frontiers in ICT

OPINION
published: 01 September 2015
doi: 10.3389/fict.2015.00018

Studying attacks to information systems using functional networks

Massimiliano Zanin[1,2] and David Papo[3]*

[1] Innaxis Foundation and Research Institute, Madrid, Spain, [2] Departamento de Engenharia Electrotecnica, Faculdade de Ciencias e Tecnologia, Universidade Nova de Lisboa, Lisbon, Portugal, [3] Center for Biomedical Technology, Universidad Politécnica de Madrid, Madrid, Spain

Keywords: complex networks, functional networks, security, information systems security and privacy, pattern recognition

OPEN ACCESS

Edited by:
Nicola Zannone,
Eindhoven University of Technology,
Netherlands

Reviewed by:
Federica Paci,
University of Southampton, UK
Silvio Ranise,
Fondazione Bruno Kessler, Italy

***Correspondence:**
Massimiliano Zanin
mzanin@innaxis.org

Specialty section:
This article was submitted to
Computer and Network Security,
a section of the journal Frontiers in ICT

Received: 05 July 2015
Accepted: 18 August 2015
Published: 01 September 2015

Citation:
Zanin M and Papo D (2015) Studying
attacks to information systems using
functional networks.
Front. ICT 2:18.
doi: 10.3389/fict.2015.00018

Information systems, i.e., the set of hardware and software tools that organizations use to collect and process data, are critical elements in any developed economy. Due to the intrinsic value of the information stored, government agencies and corporations alike are constantly on alert to heighten their infrastructure and data security due to threats from hackers and cyberterrorists. Information security (*InfoSec*) then aims at keeping information confidential, available, and at assuring its integrity, usually by detecting and preventing intrusions by external entities.

However, protecting large-scale information systems is an increasingly difficult challenge; not only do attackers tend to be technically more sophisticated but they also show a higher degree of collaboration amongst themselves. Both factors concur in producing attacks, which are often of a previously unknown nature. Ensuring resilience against unpredictable, potentially high impact events forces organizations into allocating huge amounts of resources.

Here, we propose that, in analogy with studies of functionally non-local complex systems, representing information systems as functional networks may help detecting and classifying patterns associated with different families of attacks, and constructing *proactive* defense systems capable of detecting intrusion of previously unknown characteristics. As a prototypical example, we compare *InfoSec* with the study of functional brain networks, a field that has received increasing attention in the last decade, and that presents similar challenges (i.e., the identification of "normal" and "pathological" conditions).

Information Systems: The Brain Analogy

Of all functionally non-local complex systems, the brain is probably the most studied. The human brain comprises an estimated 10^{11} neurons, each with an average of 10^4 connections with other neurons, performs on the order of 10^{15} synaptic operations/second, and has an estimated storage capacity of 10^{12} bytes (Sarpeshkar, 1998; Hofman, 2012). Each neuron can be thought of as a dynamical unit, and the transient coupling between these units is instrumental in many aspects of its function (Fries, 2005).

An information system consisting of large ensembles of computing units is in many ways comparable to a human brain. Structurally, like the human brain, information systems may consist of a great number of functional units (e.g., computers, routers, and firewalls), connected by cables. Information systems and the brain are also functionally comparable as both transfer, process, and store information. In both cases, these functions are associated both with activity within each functional unit and with coordinated activity between these units. Both systems can be understood as dynamical systems at various observation levels, from single unit to whole system level, whose events are represented by the information packets they generate and transfer at a given observation level.

Representing with Complex Networks

Over the last 10 years, scientists from very different fields have started resorting to complex network theory, a statistical physics understanding of graph theory (a much older branch of pure mathematics), to describe systems ranging from power grids to social interactions and the brain (Albert and Barabási, 2002; Boccaletti et al., 2006).

Representing a given system as a network involves associating nodes with the elements constituting the system, and pairwise connecting them when some explicit relation can be established (Costa et al., 2011). For instance, in transportation networks, airports can be represented by nodes, connected when a direct flight operates between them (Zanin and Lillo, 2013); likewise, connections in a social network may represent some relationship (e.g., friendship, co-working, etc.) between individuals (Vega-Redondo, 2007).

While it is straightforward to represent brain anatomy at micro as well as macro scales as a set of nodes connected through physical links, brain dynamics can also be endowed with a network-like structure, by identifying links connecting nodes with some relationship between their respective dynamics (Bullmore and Sporns, 2009). The resulting network is often termed *functional network*.

The structural, dynamical, and functional analogy with brain activity allows representing information systems as networks of interacting dynamical units, carrying out some function.

Representing Information Systems as Functional Networks

To represent *InfoSec* systems as functional networks, the dynamical variables that can be monitored and used to create functional links between nodes depend on the level at which one wishes to describe the system. From high to low levels of semantic, they may include firewall warnings, Intrusion Detection Systems alerts, software log files, down to Internet and Ethernet communications, or CPU and memory loads.

To connect pairs of nodes, two qualitatively different metrics can be used: *correlation* and *causality*. With the former, a link is established when both elements exhibit similar dynamics within the same time-window. For instance, two firewalls can be "functionally linked" when they simultaneously raise an alert; two computers, when the time series representing their CPU loads have a similar behavior at similar moments of the day. On the other hand, correlation does not imply causality; one may thus try to map causality relations between the elements. Thus, a computer may exhibit some specific behavior *as a consequence* of the activity of another unit – if the latter displayed a different activity, the former would then be driven to another dynamics.

From a practical view-point, a plethora of techniques are available for assessing correlations between time series: from the classical Pearson's linear correlation, to non-linear solutions as the Mutual Information. Furthermore, by expanding the concept of correlation to synchronization, problem-specific metrics have been developed, e.g., Synchronization Likelihood (Stam and Van Dijk, 2002) or Phase Locking Index (Pikovsky et al., 2003). Detecting causality between time series is a more complicated problem; nevertheless, metrics, like Granger Causality (Granger, 1988) or Transfer Entropy (Schreiber, 2000), are available.

Representing Attacks

The brain carries out complex tasks, such as perceptual binding or reasoning, by both processing information in largely segregated modules and integrating multiple information sources in a unified code or temporal sequence (Tononi et al., 1998). Correspondingly, the statistical mechanics approach underlying complex network theory allows conceiving of macroscopic brain function as emerging from the interactions of a vast number of microscopic neural units, and characterizing it in terms of topological properties that are essentially statistical in nature and do not directly derive from particular nodes or links. In turn, brain damage of various kinds can be represented in terms of deviations from these properties.

Functional non-locality implies that studying a single brain region may not be sufficient to detect clinically relevant deviations from normal behavior. For instance, monitoring activity at a single brain region may not contain enough information to understand epileptogenic dynamics. An alert system based on such information would typically generate an undue amount of false alarms.

Likewise, in large-scale information systems, on the one hand, activity is characterized by a high degree of spatial and temporal complexity and, on the other hand, an abnormal condition, e.g., an intrusion, may be associated with changes in global dynamics without necessarily affecting the microscopic level of its single constituent elements.

Attacks to information systems often involve simultaneous sub-attacks on different parts of the system, and may include hierarchical attack sequences, e.g., if the first sub-attack does not work, the intruder would try a second one, and so forth.

While false alarms typically occur when trying to fend off attacks based on local alarms, such as those generated by a single firewall, representing an information system as a dynamical networked system allows reconstructing the "topological structure" of a specific attack – i.e., the structure created by the interactions between the elements of the system (Boccaletti et al., 2006; Costa et al., 2007). Normal/typical activity may then be characterized by coordinated patterns of activity, captured by connectivity and the topological properties of the corresponding network. These patterns may turn out to be significantly altered during an attack, and the deviations can be used to reliably detect the presence of an attack. Nodes of the functional network would then represent alarms, pairwise connected when they co-occur in a real attack. The application of network theory to such structures can then yield valuable information about the attack. For instance, nodes centrality would quantify the importance of a specific element in the detection of the attack; communities (or modules) of highly interconnected elements would indicate the presence of coordinated (but independent) sub-attacks; and the temporal evolution of the network may yield information about how the attacker adapts to evolving security measures (Jeong et al., 2001).

Considering causal functional networks can yield even higher-level information about the attack. Specifically, *cascade effects* can be detected: the abnormal behavior of one element may force other elements to undesired dynamics. In these situations, a simple correlation may create a distorted representation of the system; events do not co-occur, but they are instead generated by a single, *root* node.

In *summary*, the *post hoc* analysis of attacks by means of complex networks can yield several benefits for the analyst; it facilitates the identification of spurious alarms, reduces the number of alarms to be monitored by detecting redundancies created by causality relations, and improves the meaningfulness of alarms by clarifying how they interact with each other, ultimately, strengthening the diagnosis.

Constructing Proactive Systems

In the previous section, we discussed the advantages of using functional networks in a reactive environment; past attacks can be studied and features of the corresponding models can then be used to recognize future attacks. However, this approach has an intrinsic limitation: pattern matching cannot be performed under unknown conditions, that is, new attacks cannot be recognized if they significantly deviate from past instances. The ultimate goal would then be to construct a *proactive* system, i.e., a system able to identify an attack, even for the first time it is encountered.

Functional networks can once again save the day. Consider again the example of the brain, and its activity associated with the execution of a given task. To understand such activity, it will be useful to reconstruct not only the network associated with this condition but also a convenient baseline. The network associated with brain activity under resting conditions, i.e., in the absence of external stimuli, is often used as a baseline to gage the properties of the network associated with task-induced activity. Provided it is sufficiently sampled, one can also use resting activity to understand the properties of specific perturbations (Papo, 2013). This is a consequence of the fluctuation-dissipation theorem (Kubo, 1966), which establishes a substantial equivalence between stimulus-evoked and spontaneous fluctuation correlations, provided the system is not driven too far from equilibrium.

Similarly, the functional network associated with the information system baseline activity may be reconstructed by identifying time-varying patterns of correlated activity found under normal conditions. New networks can then be created in real time, and compared with the resting one. Mismatches between these two networks would be recognized as abnormal connectivity patterns and could then be used to trigger an alarm. For instance, the appearance of a set of firewalls with synchronized activity may correspond to a distributed attack; likewise, the synchronized activity of a large set of computers may indicate the presence of a virus acting on them.

Problems to Overcome

The ideas sketched in the previous sections constitute a theoretical exercise discussing the possibility of a network-based approach to InfoSec, by analogy to what is now becoming standard practice in system-level neuroscience. Applying these conjectures to concrete InfoSec issues will require overcoming several barriers, two of which are discussed in what follows.

On the theoretical level, it is necessary to define problem-specific coupling measures between elementary system units. Existing metrics are either domain-general, or were designed to account for characteristics specific of brain dynamics. It is thus necessary to conceive of metrics incorporating the specific taxonomy of the information system components – e.g., detecting when two firewalls are yielding alerts due to similar causes, identifying the best time resolution in CPU load data, etc.

On the practical level, many available data sets for security systems benchmarking (Thomas et al., 2008; Nehinbe, 2011) are not suited to support the proposed developments. Data sets, such as DARPA or DEFCON, lack logs of the system activity under normal conditions, making it impossible to characterize its "resting" dynamics. The sampling rate must be fast enough to ensure that phenomena unfolding at all frequencies, as well as the relationships between these phenomena, can at least in principle be detected.

References

Albert, R., and Barabási, A. L. (2002). Statistical mechanics of complex networks. *Rev. Mod. Phys.* 74, 47. doi:10.1103/RevModPhys.74.47

Boccaletti, S., Latora, V., Moreno, Y., Chavez, M., and Hwang, D. U. (2006). Complex networks: structure and dynamics. *Phys. Rep.* 424, 175–308. doi:10.1016/j.physrep.2005.10.009

Bullmore, E., and Sporns, O. (2009). Complex brain networks: graph theoretical analysis of structural and functional systems. *Nat. Rev. Neurosci.* 10, 186–198. doi:10.1038/nrn2575

Costa, L. D. F., Oliveira, O. N. Jr., Travieso, G., Rodrigues, F. A., Villas Boas, P. R., Antiqueira, L., et al. (2011). Analyzing and modeling real-world phenomena with complex networks: a survey of applications. *Adv. Phys*. 60, 329–412. doi:10.1080/00018732.2011.572452

Costa, L. D. F., Rodrigues, F. A., Travieso, G., and Villas Boas, P. R. (2007). Characterization of complex networks: a survey of measurements. *Adv. Phys.* 56, 167–242. doi:10.1080/00018730601170527

Fries, P. (2005). A mechanism for cognitive dynamics: neural communications through neuronal coherence. *Trends Cogn. Neurosci.* 9, 474–480. doi:10.1016/j.tics.2005.08.011

Granger, C. W. (1988). Some recent development in a concept of causality. *J. Econom.* 39, 199–211. doi:10.1016/0304-4076(88)90045-0

Hofman, M. A. (2012). Design principles of the human brain: en evolutionary perspective. *Prog. Brain Res.* 195, 373–390. doi:10.1016/B978-0-444-53860-4.00018-0

Jeong, H., Mason, S. P., Barabási, A. L., and Oltvai, Z. N. (2001). Lethality and centrality in protein networks. *Nature* 411, 41–42. doi:10.1038/35075138

Kubo, R. (1966). The fluctuation-dissipation theorem. *Rep. Prog. Phys.* 29, 255–284. doi:10.1088/0034-4885/29/1/306

Nehinbe, J. O. (2011). "A critical evaluation of datasets for investigating IDSs and IPSs researches," in *2011 IEEE 10th International Conference on Cybernetic Intelligent Systems (CIS)*. London: IEEE.

Papo, D. (2013). Why should cognitive neuroscientists study the brain's resting state? *Front. Hum. Neurosci.* 7:45. doi:10.3389/fnhum.2013.00045

Pikovsky, A., Rosenblum, M., and Kurths, J. (2003). Synchronization: a universal concept in nonlinear sciences. Vol. 12, Cambridge University Press. doi:10.1119/1.1475332

Sarpeshkar, R. (1998). Analog versus digital: extrapolating from electronics to neurobiology. *Neural Comput.* 10, 1601–1638. doi:10.1162/089976698300017052

Schreiber, T. (2000). Measuring information transfer. *Phys. Rev. Lett.* 85, 461. doi:10.1103/PhysRevLett.85.461

Stam, C. J., and Van Dijk, B. W. (2002). Synchronization likelihood: an unbiased measure of generalized synchronization in multivariate data sets. *Physica D* 163, 236–251. doi:10.1016/S0167-2789(01)00386-4

Thomas, C., Sharma, V., and Balakrishnan, N. (2008). "Usefulness of darpa dataset for intrusion detection system evaluation," in *SPIE Defense and Security Symposium*. Orlando, FL.

Tononi, G., Edelman, G. M., and Sporns, O. (1998). Complexity and coherency: integrating information in the brain. *Trends Cogn. Sci.* 2, 474–484. doi:10.1016/S1364-6613(98)01259-5

Vega-Redondo, F. (2007). *Complex Social Networks*. New York, NY: Cambridge University Press.

Zanin, M., and Lillo, F. (2013). Modelling the air transport with complex networks: a short review. *Eur. Phys. J. Spec. Top*. 215, 5–21. doi:10.1140/epjst/e2013-01711-9

Conflict of Interest Statement: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.