P3: Toward Privacy-Preserving Photo Sharing

Moo-Ryong Ra[*]
mra@usc.edu

Ramesh Govindan
ramesh@usc.edu

Antonio Ortega
ortega@sipi.usc.edu

University of Southern California

Abstract

With increasing penetration of mobile devices, photo sharing services are experiencing a resurgence. Aside from providing storage, photo sharing services enable bandwidth-efficient downloads to mobile devices by performing server-side image transformations (resizing, cropping). On the flip side, photo sharing services have raised privacy concerns such as leakage of photos to unauthorized viewers and the use of algorithmic recognition technologies by providers. To address these concerns, we propose a privacy-preserving photo encoding algorithm that extracts and encrypts a small, but significant, component of the photo, while preserving the remainder in a standards-compatible form. These two components can be separately stored. This technique significantly reduces the signal-to-noise ratio and the accuracy of automated detection and recognition on the publicly available photo, while preserving the ability of the provider to perform server-side transformations to conserve download bandwidth usage. Our prototype privacy-preserving photo sharing system, P3, works with Facebook and Flickr, and can be extended to other services as well. P3 requires no changes to existing services or mobile application software, and adds minimal photo storage overhead.

1 Overview and Contributions

With the advent of mobile devices with high-resolution on-board cameras, photo sharing has become tremendously popular. Users can share photos either through photo sharing services like Flickr or Picasa, or popular social networking services like Facebook or Google+.

However, this development has generated privacy concerns. Private photos have been leaked from a prominent photo sharing site [2]. Furthermore, widespread concerns have been raised about the application of face recognition technologies in Facebook [1]. Despite these privacy threats, it is not clear that the usage of photo sharing services will diminish in the near future. This is because photo sharing services provide several useful functions that, together, make for a seamless photo browsing experience. In addition to providing photo storage, PSPs also perform several server-side image transformations (like cropping, resizing and color space conversions) designed to improve user perceived latency of photo downloads and, incidentally, bandwidth usage.

In this paper, we explore the design of a privacy-preserving photo sharing algorithm (and an associated system) that *ensures photo privacy without sacrificing the latency, storage, and bandwidth benefits provided by PSPs*. This paper makes two novel contributions that, to our knowledge, have not been reported in the literature. First, the design of the P3 algorithm, which prevents leaked photos from leaking *information*, and reduces the efficacy of automated processing (e.g., face detection, feature extraction) on photos, while still permitting a PSP to apply image transformations. It does this by splitting a photo into a public part, which contains most of the *volume* (in bytes) of the original, and a secret part which contains most of the original's *information*. Second, the design of the P3 system, which requires no modification to the PSP infrastructure or software, and no modification to existing browsers or applications. P3 uses interposition to transparently encrypt images when they are uploaded from clients, and transparently decrypt and reconstruct images on the recipient side.

P3 is proof-of-concept of easily deployable privacy preserving photo storage. Adoption of this technology will be dictated by economic incentives: for example, PSPs can offer privacy preserving photo storage as a premium service offered to privacy-conscious customers.

2 Results Summary

P3 is a privacy preserving photo sharing scheme that leverages the sparsity and quality of images to store most of the information in an image in a secret part, leaving most of the volume of the image in a JPEG-compliant public part, which is uploaded to PSPs. P3's public parts have very low PSNRs and are robust to edge detection, face detection, or sift feature extraction attacks. These benefits come at minimal costs to reconstruction accuracy, bandwidth usage and processing overhead.

References

[1] Facebook Shuts Down Face Recognition APIs After All, http://www.theregister.co.uk/2012/07/09/facebook_face_apis_dead/.

[2] CNN: Photobucket leaves users exposed. http://articles.cnn.com/2012-08-09/tech/tech_photobucket-privacy-breach.

[*] Moo-Ryong Ra is a Ph.D. student.