Privacy Implications of Smart TVs

J. Varmarken[1], H. Le[1], J. Al Aaraj[1], A. Shuba[1], R. Trimananda[1], A. Markopoulou[1], Z. Shafiq[2]
[1]University of California, Irvine, [2]University of California, Davis

https://athinagroup.eng.uci.edu/projects/smarttv/

Advertising and Tracking [PETS'20]

Motivation

- At least one smart TV in 82% of US TV households → large audience for (targeted) ads.
- Smart TVs contact more third parties than other IoT devices [Ren et al., IMC'19] → indicative of tracking?
- Unique opportunity for advertisers: Smart TVs can observe users' viewing habits → use this information to target ads
- Ads and tracking on web and mobile platforms studied extensively; little work on smart TVs.

Research Questions

RQ1: What does the smart TV advertising and tracking services (ATS) ecosystem look like?

Network Measurement Approach

- Smart TVs in the wild
  - Traffic from smart TVs in 41 US homes.
  - 57 smart TVs spanning 7 different smart TV platforms; 3 weeks of traffic.
- Smart TVs in a testbed
  - Automated interaction with Fire TV and Roku apps while recording traffic.
  - Tested top-1000 apps of each platform.

RQ2: How effective are existing privacy-enhancing tools?

Analyze DNS-Based Blocklists

4 sets of blocklists
Pi-hole Default (PD) The Firebog (TF)
Mother of all Ad-Blocking (MoaAB)
StopAd (SATV)

- Why blocklists: Universally applicable across all smart TV platforms
- Also used for labeling domains in our datasets as ATS/non-ATS

Results Highlights (Testbed Dataset)

ATS Domains Contacted

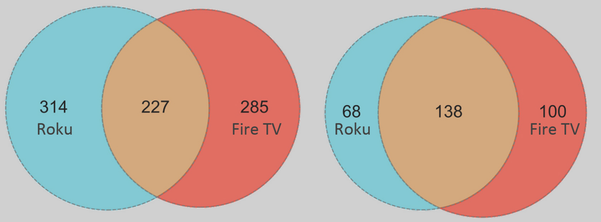

Fully Qualified Domain Names (FQDNs) Effective Second-level Domains (eSLDs)

Takeaway: Large portion of ATS domains are unique to Roku and Fire TV, respectively

Behavior of the Same App on Roku vs. on Fire TV

- Apps identified by fuzzy matching on app name (due to slight variations in app names, e.g., "TechSmart.tv" on Roku vs. "TechSmart" on Fire TV).
- Analyzed FQDNs contacted by the same app on the two platforms: portion of domains exclusive to the Roku/Fire TV versions of the app, and domains used by both versions of the app ("common" in the diagram below).

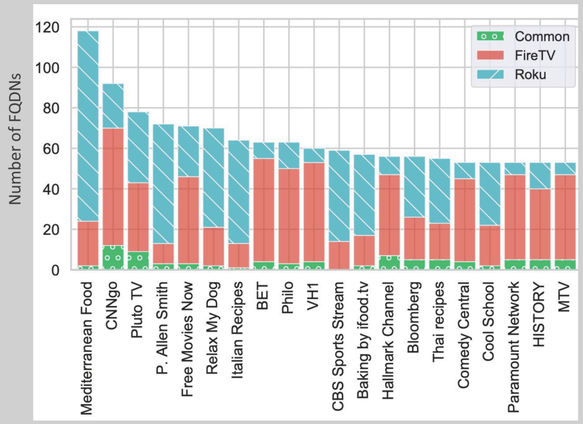

(only a subset of apps shown)

Takeaway: only minor overlap in domains accessed by Roku/Fire TV versions of the same app

Blocklists Ability to Prevent Exposure of PII

- Search packet payloads for personally identifiable information → determine if blocklists would have blocked the request ("block rate": percentage of FQDNs that received the respective PII that would have been blocked by the union of the four blocklists).
- Exposures categorized by receiving party relative to the app's developer.

	PII	1st Party		3rd Party		Platform Party	
		Apps	Block Rate	Apps	Block Rate	Apps	Block Rate
Roku	Ad ID	4	25%	263	88%	0	-
Roku	Serial Num.	48	5%	128	74%	0	-
Roku	Device ID	N/A	-	N/A	-	N/A	-
fire tv	Ad ID	17	25%	53	78%	715	71%
fire tv	Serial Num.	10	0%	51	33%	867	9%
fire tv	Device ID	19	8%	153	36%	819	14%

Takeaway: 100s of apps expose PII. The blocklists do reasonably well at preventing this on Roku but struggle on Fire TV. Joint exposures of static (serial number) and dynamic (ad ID) PII eliminate users' ability to opt out of targeted advertising by resetting their ad ID.

Network Fingerprints of Smart TV Apps [PETS'22]

Motivation

- Smartphone app usage is indicative of the user's demographics, personality, interests etc. [Zhao et al., PMC volume 59].
- Assuming this carries over to smart TVs, and considering that viewing history is regarded a cornerstone of programmatic TV advertising [Malthouse et al. IJA volume 37], smart TV app usage data is arguably a treasure trove of information for businesses engaged in targeted advertising.
- Because ISPs are known to collect and use information about their customers for advertising purposes, it is important to quantify to what extent they can collect smart TV app usage data as well.

Research Question

To what extent can smart TV apps be identified from their network traffic?

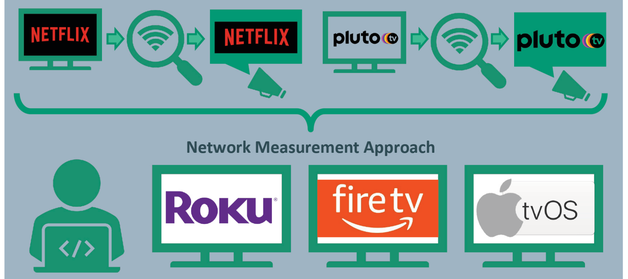

Network Measurement Approach

Use software instrumentation to repeatedly launch the top-1000 apps of the 3 smart TV platforms with the largest market shares while recording network traffic. Then examine the resulting traffic traces for network fingerprints.

Results Highlights

Prevalence and Distinctiveness

	DBF		PBF		TBF	
	Pre.	Dis.	Pre.	Dis.	Pre.	Dis.
Apple TV	96%	59%	68%	77%	95%	3%
Fire TV	88%	63%	95%	88%	86%	7%
Roku	100%	46%	100%	72%	100%	1%

Takeaway: All 3 fingerprint types prevalent on all platforms, but only DBFs and PBFs can reliably identify an app among other apps.

Sizes

	DBF		PBF	
	Median	Distinct > Median	Median	Distinct > Median
Apple TV	4	64%	2	99%
Fire TV	4	80%	4	100%
Roku	4	55%	5	94%

Takeaway: A large fraction of the fingerprints that are larger than the median size are distinct fingerprints.

Fingerprinting Techniques Considered

Domain-based Fingerprint (DBF): Set of **domains** that are consistently contacted when app *x* is launched.

Packet-pair-based Fingerprint (PBF): Set of request-reply packet exchanges that appear *N* times in total across all *N* samples of app *x*'s on-launch traffic. Two packet exchanges are identical if the packet sizes and directions match. Adopted from [Trimananda et al., NDSS'20].

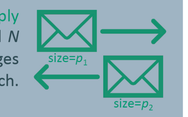

TLS-based Fingerprint (TBF): Set of **TLS fingerprints** that are consistently present in app *x*'s on-launch traffic. TLS fingerprints, due to Ristić, are based on the parameters the client sends when initiating a TLS session.

Evaluating Fingerprint Performance

Metrics Considered

Prevalence: How many apps exhibit a DBF/PBF/TBF?
Distinctiveness: How many apps exhibit a DBF/PBF/TBF that is distinct from those of all other apps (within or across platform(s))?
Sizes: How many members do DBFs/PBFs/TBFs contain? E.g., for a DBF, its "members" are the domains in the DBF.

Methodology

Form matrix *M* with apps as columns and fingerprint members as rows. *M*[*i*, *j*] := 1 if fingerprint member *i* in app *j*'s fingerprint, 0 otherwise.

	app x	app y	app z
xyz.com	1	1	1
x.com	1	0	0
y.com	0	1	0

Agglomerative clustering on columns of *M*. Cosine distance for fingerprint similarity. Distance threshold of 0 when extracting clusters → apps must have identical fingerprints to end up in the same cluster.

Performance metrics now obtainable from matrix and clustering:
- Prevalence: # columns with >0 non-zero cells.
- Distinctiveness: # singleton clusters.
- Sizes: # non-zero cells in each column.

Why Do Some Apps Share the Same DBF?

- Apps that end up in the same cluster share the same DBF.
- Examined the developers responsible for apps that clustered together.
- Examined the domains in the DBFs of apps that clustered together.

	Apps that only share its DBF with other apps from same developer
Apple TV	61%
Fire TV	69%
Roku	27%

Takeaway: Apps that share the same DBF often stem from the same developer or have been developed using the same "no code" app generation toolkit/SDK.

Broader Impact

In addition to the presentation at the Privacy Enhancing Technologies Symposium (PETS), our work on advertising and tracking on smart TVs was also presented at:

- FTC's PrivacyCon 2021
- ProperData's Research Exploration Workshop on IoT and Privacy
- a DuckDuckGo corporate event

SaTC Frontiers: Collaborative: *1956393* [1], *1955227* [2], *2103439* [3], *1956435* [4]
Institutions: UC Irvine [1], Northeastern Univ. [2], UC Davis [3], USC [4]
Lead PIs: A. Markopoulou [1], D. Choffnes [2], Z. Shafiq [3], K. Psounis [4]