# Countering click spam: Researchers test new algorithm to detect, combat fraudulent clicks online

November 4 2013

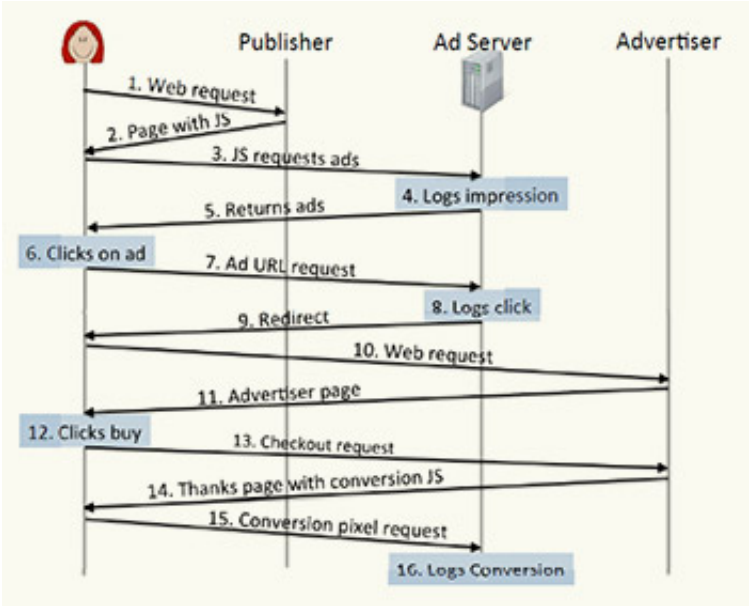


Anatomy of a click

(Phys.org) —When is a click not a click? When an advertising network registers a click on one of their online advertisements, how can it be sure that a single consumer – a "pair of eyeballs" in Madison Avenue jargon – and not a malware computer program, is behind that one click? Or that the viewer's click was intentional, not induced by deceptive or misleading advertising?

Click-spam has become a little-known way of life on the Internet. Little known, compared to other types of spam, because much of the fraud is targeted at the advertising networks, rather than at consumers directly. So what happens when an automated system can "click" on hundreds of ads in less than a second?

Given that U.S. online advertising tops $36 billion a year, even a small fraction of fraudulent clicks on Web advertisements adds up.

"Hundreds of millions of dollars are siphoned off in ad revenues based on illicit click-spam schemes," said University of California, San Diego computer science and engineering postdoctoral researcher Vacha Dave. "We knew click-spam was out there, but the hard part was how to prove the fraud scientifically. So we came up with an approach based on what the most frequent scams have in common."

"Vacha has become a real expert on click fraud in Web advertising," said CSE professor Geoffrey Voelker, who is affiliated with Calit2's Qualcomm Institute. "The approach she designed was recently put in place by a major ad network and has had an immediate major impact."

In a paper to be presented November 7 in the "Web Attacks" session of the ACM Conference on Computer and Communications Security in Berlin, Germany, Dave (pronounced dah-veh) will spell out a new approach to fighting click-spam. She and her co-authors Yin Zhang at the University of Texas at Austin and Saikat Guha at Microsoft Research India came up with a catchy name for the algorithm they created to catch click-spam in search ad networks. They call it ViceROI, and it's designed to be deployed at the ad network where it has visibility into all ad clicks.

"We designed ViceROI based on the intuition that click-spam is a profit-making business that needs to deliver higher return on investment – ROI

– for click-spammers than other ethical business models in order to offset the downside risk of getting caught," said the researcher. "Click-spam publishers should therefore have inordinately high return on investment." Figuring out actual ROI can be difficult because ad networks jealously guard their data, so the researchers employed revenue-per-user estimates as a close proxy for ROI.

Dave was in a unique position with her colleagues to test what they call the "simple-but-general ViceROI approach." They were given access to real-world data from a large ad network.

Until now, the UC San Diego researcher said, ad networks typically responded to click-spam reactively. They would react after an advertiser complained about being billed excessively because of click-spam, e.g., if it was getting thousands of clicks from the same IP address and none of the clicks led to paid transactions. The ad network would simply block or filter clicks from that IP address going forward. But the lack of transparency has often led to click-spam not being uncovered for years at a time (in one case cited by the paper, four years and $14 million in fraudulent clicks, uncovered after the fact). In other cases, an offender could get around a filter by using a distributed botnet to make the same number of clicks, but all from different IP addresses to avoid detection.

According to Dave, she tested the ViceROI algorithm on a real-world dataset. Out of the tens of thousands of publishers in the data set, ViceROI flagged several hundred which apparently were resorting to click-spam of various sorts. The researchers then investigated manually roughly one-hundred of the flagged publishers. The result: they found at least three spam publishers in each of six types of click-spam: conversion fraud; ad injection; search hijacking; malware; arbitrage; and parked domains.

"The ViceROI approach flags click-spam through all these mechanisms

and… is resilient against click-spammers using larger botnets over time," reported the paper's authors, adding that their approach "ranked among the best existing filters deployed by the ad-network today while being far more general."

As part of their research, the team also placed so-called "bluff" ads. Bluffs are nonsensical and therefore highly unlikely to be clicked on by consumers. So if the ad started attracting clicks, the assumption was that they were most likely coming from click-spammers – allowing the researchers to assess the accuracy of their algorithm by using bluff ads as benchmarks for comparison.

Another challenge, according to Dave, is tracing where fraudulent clicks come from. "Botnets and botmasters make it very difficult to be certain about the source of click-spam," said the UCSD researcher. "Even ad networks are reluctant to talk openly about what's being done to combat fraud in this area, because it will inevitably lead spammers to find new ways around new technologies put in place at the ad-network level."

The Conference on Computer and Communications Security (CCS) runs Nov. 4-8. It's the flagship annual conference of ACM's Special Interest Group on Security, Audit and Control (SIGSAC). Information security researchers, practitioners, developers and users worldwide attend CCS to explore cutting-edge ideas and results.

**More information:** www.sigsac.org/ccs/CCS2013/index.html

Provided by University of California - San Diego

https://phys.org/news/2013-11-countering-click-spam-algorithm-combat.html