Summary – Flash-DDoS Defense via Human Behavior Modeling

Application-level, aka "flash-DDoS" attacks are the most challenging form of distributed denial of service (DDoS). They flood the victim with legitimate-like service requests generated from numerous bots. There is no defense today that is even remotely effective against flash-DDoS attacks. Traditional DDoS defenses that look for malicious content fail because flash-crowd attacks generate legitimate service requests. Defenses that detect spoofing also fail because flash-crowd attacks do not usually spoof their traffic. Defenses that look for aggressive senders may fail because the size of today's botnets enables attackers to spread the attack power among many bots, with each sending at a low rate. Defenses that model traffic feature prevalence or sender synchronization can be easily defeated by the attacker varying their attack content and sending dynamics. A few defenses have been proposed that detect flash-crowd attacks, but none can separate attack from legitimate traffic. Thus, flash-crowd DDoS attacks are today a serious and unmitigated threat to any server.

Intellectual Merit: Our key insight is that most service requests today are generated by humans, and that humans have unique ways they interact with and process content that bots cannot easily replicate. We propose to develop novel technologies called ASTUTE (pASsive TUring TEsts) to distinguish bots from human users, by modeling three aspects of human user behavior: (1) dynamics of human-server interaction, (2) human preference for server content, and (3) human processing of visual and textual cues. IP addresses of detected bots will be blacklisted and their traffic will be dropped during server overload.

ASTUTE technologies model human behavior without conscious human participation, thus performing Turing tests (human vs machine differentiation) transparently to humans. Because of this transparency they can be triggered upon each user request, continuously testing for bot presence. This ensures that bots who fly under the radar in one test may be detected upon their very next request. ASTUTE technologies are orthogonal to each other and can be combined for synergistic defense. They are superior to current DDoS defenses, because they can defend against stealthy flash-DDoS attacks that bypass current defenses. Another common approach to verify a human presence is an active Turing test using graphical puzzles that humans can easily recognize but machines cannot. These tests fail for flash-DDoS defense for a number of reasons explained in the proposal.

In our preliminary work, we have developed simple prototypes of three proposed ASTUTE technologies, and we have evaluated them using theoretical analysis and Web log analysis. Our findings show that we can detect bots with extremely low false positive and false negative rates ($<$1%) thus forcing attackers to use very large botnets that are beyond reach of most of them.

In the proposed work we seek to improve our technologies to greatly increase their sophistication, thus increasing sensitivity and accuracy, and raising the bar for the attacker even higher. We also seek to combine these technologies into the MAXIM system (Multi-AXes Intelligence Models), and exploit their synergy. We believe that this combined system would make any sustained flash-DDoS attack impossible, because the required botnet size would exceed the number of computers that exist today. We further propose to investigate the applicability of our technologies to protection of popular, non-Web services, such as DNS, where some request traffic is machine-generated but the underlying process driving this generation is originated by humans. Finally, we propose to investigate the applicability of our technologies to non-DDoS security threats where it is important to differentiate human-generated vs bot-generated traffic, such as click-fraud.

Broader Impact: Success in this project would be a game-changer in the application-level DDoS field. Virtually any server today is vulnerable to flash-DDoS attacks and there are no effective defenses that it could deploy. Our ASTUTE technologies will provide effective and efficient defenses to any Web server, thus protecting the most popular and one of key services in the Internet. If we are successful in applying our techniques to protection of other critical services, such as DNS, this would greatly fortify Internet security against one of its most vicious threats.

We will implement all our code as extensions of popular open-source server platforms, such as Apache (for Web) and bind (for DNS). At the end of this work we will deliver working prototypes of these extensions, thus our research will directly transition into practice for any interested party at no cost to them. All our code will be released as open-source under the GNU GPL v3 license.

The PI will actively seek to involve undergraduates and students from underrepresented groups in research and educational activities within this project.