discretionary access control (DAC) [6], [8], [14] enforces controls on sharing information at the discretion of the “owner” of the information but fails containment completely by allowing unprotected copies to be made. (Originator Control or ORCON [3], [9], [16], [18] attaches policies from the original to the copies to fix this defect, but does not directly address the policy challenge.) Lattice-based access control (LBAC) [4], [5], [6], [22] restricts information to flow in one direction in a lattice of security labels. Copies inherit the least upper bound of labels from the originals and remain contained. Information sharing in LBAC is essentially preordained in that information is either not shared or shared with everyone who has a sufficiently strong clearance. Any deviation from this pattern requires creation of a new label, which is not supported in existing LBAC models and breaks their existing mathematical foundations. Role-based access control (RBAC) [7], [24] is designed to facilitate assigning permissions based on job function and such considerations. Although RBAC can be configured to enforce DAC and LBAC [17] it is not designed with information sharing in mind, so it does not directly address the containment or policy challenges.

Attribute-based access control (ABAC) models such as UCON [19] and XACML [1] have the virtue of flexible policy specification by using general attributes in addition to roles and security labels. Where information sharing is static, ABAC can be configured to facilitate the process. For instance, if the structure of figure 1 is static, it is relatively straightforward to set up attributes and configure an ABAC policy to enforce the illustrated sharing scenario in the figure. However, the SIS problem is highly dynamic and it is often difficult to predict how the structure will change in the future. Below, this issue is illustrated using a minor extension to our earlier scenario.

Expanding on the life-cycle of a cyber security incident scenario in figure 1, additional incident groups may be spun out from the core group. In figure 2 incident groups g1, g2 and g3 are created to handle different aspects of a related incident. For instance, the banking sector personnel in the core may create g1, city agencies in the core create g2, etc. Each incident group focuses on a specific aspect of a similar incident in their own domain. (Groups may also be established as per a related incident such as port scans, uncharacteristic failed logins, etc.) When appropriate, an incident group may want to share certain information with other incident groups. For instance, g2 may want to allow g1 to read some of its objects for a brief period

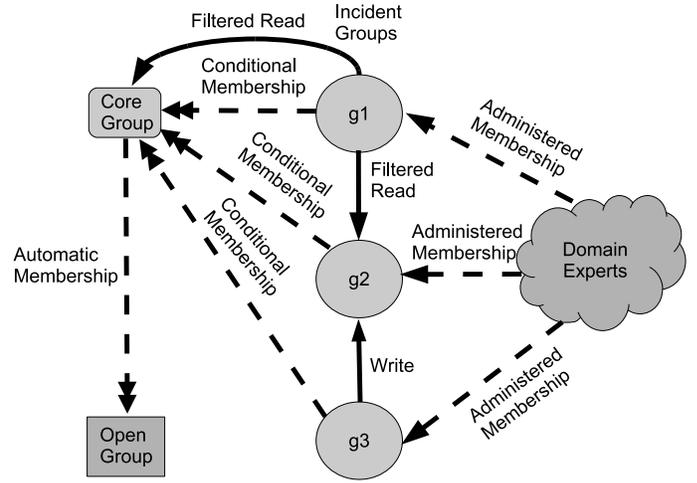


Fig. 2. Dynamic community cyber incident management.

(end of day reports in g2, for example). This is achieved by the filtered read relationship between g1 and g2. Similarly, g3 may want to share some information with g2 (we want your users to be aware of this, for example), which is achieved by the write relationship between g3 and g2. (A filtered write relationship may be desirable if information flow from g3 to g2 is a major concern.) The read and write relationships allow for “pull” and “push” models of information sharing respectively.

Clearly, user and object group membership and inter-group relationships are highly dynamic even in simple SIS scenarios as illustrated in figure 2. This bears out the claim that new models are needed for modern SIS scenarios. At the same time successful classical models such as DAC, LBAC and RBAC embody intuitions and principles that are likely to be vital to a comprehensive solution. Furthermore, new models should facilitate ease of administration to enable setting up structures dynamically with minimal user intervention. This calls for accompanying usable administrative models. Recognizing the need for new models, our current research builds upon our initial foundational results in developing new g-SIS models.

III. GROUP-CENTRIC SECURE INFORMATION SHARING

Figure 3 shows our roadmap for developing group-centric models for secure information sharing. There are two classes of g-SIS models: isolated, undifferentiated ($g\text{-SIS}^i$) and connected, undifferentiated ($g\text{-SIS}^c$). The groups are undifferentiated in the sense that attributes other than group membership have no implication on authorization in the group. For convenience, we henceforth avoid explicit mention of the term undifferentiated and simply call these two classes isolated and connected respectively. In our prior work [10], [11], [12], we focussed primarily on isolated g-SIS models. In $g\text{-SIS}^i$, groups are isolated in the sense that they do not directly interact with each other. For instance, a user’s membership in one group has no implication on her authorizations in other groups. Likewise an object’s availability to one group has no dependence or implication on availability in a different group.

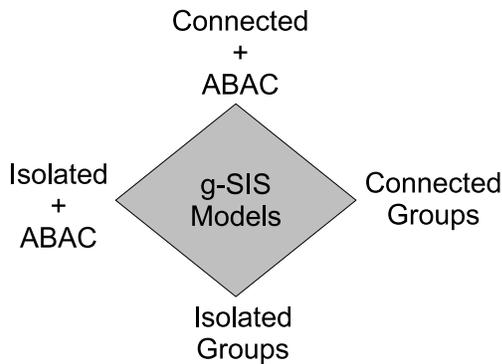


Fig. 3. g-SIS Research Roadmap.

From here, our research pursues two tracks (see figure 3). On the left, our goal is to integrate our prior work on isolated g-SIS models with attribute based access control. In practice, a realistic system would require groups combined with additional static attributes such as security clearances and roles for meaningful information sharing. Although we anticipate this extension to be straight-forward, we believe that such an isolated group + ABAC model would be of pragmatic value in many scenarios. On the right, our goal is to develop connected group models. The connected group model is concerned about developing useful inter-group relationships that can be set up and torn down dynamically such as in scenarios discussed earlier. Since relationships are temporal in nature, these models need rigorous mathematical foundations in order to understand and control information flow amongst groups. Furthermore, for membership management, effective administrative models are important. This involves analysis of existing administrative models and developing new ones if necessary. These models should facilitate easy setup of a group structure using various relationships. For instance, membership in a group could be completely discretionary in the sense that the owner of the group can select its members. On the other hand, flexible admin models may provide more sophisticated capabilities such as the ability of a user to establish a group whose members are derived from groups she currently has access to [28]. In our preliminary research [25], we have identified a number of useful inter-group relationships. Membership relationships include conditional, mutual exclusion, cardinality constraints, etc. Permission relationships include read subordination, write subordination, subject create and move subordination, etc. These two tracks culminate in a final model of connected g-SIS + ABAC which brings in attribute based constraints in addition to the policy specified by g-SIS^c. This model would incorporate important principles from prior models including owner control, one directional information flow and role based permissions management into g-SIS^c models.

IV. CONCLUSION

In this paper, we motivated new research directions in secure information sharing using community cyber security as an example domain for dynamic inter-group relationships. The

scenarios presented were based on the insights gained by the decade long experience at UTSA in cyber security exercises and training in communities spread across the USA. We have outlined a roadmap of our research on group-centric models for secure information sharing. Our final goal is to develop a unified model with connected g-SIS for handling dynamic information sharing scenarios and constraints based on static attributes for finer grained access control within groups.

Acknowledgement

The authors are partially supported by grants from AFOSR MURI, State of Texas Emerging Technology Fund and the Department of Homeland Security.

REFERENCES

- [1] OASIS eXtensible Access Control Markup Language. www.oasis-open.org/committees/xacml/.
- [2] TCG specification architecture overview. <http://www.trustedcomputinggroup.org>.
- [3] M. Abrams, J. Heaney, O. King, L. LaPadula, M. Lazear, and I. Olson. Generalized Framework for Access Control: Towards Prototyping the ORGCON Policy. *Proceedings of the 14th National Computer Security Conference*, pages 257–266, 1991.
- [4] D. Bell and L. La Padula. Secure computer systems: Unified exposition and multics interpretation.
- [5] D. Denning. A Lattice Model of Secure Information Flow. *Communications of the ACM*, 19(5):236–243, 1976.
- [6] DoD National Computer Security Center (DoD 5200.28-STD). *Trusted Computer System Evaluation Criteria*, December 1985.
- [7] D. Ferraiolo, R. Sandhu, S. Gavrilu, D. Kuhn, and R. Chandramouli. Proposed NIST standard for role-based access control. *ACM Transactions on Information and System Security (TISSEC)*, 4(3):224–274, 2001.
- [8] G. Graham and P. Denning. Protection-principles and practice. *Proceedings of the AFIPS Spring Joint Computer Conference*, 40:417–429, 1972.
- [9] R. Graubart. On the Need for a Third Form of Access Control. *Proceedings of the 12th National Computer Security Conference*, pages 296–304, 1989.
- [10] R. Krishnan, R. Sandhu, J. Niu, and W. H. Winsborough. A conceptual framework for group-centric secure information sharing. In *ASIACCS '09: Proceedings of the 4th International Symposium on Information, Computer, and Communications Security*, pages 384–387, New York, NY, USA, 2009. ACM.
- [11] R. Krishnan, R. Sandhu, J. Niu, and W. H. Winsborough. Foundations for group-centric secure information sharing models. In *SACMAT '09: Proceedings of the 14th ACM symposium on Access control models and technologies*, pages 115–124, New York, NY, USA, 2009. ACM.
- [12] R. Krishnan, R. Sandhu, J. Niu, and W. H. Winsborough. Towards a framework for group-centric secure collaboration. In *Proceedings of 5th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom)*, 2009.
- [13] R. Krishnan, R. Sandhu, and K. Ranganathan. PEI models towards scalable, usable and high-assurance information sharing. In *SACMAT '07: Proceedings of the 12th ACM symposium on Access control models and technologies*, pages 145–150, New York, NY, USA, 2007. ACM.
- [14] B. Lampson. Protection. *ACM SIGOPS Operating Systems Review*, 8(1):18–24, 1974.
- [15] R. Levin, E. Cohen, W. Corwin, F. Pollack, and W. Wulf. Policy/mechanism separation in Hydra. In *5th ACM Symposium on Operating Systems Principles*, pages 132–140, 1975.
- [16] C. McCollum, J. Messing, and L. Notargiacomo. Beyond the pale of MAC and DAC - defining new forms of access control. *Proceedings of the 1990 IEEE Symposium on Security and Privacy*, pages 190–200, 1990.
- [17] S. Osborn, R. Sandhu, and Q. Munawer. Configuring Role-Based Access Control to Enforce Mandatory and Discretionary Access Control Policies. *ACM Transactions on Information and System Security*, 3(2):85–106, 2000.

- [18] J. Park and R. Sandhu. Originator control in usage control. *Policies for Distributed Systems and Networks, 2002. Proceedings. Third International Workshop on*, pages 60–66, 2002.
- [19] J. Park and R. Sandhu. The UCON_{ABC} usage control model. *ACM Transactions on Information and System Security*, 7(1):128–174, February 2004.
- [20] S. Rafaeli and D. Hutchison. A survey of key management for secure group communication. *ACM Computing Surveys*, pages 309–329, September 2003.
- [21] J. Saltzer and M. Schroeder. The protection of information in computer systems. *Proceedings of IEEE*, 63(9):1278–1308, 1975.
- [22] R. Sandhu. Lattice-based access control models. *IEEE Computer*, 26(11):9–19, November 1993.
- [23] R. Sandhu. The PEI framework for application-centric security. In *Proceedings of 5th International Conference on Collaborative Computing: Networking, Applications and Worksharing (CollaborateCom)*, 2009.
- [24] R. Sandhu, E. Coyne, H. Feinstein, and C. Youman. Role-based access control models. *IEEE Computer*, 29(2):38–47, February 1996.
- [25] R. Sandhu, R. Krishnan, J. Niu, and W. Winsborough. Group-centric models for secure and agile information sharing. In *Computer Network Security, Communications in Computer and Information Science*. Springer, 2010.
- [26] R. Sandhu, K. Ranganathan, and X. Zhang. Secure information sharing enabled by trusted computing and PEI models. In *ASIACCS '06: Proceedings of the 2006 ACM Symposium on Information, computer and communications security*, pages 2–12, New York, NY, USA, 2006. ACM.
- [27] R. Sandhu and P. Samarati. Access control: Principles and practice. *IEEE Communications*, 32(9):40–48, 1994.
- [28] R. Sandhu and M. Share. Some owner-based schemes with dynamic groups in the schematic protection model. In *Proceedings of IEEE Symposium on Security and Privacy*, pages 61–70, Oakland, CA, April 1986.
- [29] Wikipedia. Analog hole, September 2009. [Online; accessed Dec-15-2009].