

Helping TicketMaster

Changing the Economics of Ticket Robots
with Geographic Proof-of-Work

Ed Kaiser

Wu-chang Feng



Portland State
UNIVERSITY

The Ticketing Problem

- Popular event tickets are sold **below market rate**
 - maximize attendance
 - wider access to event
- Scalpers employ **automation** to buy many tickets
 - for resale at greatly inflated prices



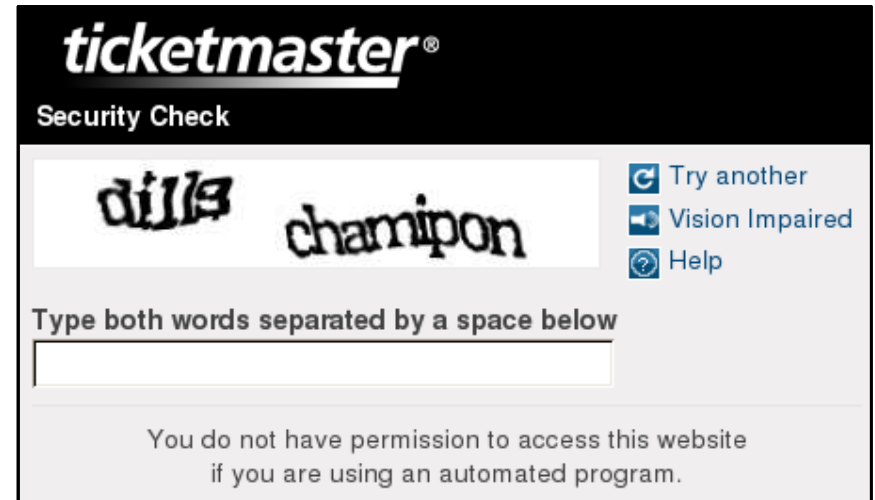
“tickets sold out in 5 minutes”

tour has been a major moneymaker for the online ticket exchange StubHub. represented among sellers on StubHub and other online ticket exchanges, have been so successful in snagging Hannah Montana tickets in plentiful quantities.

Ticket sales for big-name concerts now follow a distressingly consistent pattern: At 10 a.m. on a Saturday, tickets go on sale, and by 10:05 a.m., all tickets are sold. Yet by 10:05, StubHub and other ticket exchanges already have a plenitude of tickets listed for the sold-out event — only now, they cost much more.


CAPTCHAs are Current Solution

- Vendors use **CAPTCHAs to ward off ticket scalpers**
 - thwarts completely automated approaches
- Yet scalpers adopt hybrid **botnet–human approaches to circumvent them**
 - outsourcing costs less than 1¢ per solution
 - while profits are on order of dollars



Proof-of-Work Approaches

- **Computational puzzle for client machines**
 - difficulty set **individually** per-client

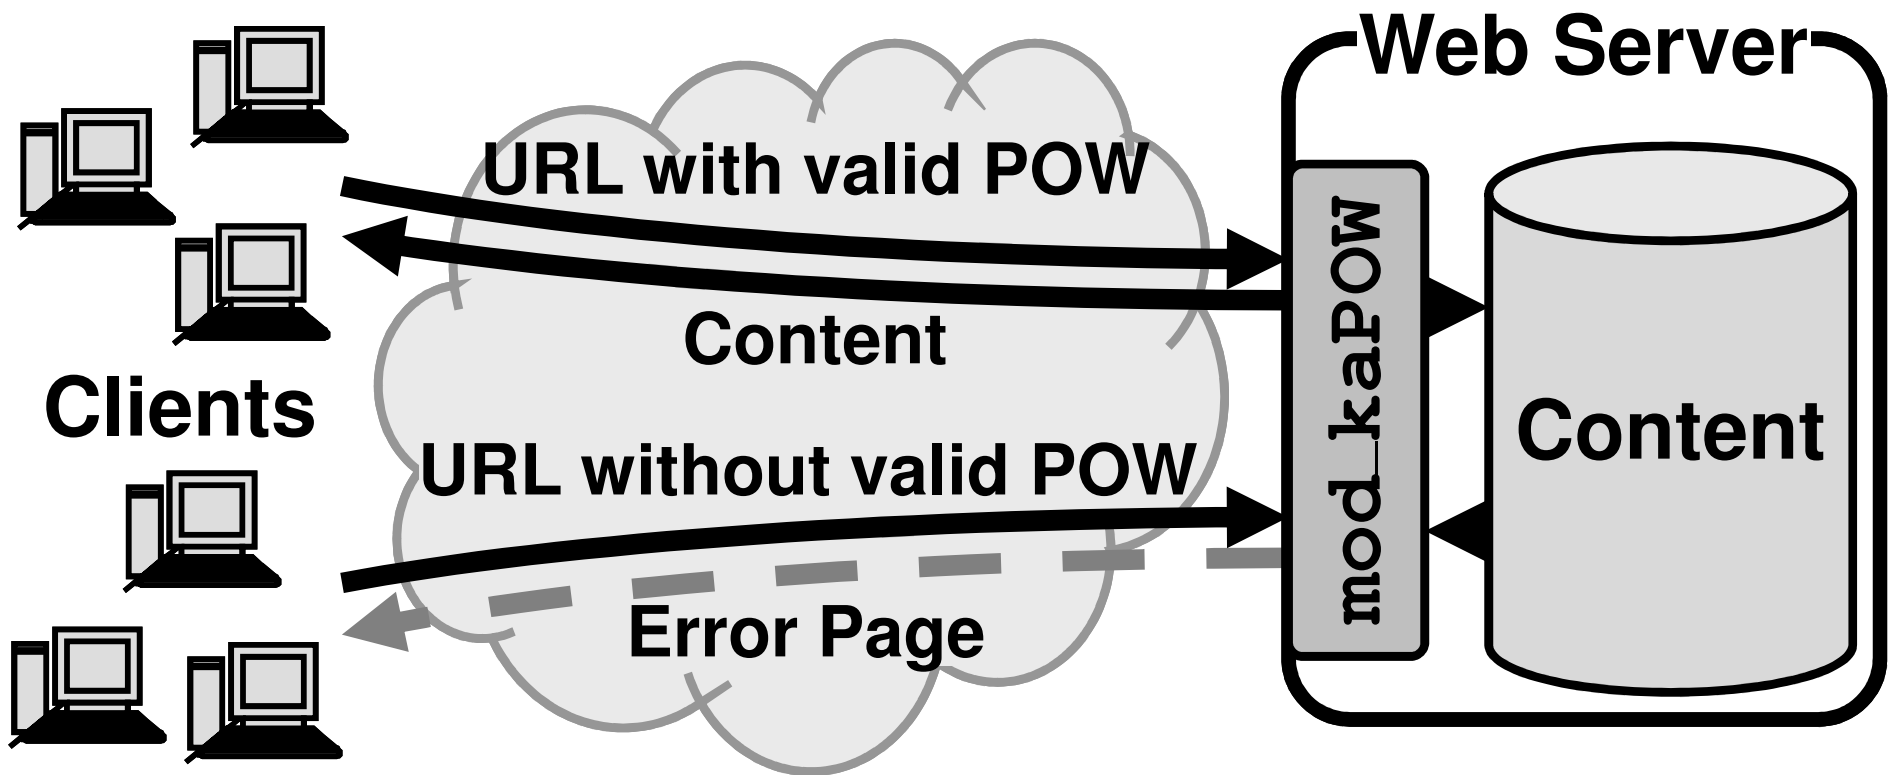


- **Unobtrusive**, not prohibitive to users
- **Prevents concurrent attacks** by adversaries who target multiple events
 - altruistic yet can also be self-serving

The mod_kaPOW Approach [GI08]

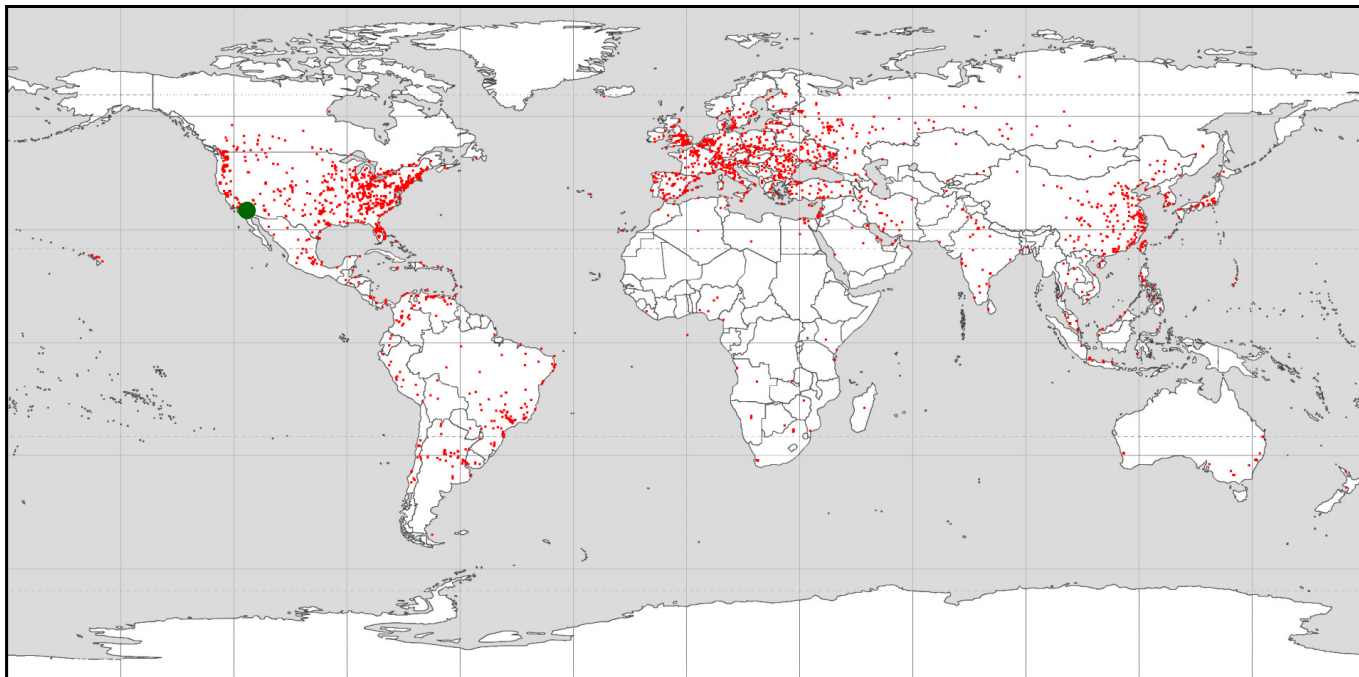
- **Transparent Proof-of-Work**

- leveraging ubiquitous JavaScript (96% of clients)
- appends answers as URL query data
- load-based difficulty



Hypothesis

- **Geographically-driven puzzle difficulty can make ticket scalping more expensive**
 - legitimate clients are often local
 - while botnets are geographically scattered



Geographic Databases


- **Reasonable overhead**
 - cached locally, periodically updated
 - not directly targetable by attack
- **Becoming very accurate**
 - continually improving
 - most misclassifications are not drastically wrong
 - augment with other geographic information like credit card billing addresses
- **Applicable to other localized services**
 - e.g., local online newspapers

Disincentivizing Ticket Botnets

- **Botnet may be compromised machines**
 - without physical owner's knowledge
 - bot herders do not want to be discovered
 - rent by the hour / geography
 - naïve approach rents local and distant bots
 - excessive computation may be detected leading to repair
 - more complex approach rents only local bots
 - require many more machines across all event centers
 - increased chance that local bots will conflict with physical owner leading to detection and repair
- **Botnet may be owned machines**
 - more secure but much more costly

PHP Prototype

- Requires **fewer modifications** to adopt
 - no module to be compiled into Apache
 - change the content rather than the infrastructure
 - `http://kapow.cs.pdx.edu/geotickets/`
- The question we answer:
How to best set the per-client difficulty based on distance?



Evaluation

- **Built a simulator in C++**
 - validated against prototype (100 clients, 100 adversaries)
- **Measured access vs. distance**
 - with large adversary populations (250,000 adversaries)
 - across 25 largest US metropolises
- **Tested adversary geographic distribution**
 - fixed populations (2,500 clients, 2,500 adversaries)
 - determined domestic distribution optimal
- **Explored difficulty as function of distance**

Simulator Assumptions

- **Clients & adversaries buy any and all tickets allowed per transaction**
 - tickets are then sold in uniformly sized lots
 - fungible (i.e., “general admission”)
- **Diehard fans F5’ing ticket sale webpage**
 - legitimate clients are present at nearly the moment tickets go on sale
- **Legitimate demand sufficient to sell out**
 - scalpers focus on event tickets that are likely to resell at substantial markup (lose \$ on tickets not resold)
 - $\#_{\text{tickets}} = \#_{\text{clients}}$, so adversary success \rightarrow client loss


Theoretical Bounds

- **Bounds on adversary success rate**
 - as percentage of total tickets acquired

$$\textit{without geographic defense} \approx \frac{|A|}{|A| + |C|}$$

$$\textit{optimal geographic defense} \approx \frac{|A_{local}|}{|A_{local}| + |C|}$$

Difficulty Algorithms



Conclusions

- **Scalpers automate ticket acquisition**
- **CAPTCHA defenses are inadequate**
- **POW approaches are advantageous**
 - puzzle difficulty can be individually set per-client

Evaluated geographic distance as metric for setting puzzle difficulty

- can be applied to other problems