Syntropic Authentication

Suresh Chari, Pau-Chen Cheng, Larry Koved, Ian Molloy, Youngja Park
IBM T. J. Watson Research Center, P.O. Box 218, Yorktown Heights, NY 10598
{schari, pau, koved,molloyim, young_park} @ us.ibm.com

ABSTRACT

Verifying the identity of a computer user has not changed much since the earliest days of interactive computing. With limited exceptions, userid and password have dominated, occasionally supplemented with a secondary factor. We propose an alternative approach to continuous user authenticate in an interactive system based on the composition of behavioral authentication techniques.

We outline what we refer to as *syntropic authentication*, where multiple behavioral authentication techniques are composed to produce an overall score about whether the claimed identity is the user of the system. We examine continuous authentication in the context of a user using a desktop/laptop computer, although the approach can be applied to mobile, augmented reality, and other computing environments. A syntropic authentication system complements other kinds of authentication, including traditional and biometric authentication techniques.

1. INTRODUCTION

Authentication typically verifies the identity of a principal. Passwords have been long entrenched in computer authentication due to the simplicity of implementation. There is extensive literature on the many known limitations of the userid/password paradigm, notwithstanding that the user may not be present at the computer at the time of use. Knowing whether the authenticated user is present remains a challenge to the security community. We outline an approach that incorporates elements of behavioral biometrics and observable system-level behaviors to determine whether the user of the computer or computing services is in fact the claimed user *as they are using the system*. While the descriptions in this paper are centered on a user using a desktop/laptop system, we believe that the approach is generalizable to other computing devices, systems and services, including mobile computing and augmented reality.

Our approach is based on observations of *computer system windowing events*, *network footprints of user actions*, *application specific user behaviors*, and *analysis of linguistic content and patterns* augmented with the processes that generate them. Events generated as a result of user activities are used to infer models of the tasks the user is performing, the user's expertise in these applications and technologies, possible roles/responsibilities of the user, as well as personal idiosyncrasies, style and personal preferences. Our thesis is that behavioral biometrics based on these modalities, when combined with existing work on input device biometrics, will yield an accurate fingerprint of user interaction with the system. We believe that no two users will have precisely the same ensemble of attributes, and we can verify whether the claimed user is present.

Who Are You? Adventures in Authentication Workshop, Symposium on Usable Privacy and Security (SOUPS), July 9-11, 2014, Menlo Park, CA.

2. MOTIVATION

To motivate our approach, we looked at popular models from Computer-Human Interaction including GOMS (Goals, Operators, Methods and Selection rules) [1], ACT-R [2] and Activity Recognition [3-10]. A GOMS-like model can be viewed as a high level organization template. We can view the user as starting with a set of goals and interaction tasks/activities. These can result from the user's organizational role or expertise with using the applications at their fingertips and personal preferences. Starting with a high level goal, the user considers the various sequences (the *methods* and *selection* rules in the GOMS model) of elementary operators (e.g., programs) and chooses one or more of these possibilities. The choices that the user makes again reflect the same attributes such as expertise, personal idiosyncrasies, style or preferences. User activity is observed at various levels of the system. The sequences of operators that the user chooses often directly correspond to sequences of events at the operating system level (e.g., application invocation, application termination, opening a file, editing operations). Similarly, invoking certain operators within an application results in the application contacting other servers or computers often leaving behind a large footprint at the network level. Finally, once the user interacts with applications, we have artifacts such as linguistic fragments and other attributes of how various pieces of text were produced. These could include keystrokes indicative of editing operations, idiosyncratic slips and errors as well as the context, frequently misspelled words and subsequent corrections, application used, recipient of the email etc. Our approach aims to use data from the network, system and applications and to infer enough unique attributes of the user in order to perform identity verification.

A *syntropic profile* is a *cognitive fingerprint* of a user performing tasks that are consistent with their job role, background and skills. If we were to be randomly observing a user, we would see a largely undifferentiated stream of keystrokes and mouse events that would be difficult to correlate with system and network-level events. But in reality, the events are well coordinated in order to achieve higher-level goals related to their job and personal interests.

3. BEHAVIORAL BIOMETRICS FOR SYNTROPIC AUTHENTICATION

We briefly outline the approach to model each of the inputs to an overall ensemble score that verifies the user identity.

3.1 Windowing Event Sequences

Users interact with the computer windowing systems in stereotypical ways that partially distinguish them from other users. Based on roles within the organization, goals, available software tools, training, experience and expertise, a user creates a digital fingerprint of how they interact with windowing system artifacts, including window manipulation, menu access and navigation, application launch sequence, etc.

Since there is more than one possible sequence of operations to achieve a goal, the user relies on prior knowledge and preference for performing the method of starting an application. Possible operations include the user double clicking on a desktop icon, single clicking an icon on the taskbar, clicking on the Windows™ Start icon/menu, or selecting and navigating through the programs menu in the Start menu. These interactions to start the application can be done via the mouse, keyboard, touch screen, or a combination of one or more techniques. Similarly, termination of the application can be done via a variety of techniques, including the application menu bar, a combination of keystrokes, mouse interaction with the window frame / title bar, or finger swipe on a screen. With the widespread adoption of tablet and mobile computing, augmented reality, among other interaction technique, the range of interaction possibilities increases. Within this wide range of options, there are multiple variations as well.

There are several machine learning techniques ca be used to build models of users, including Hidden Markov Models (HMMs) and Hierarchical HMMs.

3.2 Generative Models of user actions

We propose to model a user's actions as digital manifestations of the cognitive goals and operations the user performs, as exhibited by the files and resources accessed, the methods and modes in which they are accessed, and the applications the user chooses to perform these operations. Users can be fingerprinted through their choice of operations to perform the underlying (sub)tasks, identifying the roles in which the user is acting in each context.

There are multiple levels of granularity at which such user fingerprinting and role identification tasks can be performed. At a coarse grained level, we can measure the applications used, and the amount of time or number and frequency of operations the user performs in each application. For example, some users will primarily edit documents in text editors while others will read PDF files, manipulate spreadsheets in a spreadsheet, or use a web browser. The choice of application to use depends on the user's current cognitive tasks, training, and expertise with the available applications.

Similarly, we can observe the documents and other resources the user accesses. These include the documents in the examples above, but also text files, databases, images and videos, remote servers, etc. These resources can often be assigned attributes explicitly, through keywords, tags, or other metadata, or implicitly through the file system hierarchy. This can be used to cluster certain resources into accounts or case files, key attributes that provide strong indications of the current user tasks.

At finer levels of granularity, we will measure *what* the users do in each application on the given resources, including the features, commands, and functions executed. In a text document, certain features, such as the ability to produce a table of contents, cross reference sections or insert keyed references, will depend on the task and skill level and training of the user. As software becomes increasingly more powerful, a user's skill level and work requirements will only necessitate a small subset of the total features, resulting in a large, sparse relation between application features, and the users that understand and leverage those features. At this level, we will measure the time and frequency the user executes each application feature.

Finally, in each application we can measure *how* the user invokes each command. Increasingly more complex software introduces more ways to leverage each application feature. Modern applications typically contain hierarchical menus, toolbars, keyboard shortcuts, contextual menus from right clicking, and context specific toolbars that appear and disappear given the current context, such as mouse position or cursor selection. We believe the methods by which each user will leverage the abilities of the applications can be used as an indication of the expertise and familiarity, as discussed in the previous section.

Role mining [11-13], generative models, partially observable Markov models and granger models appear to be the most promising approaches.

3.3 Monitoring Language Patterns

Users of computers often use language to generate emails, reports, and other textual content. Linguists have long believed that individual people have distinctive ways of writing and speaking (i.e., idiolect), and, these idiosyncratic attributes can be used to distinguish an individual from others. Recently, we witness increasing adoption of authorship attribution and forensic linguistics for intelligence, criminal investigation and plagiarism detection [14].

The state-of-the art techniques used in automatic authorship attribution and forensic linguistics rely on linguistic characteristics at every level – character [14], lexical, syntactic and stylistic [15] – and apply a classification tool to determine the author from multiple candidate authors.

We propose enhancing forensic linguistics in two directions. First, we exploit additional behavioral and contextual features as well as linguistic and stylistic features for active authentication. Second, we apply techniques to more accurately capture fine-grained knowledge on the user and the user's evolving linguistic behaviors. Specifically, we will apply multi-view learning algorithms and on-line learning approaches.

3.4 Network Activity Monitoring

A user's activities can directly initiate or indirectly trigger many network activities. We establish a *network fingerprint* of a user's interactions with other entities on the network. These entities include, but are not limited to services/applications, servers, and helper services (such as DNS). Such a fingerprint will mainly consist of statistical profiles of features extracted from network activity.

Network activities include many activities, including seemingly non-network related activities, such as editing a file, and may indirectly trigger network activities such as accessing the file on a network file/storage server, or even a cloud--based storage service. These activities usually trigger DNS queries. Network activity features can be used to build a profile on how a user interacts with the network and other entities on the network. Techniques similar to those used for G*enerative Models of user actions,* as described above, are likely relevant.

4. Summary

We presented a continuous authentication technique base on user activities. Features collected at multiple levels capture a range of knowledge, capabilities, personal preferences and idiosyncratic behavior; a user behavior model is based on these features.

References

[1] S. K. Card, A. Newell and T. P. Moran, The Psychology of Human-Computer Interaction, Hillsdale, NJ: L. Erlbaum Associates Inc., 1983.

[2] A. Newell, Unified Theories of Cognition, Cambridge, MA: Harvard University Press, 1990.

[3] J. A. Ward, P. Lukowicz and H.-W. Gellersen, "Performance metrics for activity recognition," *ACM TIST*, vol. 2, no. 1, 2011.

[4] S. McKeever, J. Ye, C. J. Bleakley and S. Dobson, "Activity recognition using temporal evidence theory," *JAISE*, vol. 2, no. 3, pp. 253-269, 2010.

[5] J. Yin, Q. Yang and Z.-N. Li, "Activity recognition via user-trace segmentation," *TOSN*, vol. 4, no. 4, 2008.

[6] D. L. Vail, M. M. Veloso and J. D. Lafferty, "Conditional random fields for activity recognition," in *Proceedings of the 6th international joint conference on Autonomous agents and multiagent systems*, Honolulu, Hawaii, 2007.

[7] J. Shen, L. Li and T. G. Dietterich, "Real-time detection of task switches of desktop users," in *Proceedings of the 20th international joint conference on Artifical intelligence*, Hyderabad, India, 2007.

[8] J. Shen, E. Fitzhenry and T. G. Dietterich, "Discovering frequent work procedures from resource connections," in *Proceedings of the 14th international conference on Intelligent user interfaces*, Sanibel Island, Florida, 2009.

[9] F. Yang and P. A. Heeman, "Context restoration in multi-tasking dialogue," in *Proceedings of the 14th international conference on Intelligent user interfaces*, Sanibel Island, Florida, 2009.

[10] O. Brdiczka, N. M. Su and J. B. Begole, "Temporal task footprinting: identifying routine tasks by their temporal patterns," in *Proceedings of the 15th international conference on Intelligent user interfaces*, Hong Kong, China, 2010.

[11] I. Molloy, Y. Park and S. Chari, "Generative Models for Access Control Policies: Applications to Role Mining Over Logs with Attribution," in *SACMAT '12: Proceedings of the 17th ACM symposium on Access control models and technologies*, Newark, NJ, 2012.

[12] I. Molloy, J. Lobo and S. Chari, "Adversaries' Holy Grail: Access Control Analytic," in *Building Analysis Datasets and Gathering Experience Returns for Security (BADGERS)*, Salzburg, Austria, 2011.

[13] I. Molloy, N. Li, T. Li, Z. Mao, Q. Wang and J. Lobo, "Evaluating Role Mining Algorithms," in *SACMAT '09:Proceedings of the 14th ACM symposium on Access control models and technologies*, Stresa, Italy, 2009.

[14] R. N. Totty, R. A. Hardcastle and J. Pearson, "Forensic linguistics: the determination of authorship from habits of style," *Journal of the Forensic Science Society*, vol. 27, no. 1, pp. 13-28, 1987.

[15] H. van Halteren, "Linguistic profiling for author recognition and verification," in *Proceedings of the 42nd Annual Meeting on Association for Computational Linguistics*, 2004.