Cloud and Mobile Computing

Protect Privacy in Offloading

Yung-Hsiang Lu
Electrical and Computer Engineering
Purdue University

Technological Trends

- Mobile systems become primary computing platforms for most people.
- Mobile systems have limited resources: storage, performance, network bandwidths, and battery life.
- Cloud computing gains popularity replacing in-house servers and desktop applications.
- Existing cloud services: webmail, video hosting, social networks …
- Can the cloud help mobile users? Can mobile systems help the cloud?
  - Mobile systems need resources.
  - The cloud needs (real-time) data.

"Cheap" Massive Parallelism

- allow users to rent many computers for only hours
- meet the increasing need for storage, organization, analysis of large amounts of data (image, video, audio, document…) by end users
- respond to events by allocating resources quickly (for example, simulations in emergency)

Computation Offloading

mobile system

heavy computation

(e.g. image search and recognition)

high-performance server

Demo: Computation Offloading for Robot

Challenges

- offer fine-grained offloading services (for seconds or minutes, not hours or months)
- schedule real-time tasks with high parallelism (such as image processing and object recognition)
- provide easy-to-program interface with automatic parallelism detection and scalability
- tolerate bandwidth fluctuations with multiple levels of details / accuracy
- design programming languages for applications whose executions may be migrated easily
- **protect privacy**

Are you willing to put private information in the cloud?

Will Cloud Computing Kill Privacy?

Erik Larkin, PCWorld Jan 27, 2010 2:15 pm

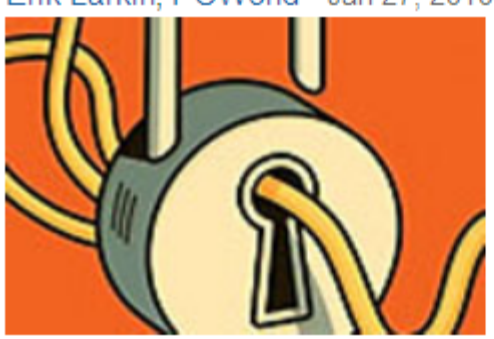

As cloud computing speeds ahead, privacy protections are too often being left in the dust.

Loosely defined, cloud computing involves programs or services that run on Internet servers. Despite the buzz surrounding it, the idea isn't new--think Webmail. But huge benefits, such as being able to gain access to your data from anywhere and not having to worry about backups, have led more people to leap to the Internet to do everything from writing documents and watching movies to managing their businesses. Unfortunately, privacy is often still stuck at home.

Behind the Times

Archaic laws that focus on where your information is, rather than what it is, are part of the

Can security concerns kill cloud computing? | IT PRO - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://www.itpro.co.uk/610299/can-security-concerns-kill-cloud-computing

Google

Can security concerns kill cloud com...

In this channel: News Features Reviews How Tos & Tuto

Home : Security : Features

Can security concerns kill cloud computing?

The IT industry has identified cloud computing as a trend for the future. But how much of a barrier to it development and adoption will security be?

By Miya Knights, 23 Apr 2009 at 16:22

It seems as though not much new is happening in enterprise IT developm doesn't involve the cloud.

The New Economy

Privacy issues hit Facebook again

Privacy issues are again at the center of debate over whether Facebook is doing enough to protect its users.

Facebook CEO Mark Zuckerberg speaks during a session at the Cannes Lions 2010 International Advertising Festival in Cannes, June 23. The company faces questions about privacy issues after a security researcher Wednesday compiled and released personal information on more than 100 million Facebook users.

Sebastien Nogier/Reuters/File

Enlarge

By *Alissa Figueroa*, Correspondent / July 30, 2010


The darker side of Webmail - Computerworld - Mozilla Firefox
File Edit View History Bookmarks Tools Help
http://www.computerworld.com/s/article/9078638/The_darker_side_of_Webmail
Google
The darker side of Webmail - Compu...
Home > Networking
The darker side of Webmail
Web-based e-mail may be exposing you to privacy and security problems you didn't expect
By Tam Harbert
April 28, 2008 12:00 PM ET
Comments (13)
Recommended (297)
Share
Computerworld - Web-based e-mail is booming. Services such as Gmail, Yahoo Mail and Hotmail are convenient, accessible and, best of all, free. Many of us have come to rely on them without giving it a second thought.
But second thoughts may be in order, according to security experts, privacy advocates and some Webmail users. Few consider the fact that Webmail is inherently different than POP3 e-mail. It differs in who administers it and how, in the ways it may be vulnerable to hacking, and in the type of help you can expect when you have a problem.
You may not think these differences matter. And they don't -- unless they
Top
Irar
Vis
Elg
Chi

Google Privacy Blunder Shares Your Docs Without Permission - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://techcrunch.com/2009/03/07/huge-google-privacy-blunder-shares-your-docs-without-perm

TC Google Privacy Blunder Shares Your ...

Tech Gadgets Mobile Enterprise GreenTech CrunchBase Disrupt SF More

TechCrunch

What's Hot: Android Apple Facebook Google Microsoft Twitter Yahoo Zynga Subs

TechCrunch Disrupt: The Agenda Of Awesome, Last Day for Discounted Tickets »

Google Privacy Blunder Shares Your Docs Without Permission

Jason Kincaid

Mar 7, 2009

116 Comments

In a privacy error that underscores some of the biggest problems surrounding cloud-based services, Google has sent a notice to a number of users of its Document and Spreadsheets products stating that it may have inadvertently shared some of their documents with contacts who were never granted access to them.

According to the notice, this sharing was limited to people "with whom you, or a collaborator with sharing rights, had previously shared a document" – a vague

Google Searches Used To Convict Hit-And-Run Driver | Techdirt - Mozilla Firefox

File Edit View History Bookmarks Tools Help

http://www.techdirt.com/articles/20090115/0559143421.shtml

Google Searches Used To Convict Hit...

< Police In Mumbai Shutting Down Open WiFi

Carl's Jr. Apparently Unaware That People Share... >>

Legal Issues

by Mike Masnick

ri, Jan 16th 2009 1:53pm

Share This

Filed Under: nviction, google earches, hit and run

0 tweets

retweet

0

Google Searches Used To Convict Hit-And-Run Driver

from the *google-searches-in-a-court-of-law* dept

In the past, we've noted various lawsuits where Google searches done by the accused were used against them in a court of law. There was the guy who searched on "neck snap break," days before his wife was murdered, and then there was the woman who searched on "how to commit murder" and other rather damning phrases like "instant poison" and "undetectable poisons," before her husband was murdered. In yet another such case, an investment banker has been convicted of a hit-and-run that killed a woman, after his Google searches soon after the accident turned up the phrase "hit and run." The guy had claimed that he believed he hit a deer, but his Google searches suggested he knew it was a person. Beyond just searching for the phrase hit and run, he also did searches on: "auto glass reporting requirements to law enforcement," "auto glass, Las Vegas," auto parts, auto theft, and the Moraga Police Department. Since the incident was in California, the thinking was he was looking to get the damage to his car repaired out of state to avoid any suspicion from the auto repair place. While the guy appealed the ruling saying that even with those searches he didn't have any actual knowledge he had hit a person, the appeals court didn't find that to be very convincing.

14 Comments | Leave a Comment..

Search

Searc

And no

See

Energy Savings in Privacy-Preserving Computation Offloading with Protection by Homomorphic Encryption

Jibang Liu and Yung-Hsiang Lu
Electrical and Computer Engineering
Purdue University

HotPower 2010, Vancouver, Canada

Save Mobile Energy by Offloading

mobile system → heavy computation (e.g. image search and recognition) → high-performance server

Technology to Protect Data

- Anonymize
- Erase history
- Steganography ⇨
- Watermark
- Encryption
- …

inside

image source: Wikipedia

- **Contribution**: first paper showing how to use homomorphic encryption for offloading and saving mobile energy.

Is it possible to perform computation on encrypted data?

Rivest, Adleman, Dertouzos 1978

Process without Access

Darkroom

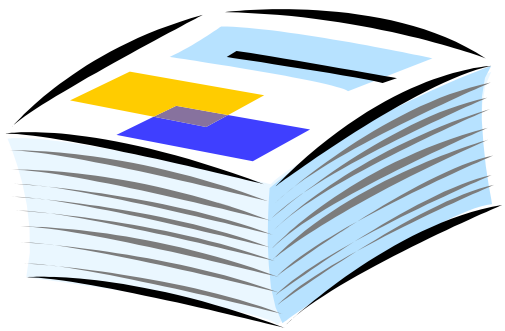

Homomorphic Encryption

x: plaintext	r: result
y: ciphertext	e(x) = y
e: encryption	d(y) = x
d: decryption	f(x) = r
f: operation	**d(f(y)) = r**

Homomorphic encryption is **not** an encryption algorithm (like RSA, DES, or AES). Instead, it says that **some** encryption algorithms allow operations on the encrypted data.

Example (Addition)

- Suppose p and q are two prime numbers, n = pq.
- y = e(x) = (x + pr) mod n, r is a user-chosen integer
- x = d(y) = y mod p
- x must be smaller than p
- p = 7, q = 5, x1 = 2, x2 = 1
- y1 = (2 + 2 × 7) mod 35 = 16, choose 2 for r
- y2 = (1 + 6 × 7) mod 35 = 8, choose 6 for r
- x1 + x2 = 3
- y1 + y2 = 16 + 8 = 24
- 24 mod 7 = 3

Example (Multiplication)

- $y = e(x) = (x + pr) \bmod n$, r is a random integer
- $x = d(y) = y \bmod p$
- p = 7, q = 5, x1 = 2, x2 = 1
- $y1 = (2 + 2 \times 7) \bmod 35 = 16$, choose 2 for r
- $y2 = (1 + 6 \times 7) \bmod 35 = 8$, choose 6 for r
- $x1 \bullet x2 = 2$, $y1 \bullet y2 = 16 \bullet 8 = 128$
- $128 \bmod 7 = 2$
- $(16 \bullet 16 \bullet 16) \bmod 7 = 4096 \bmod 7 = 1$ **wrong ("overflow")**
- $2 \bullet 2 \bullet 2 = 8$
- In practice, p and q are very large.

Nondeterministic Encryption

f(x) = r

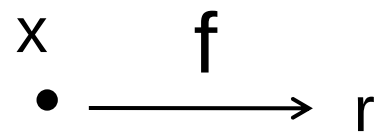

f is + and • in the examples

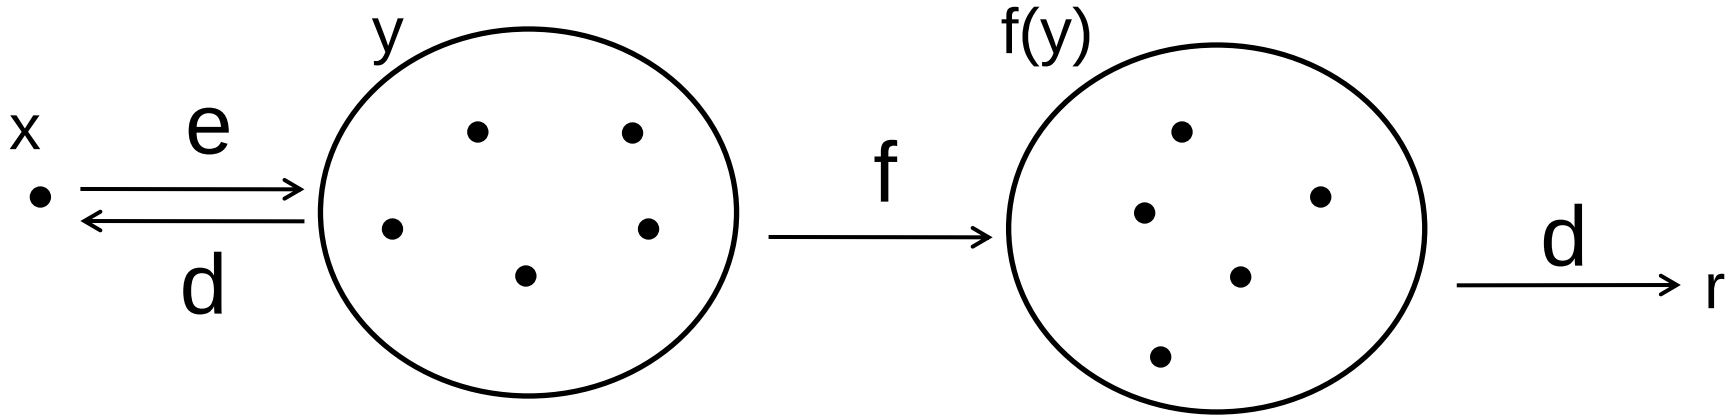

sometimes

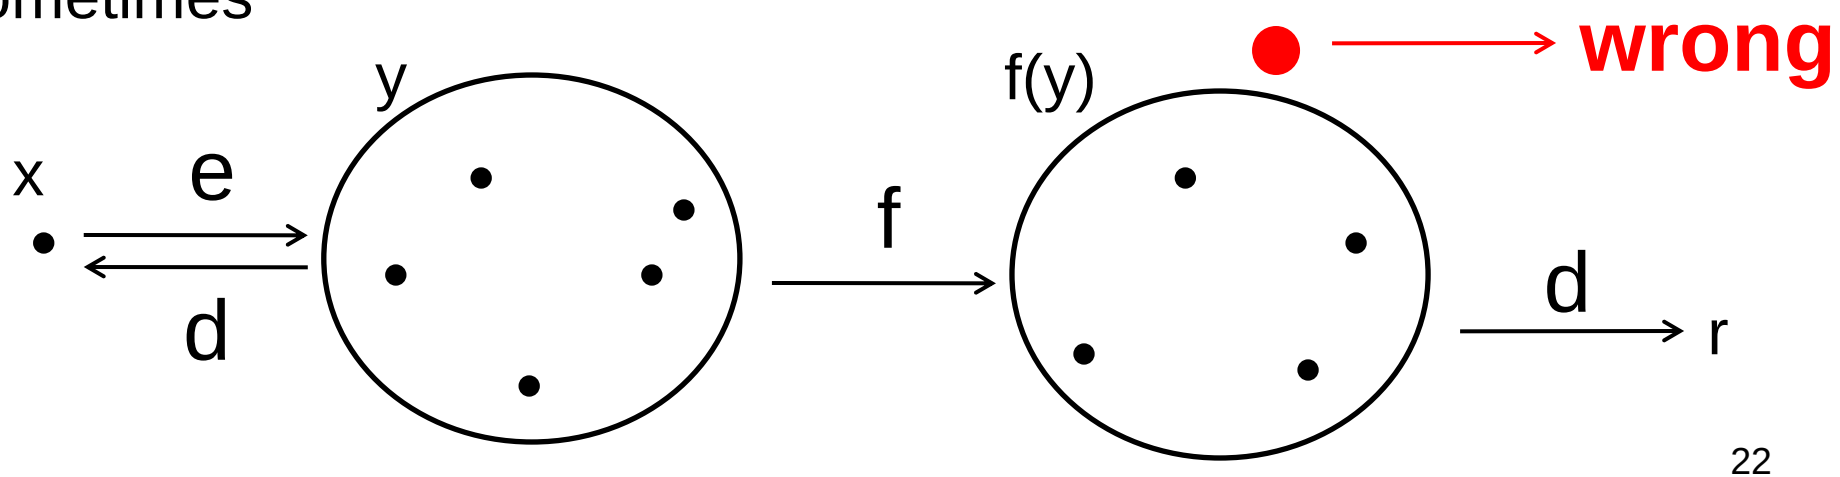

"Noise" in Homomorphic Encryption

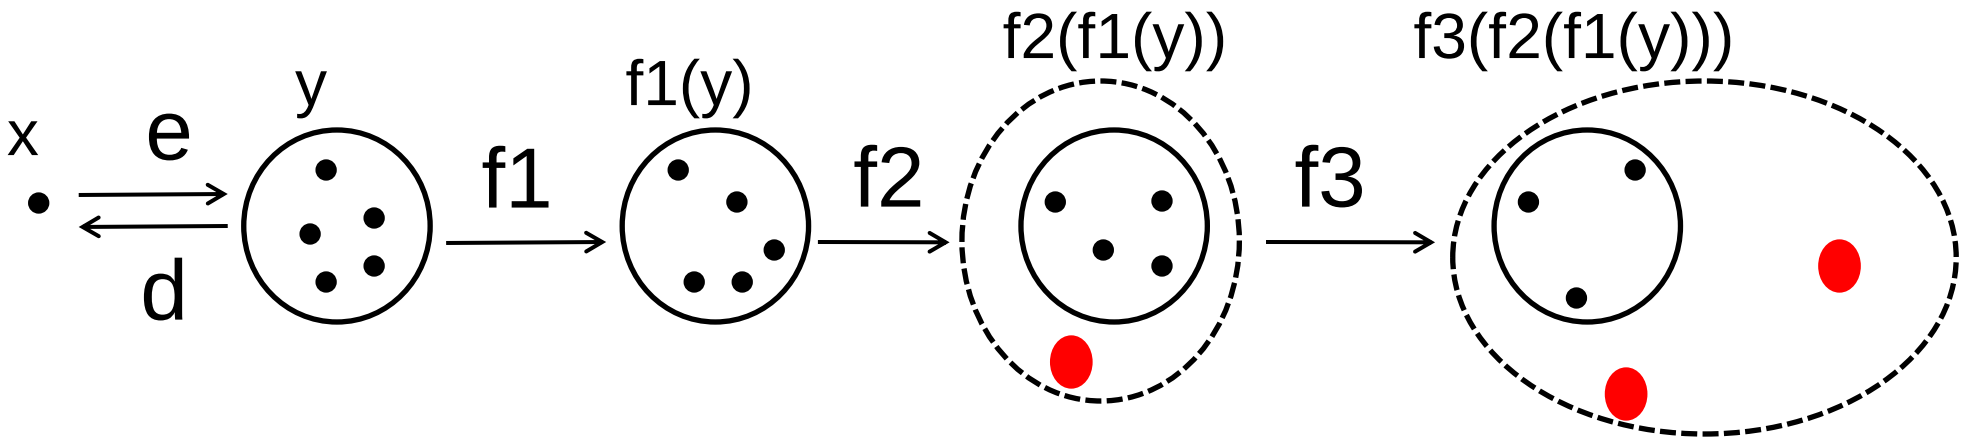

- more operations ⇒ noise accumulates ⇒ eventually falls outside the region.
- The region's size (i.e. tolerance of noise) depends on the encryption key (larger key, better tolerance)
- Frequent denoising is needed but no efficient solution was discovered until [Gentry STOC 2009]
- This paper does **not** consider denoising.

Save Mobile Energy by Offloading

mobile system → heavy computation (e.g. image search and recognition) → high-performance server

Gabor Filtering
(Insensitive to Rotations and Noise)


Original


4-pixel average

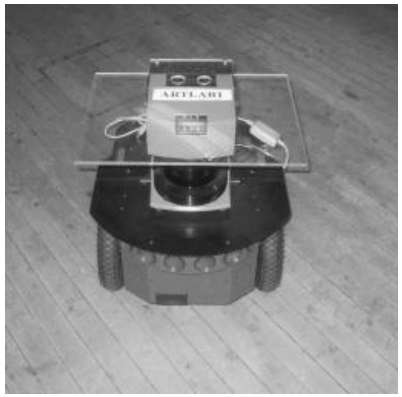
gray level


shuffle 4-pixel block


45-degree rotation


90-degree rotation


180-degree rotation


motion blur


zoom blur


noise

Step 1/2 (S1) Scale + Convolution

$$G(x,y) = \frac{1}{2\pi\sigma^2}\exp(-\frac{x'^2 + y'^2}{2\pi\sigma^2})\cos(2\pi fx')$$

$$x' = x\cos\theta + y\sin\theta$$

$$y' = -x\sin\theta + y\cos\theta$$

σ = 1, 10, 100, 200
in our evaluation

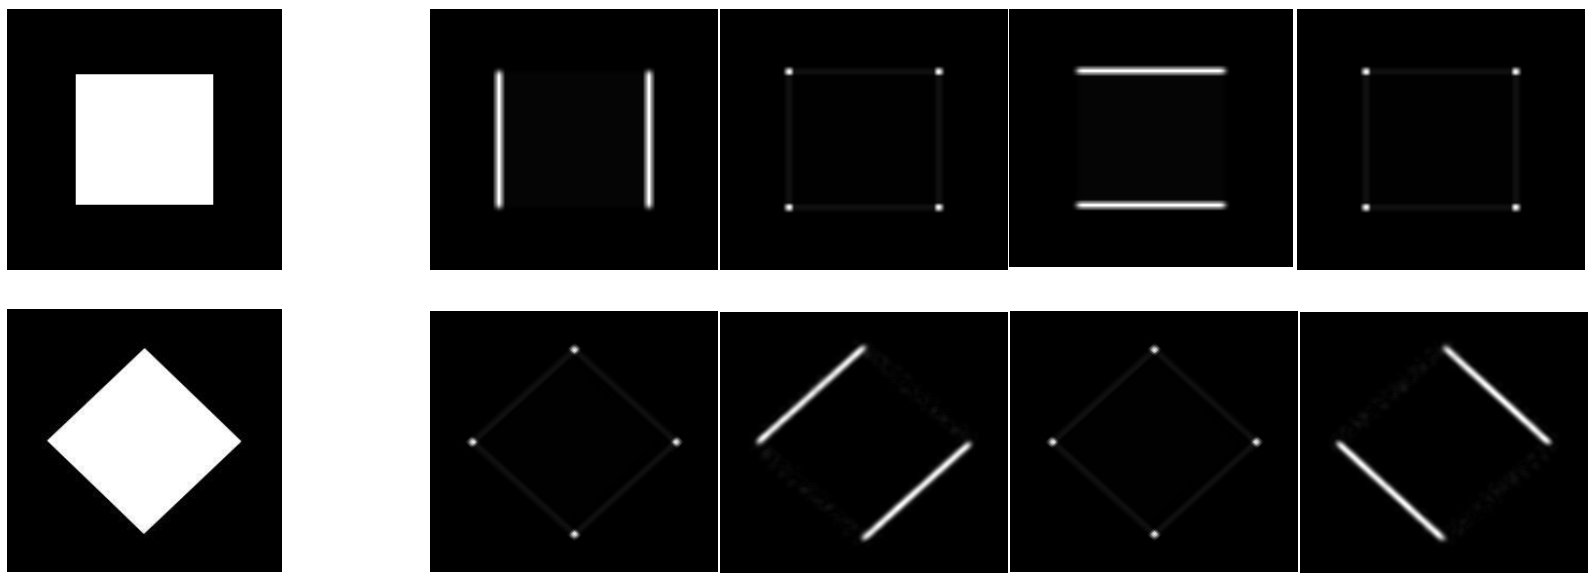

$\theta = 0$ $\pi/4$ $\pi/2$ $3\pi/4$

scale: convert floating point values to integers
by multiplying a large number (such as 10,000)

Step 2/2 (S2) Features

- computing "features" using means, standard deviations, and distributions of these images
- compare the features to find similar images

	S1	**S2**
computation intensive	Y > 99%	N <1%
operations	+ and ×	+ and ÷ and $\sqrt{}$
efficient operations on encrypted data	Y	N
offload	Y	N

Offloading Options

(a) images ⇨ S1 ⇨ S2 ⇨ result — mobile

(b) images ⇨ S1 ⇨ S2 ⇨ result — server

(c) images ⇨ encrypt ⇨ S1 ⇨ decrypt ⇨ S2 ⇨ result

	Save Energy	Protect Privacy
(a)	N	Y
(b)	Y	N
(c)	Y	Y

Parameters and Protections

- encryption key ↑
  - + protection ↑, accuracy ↑
  - - energy saving ↓
- scaling factor ↑
  - + accuracy ↑
- attacks
  - ciphertext-only attacks ⇦
  - known plaintext
  - chosen plaintext
  - adaptively-chosen plaintext

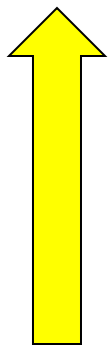

more difficult for an attacker

Experimental Setup

- Mobile: HP iPAQ 6954 PDA
- Server: 2GHz CPU, 3GB memory
- power measurement: National Instrument data acquisition, power from battery,1 KHz sampling
- accuracy measured by recall
  - L: number of returned images (20 in our evaluations)
  - Y: number of similar images (10 in our evaluation)
  - X: number matched images

$$recall(L) = \frac{X}{Y}$$

Results

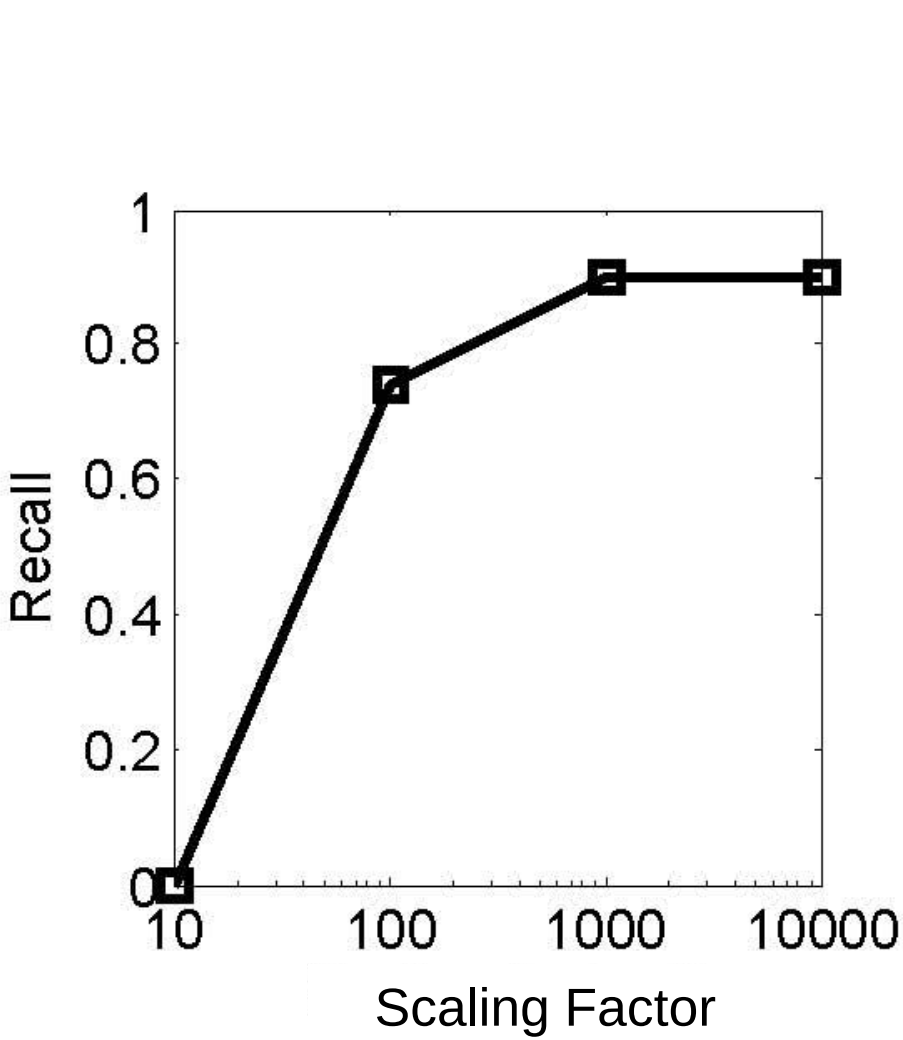
Recall
1
0.8
0.6
0.4
0.2
0
10
100
1000
10000
Scaling Factor

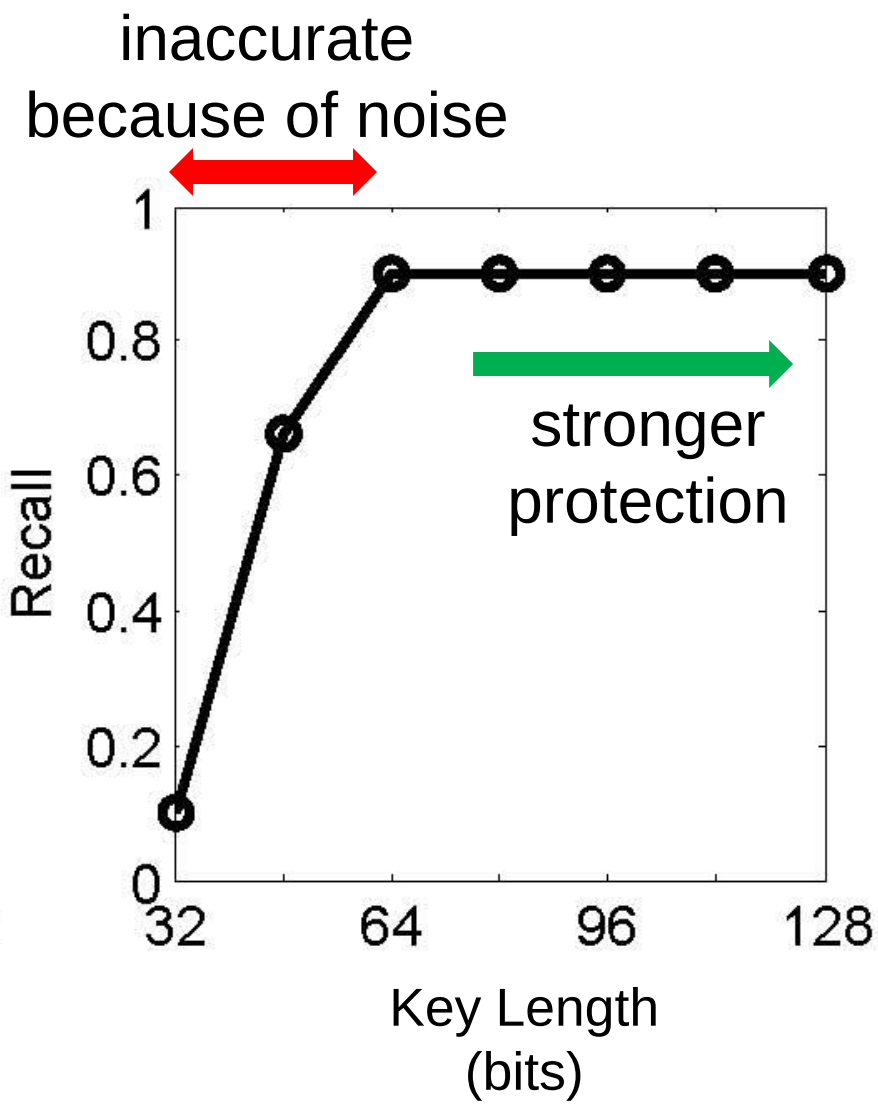
inaccurate
because of noise
stronger
protection
Recall
1
0.8
0.6
0.4
0.2
0
32
64
96
128
Key Length
(bits)

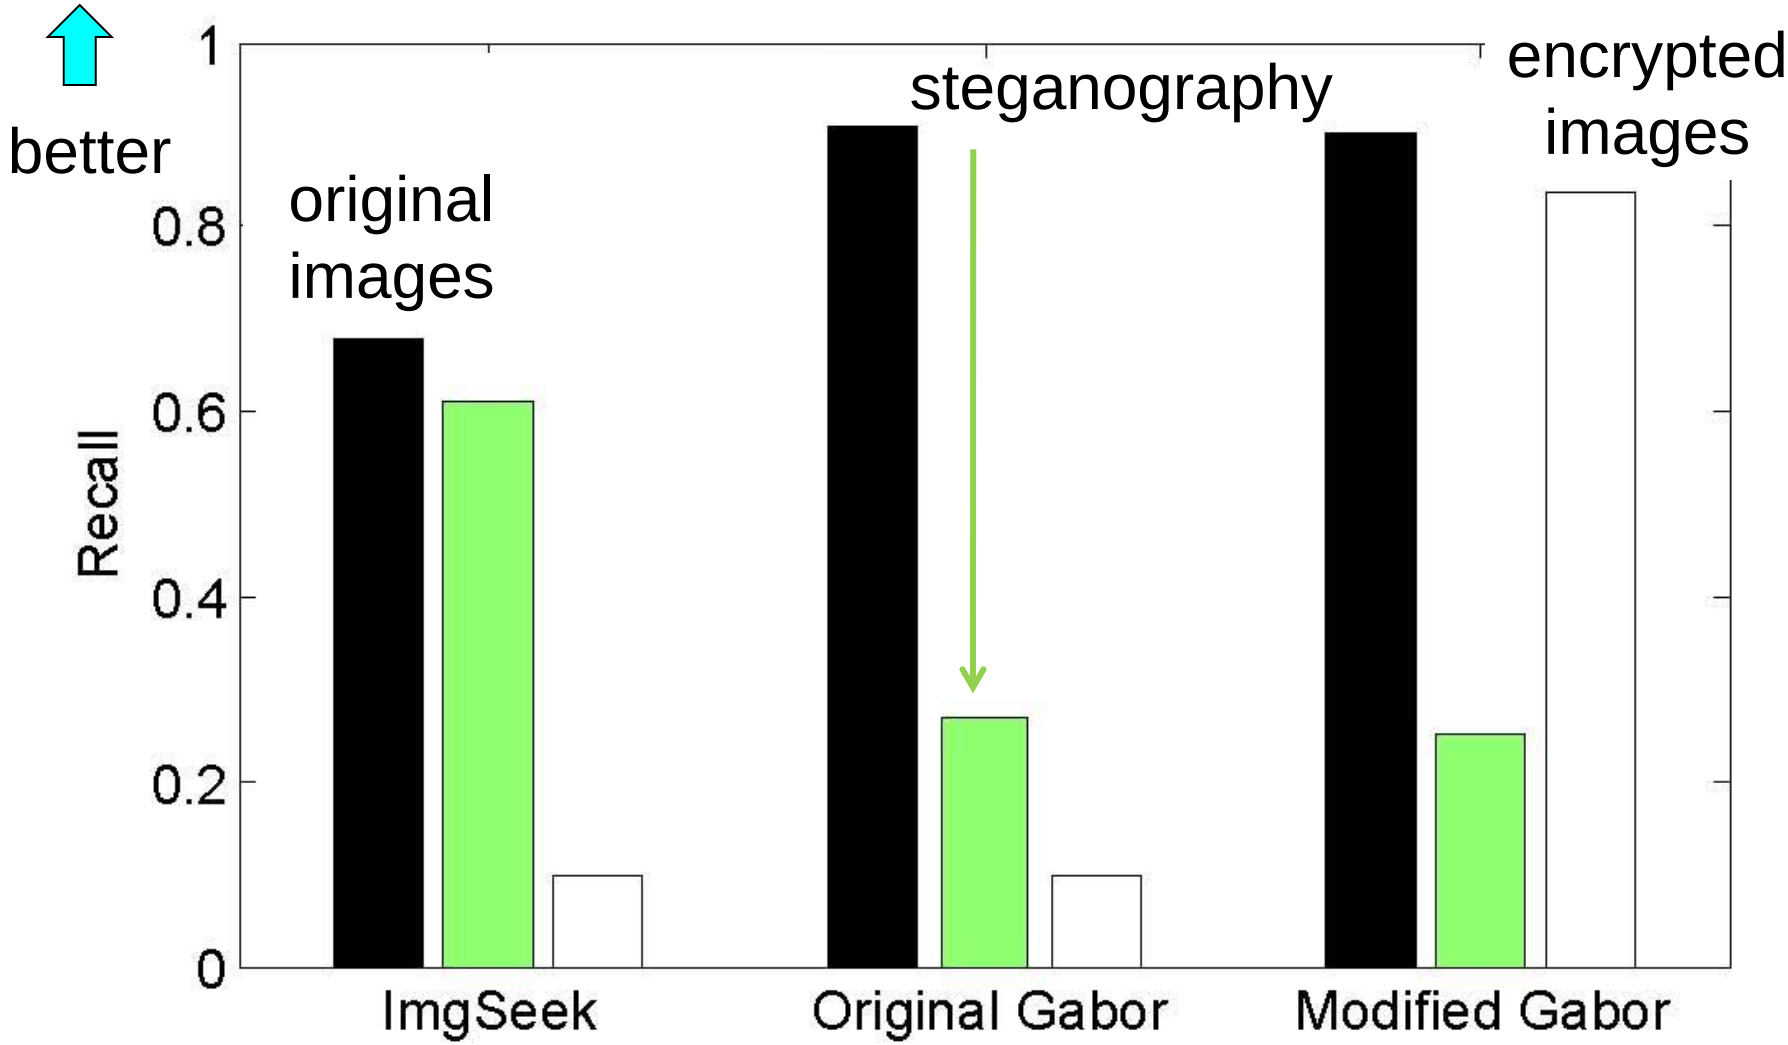

- compared with steganography [ISLPED 2010]
- The modified Gabor filter can handle both original and encrypted images.

10,000 images from Flickr evaluated on the server

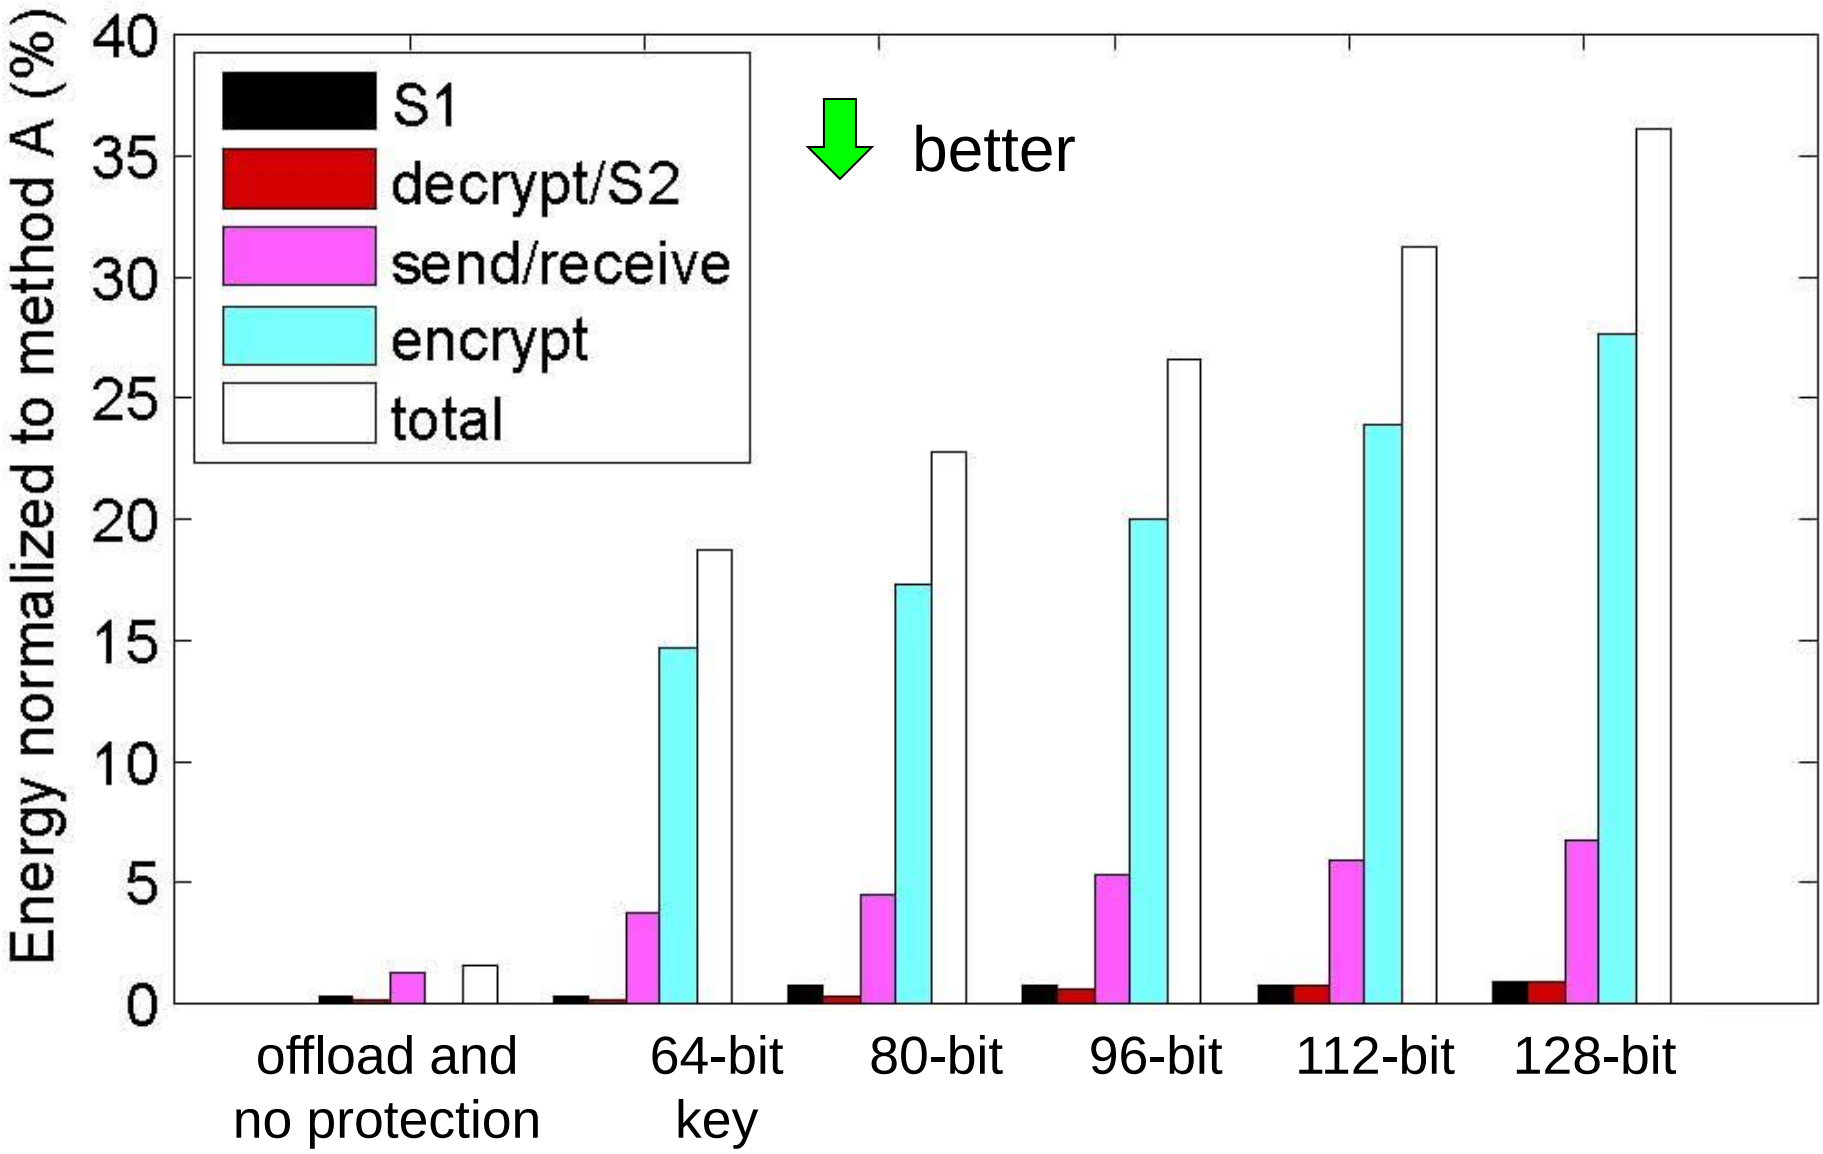

- 1000 images on the PDA, sent to server at run time
- baseline: no protection, on PDA, ~2 hours

Future Work

- migrate to newer mobile systems and re-evaluate the parameters
- extend the Gabor filter to handle zooming
- use SIFT (scale invariant feature transformation) to handle distortion
- implement denoising and evaluate the effects on energy consumption and performance
- compare different encryption algorithms

Conclusion

- Present a method to offload image retrieval with data protected by homomorphic encryption.
- Obtain accuracy comparable to no protection.
- Modify Gabor retrieval algorithm that can handle unprotected and protected images.
- Evaluate the effects of key size and scaling factor.

Acknowledgements: This work is supported in part by NSF grants CNS-0716271. Any opinions, findings, and conclusions or recommendations in the paper are those of the authors and do no necessarily reflect the views of the sponsor.

Encryption and Decryption

- $e(x) = (x + pr) \equiv y \bmod n$
- $x + pr = bpq + y$
- $y = x + pr - bpq$
- $d(y) = y \bmod p$
- $(x + pr - bpq) \bmod p \equiv (x + p(r - bq)) \bmod p = x$