Progress Is Infectious

Daniel E. Geer Jr. | In-Q-Tel
Daniel B. Larremore | Harvard School of Public Health

> When I began writing science fiction in the middle '60s, it seemed very easy to find ideas that took decades to percolate into the cultural consciousness; now the lead time seems more like 18 months. —Vernor Vinge

Communicable diseases spread among contacts. So do ideas, rumors, fashions, and opinions, all without central orchestration. Decentralized and infectious phenomena can't be easily extinguished by using centralized and formalized means. Although the phenomena listed above use a network's structure to spread efficiently, our centralized approaches to cybersecurity do not.

We need efficient and rapid immunization of networks against an attack in progress. We're headed in the other direction; Kelly Ziegler estimates that patching a fully deployed smart grid would take an entire year to complete (tinyurl.com/9cgl7so). For that and other critical infrastructure, the stakes are too high to ignore, too high to just "try harder." The literature of network science and statistical physics might hold some interesting solutions.

Immunization of networks usually asks the following question: How many nodes must be immunized before any small initial infection is prevented from growing to infect the whole network? Assuming the immunization process for a single node is expensive, difficult, or time intensive, the goal becomes to specify an immunization program that damps out an epidemic at minimum immunization cost. In other words, who gets the vaccine first? Here are three strategies, differing in approach depending on how much is known about the transmission network. In each case, immunizing a node makes all links to and from that node useless to a virus. The diagrams provided are meant to convey the general idea of each method, albeit imprecisely. We start with a simple network in which all hosts are susceptible:

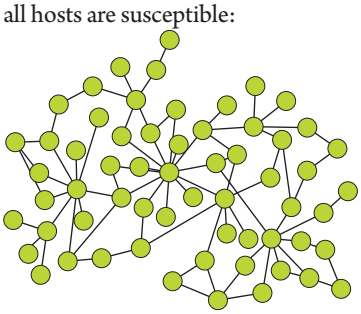

If the topology of the network to be immunized is completely unknown, then the "acquaintance immunization" technique of Reuven Cohen, Shlomo Havlin, and Daniel ben-Avraham statistically decreases the number of immunizations required to achieve collective immunity by iteratively choosing a node at random (red) and immunizing one of its network neighbors instead (blue), not the randomly chosen node itself (tinyurl.com/7l6ec4o). Why? Because a random acquaintance is likely to be better connected than the node itself:

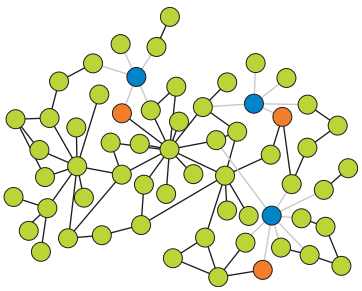

If the network's structure is known but not all links in the network are of equal strength (some are of higher bandwidth than others), then the dynamical importance measure suggested by Juan Restrepo, Edward Ott, and Brian Hunt provides criteria by which nodes that are most important to the dynamics of the infection (blue) are singled out for priority immunization (tinyurl.com/bmzg765). This approach is similar to the search engine page rank calculations and thus might be somewhat straightforward to implement.

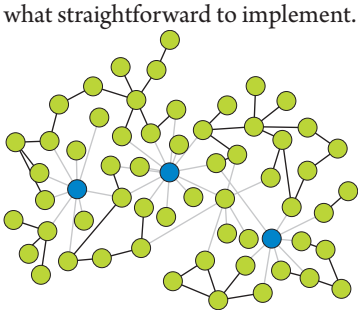

If the network's structure is exactly known and all links are of equal strength, then Jeremy Hadidjojo and Siew Ann Cheong suggest that the real goal is to split the network in two

 Copublished by the IEEE Computer and Reliability Societies

by removing nodes: find the set of nodes that creates the optimal partition of the network (blue) and then cut along the dotted line, repeating until the network fragments available to the virus can be individually secured in an outbreak (tinyurl.com/9hzbfa5). It's an explicit divide-and-conquer approach that does the dividing up front so the conquering is easier later. This is a clever idea:

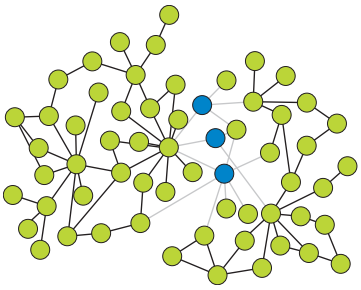

But all three techniques are still orchestrated and planned in a methodological way, and they assume at the outset that all nodes are accessible for immunization. It would be better if a patch could be loosed upon the network to spread as the network itself dictates. The concept of a benevolent virus isn't new, of course—nature got there first. For instance, rosy apple aphids infected with densovirus produce many more winged morphs than their uninfected counterparts, essential for aphid dispersal and survival and convenient for the virus, too (tinyurl.com/8fvds5h). In another case, squash plants infected with a nonfatal zucchini yellow mosaic virus become unattractive to beetles that carry a fatal bacterial wilt disease (tinyurl.com/8tkcoow).

If a patch were embedded in an infectious but otherwise benign vector, it would have to compete against the malware that it counters. Brian Karrer and Mark Newman analyzed a system of two competing epidemics using a susceptible-infected-recovered model from epidemiology, in which infection and recovery from one epidemic granted immunity to the other and vice versa (tinyurl.com/cxv8xw2). The researchers found that even in very large networks, the first epidemic to infect a large proportion of the population eclipsed the other, even if that first epidemic was less infectious. This suggests that a patch traversing the network and removing a vulnerability must be released as soon as possible to prevent the virus from gaining a significant foothold.

Sound outlandish? After exploiting a vulnerability, botnet malware will often patch multiple vulnerabilities, including the one that allowed its entrance, effectively excluding the competition. As has been suggested by Chris Wysopal, botnet malware that uses CPU cycles only during computer idle time and excludes other botnets might thereby be beneficial to the overall network's health. If the police are slow, paying the cost of mafia protection will keep the more violent gangs away.

Benevolent vigilante worms have been written,[1,2] and although condemned by some observers,[3] time marches on. The more "computers" there are that are networked, autonomous, and not under the observant control of a human of whom permission might be politely asked, the more important it is to give renewed thought to ideas like these. At least since NIMDA, malware has exploited multiple attack vectors, yet cybersecurity still relies on fully centralized patch delivery systems. The bad guys get in through the windows and air vents, but the police have to walk up the driveway and knock politely at the door. This might be particularly apropos when tackling the problem of pirated copies of Windows, which are common, don't get security updates, and yet remain part of the general Internet. The piracy is Microsoft's problem, but ubiquitous and vulnerable machines breed systemic problems, just as having more than 50 percent of Apple's installed base no longer getting security updates or the well over 500 versions of the Android OS.

It would be interesting to see Karrer and Newman's work extended with a vector-borne patch and cure as one epidemic and malware as the other. Assuming any malware has a first-mover advantage—you can't encode the cure before the disease—could clearance times be improved for even slow-moving cures?

The safety of a watchful neighborhood beats the occasional patrol car. Its mechanism lies outside formality, yet is highly effective. The mechanisms of cybersecurity could be, too. Models in network science and physics, as well as approaches from public health and epidemiology, can and should inspire developments in cybersecurity, but could also inspire nefarious players. It would be wise to explore this in future research sooner rather than later; humans managing computers they own is becoming a quaint anachronism, however sad that might be. ■

References

1. K. Kleiner, "Viral Cure Could 'Immunise' the Internet," *New Scientist*, 1 Dec. 2005; tinyurl.com/d9t75e8.
2. P. Roberts, "New Variant of Blaster Worm 'Fixes' Infected Systems," *Computer World*, 18 Aug. 2003; tinyurl.com/8zex7cq.
3. B. Schneier, "Benevolent Worms," 5 Dec. 2005; tinyurl.com/9ssjwfz.

Daniel E. Geer Jr. is the chief information security officer for In-Q-Tel. He was formerly vice president and chief scientist at Verdasys, and is a past president of Usenix. Contact him at dan@geer.org.

Daniel B. Larremore is a postdoctoral research fellow at the Harvard School of Public Health. Contact him at larremor@hsph.harvard.edu.