

Inside Dropbox: Understanding Personal Cloud Storage Services

- *Idilio Drago*
- Marco Mellia
- Maurizio M. Munafò
- Anna Sperotto
- Ramin Sadre
- Aiko Pras

- **Personal cloud storage** services are already popular
- Dropbox in 2012
 - *“the largest deployed networked file system in history”*
 - *“over 50 million users – one billion files every 48 hours”*
- **Little public information about the system**
 - How does Dropbox work?
 - What are the potential performance bottlenecks?
 - Are there typical usage scenarios?

Methodology – How does Dropbox work?

- Public information
 - **Native client**, Web interface, LAN-Sync etc.
 - Files are split in **chunks of up to 4 MB**
 - Delta encoding, deduplication, **encrypted** communication
- To understand the client protocol
 - **MITM** against our own client
 - Squid proxy, SSL-bump and a self-signed CA certificate
 - **Replace a trusted CA certificate in the heap at run-time**
- Proxy logs and decrypted packet traces

How does Dropbox (v1.2.52) work?

- Clear separation between **storage** and meta-data/client **control**
- Sub-domains identifying parts of the service

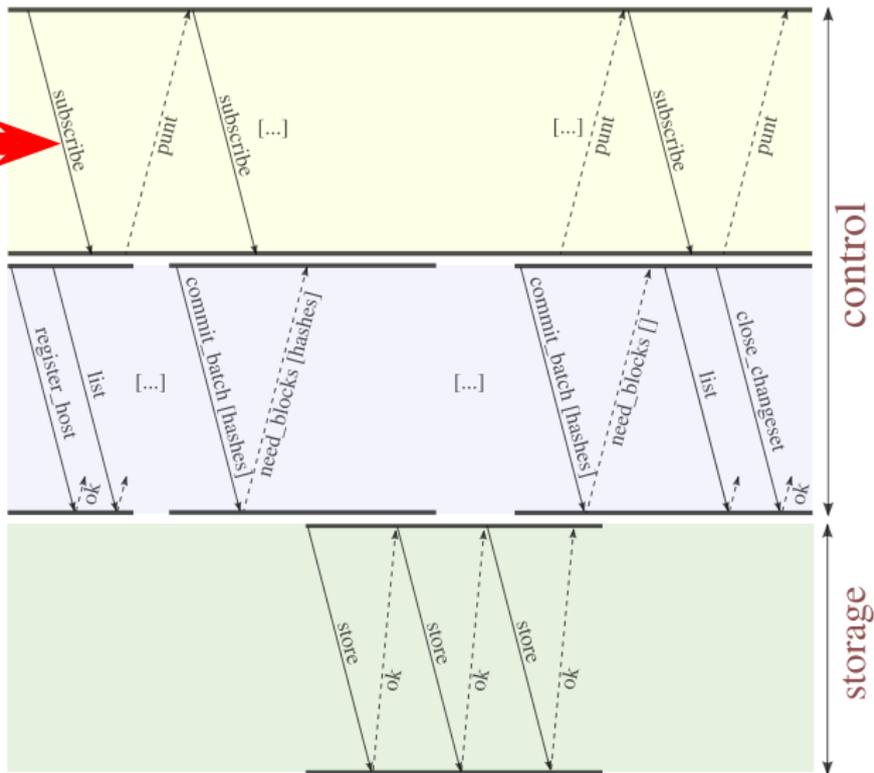
sub-domain	Data-center	Description
client-lb/clientX	Dropbox	Meta-data
notifyX	Dropbox	Notifications
api	Dropbox	API control
www	Dropbox	Web servers
d	Dropbox	Event logs
dl	Amazon	Direct links
dl-clientX	Amazon	Client storage
dl-debugX	Amazon	Back-traces
dl-web	Amazon	Web storage
api-content	Amazon	API Storage

- **HTTP/HTTPS** in all functionalities

How does Dropbox (v1.2.52) work?

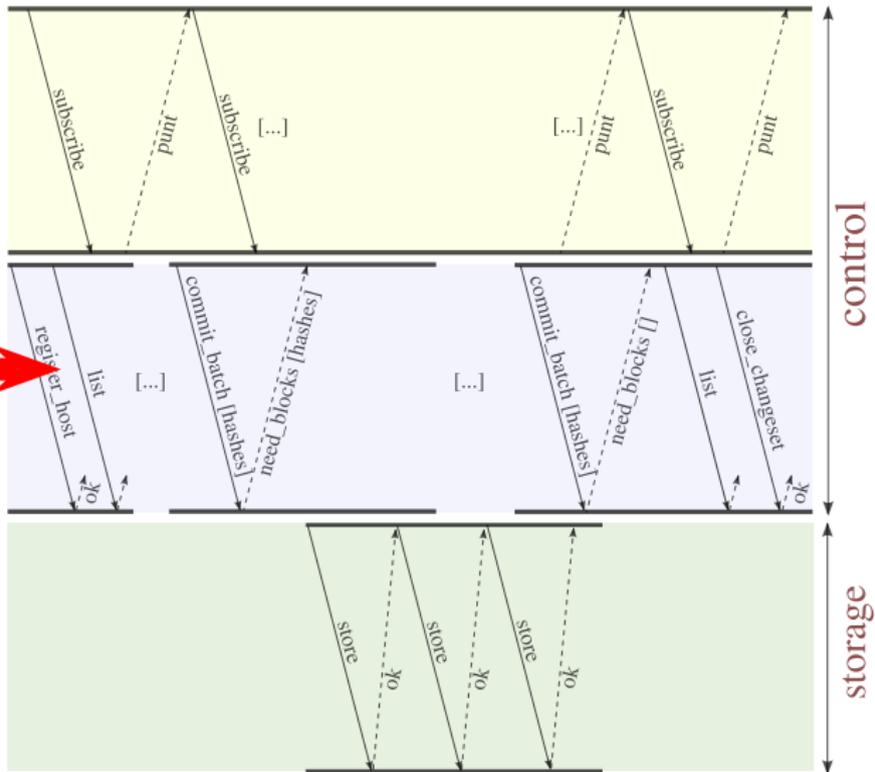
■ Notification

- Kept open
- Not encrypted**
- Device ID**
- Folder IDs

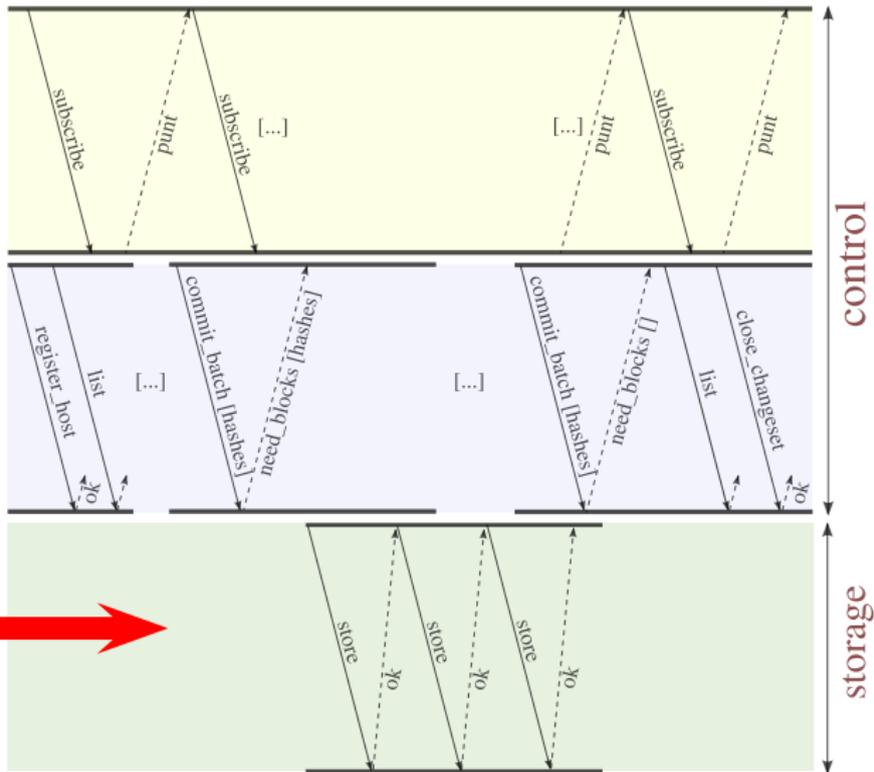


How does Dropbox (v1.2.52) work?

- Client control
- Login
- File hash
- Meta-data



How does Dropbox (v1.2.52) work?



■ Storage

- Amazon EC2
- **Retrieve vs. Store**
- **Sequential ACKs**

- Rely on **Tstat**¹ to export **layer-4 flows**
- Isolate Dropbox flows
 - **DN-Hunter**², TSL/SSL certificates, IP addresses
- **Device IDs** and folder IDs
- Use the knowledge from our own decrypted flows to
 - **Tag Dropbox flows** – e.g., *storing* or *retrieving* content
 - Estimate the **number of chunks** in a flow

¹<http://tstat.polito.it/>

²DNS to the Rescue: Discerning Content and Services in a Tangled Web

	Type	IP Addr.	Dropbox		
			Flows	Vol. (GB)	Devices
Campus 1	Wired	400	167,189	146	283
Campus 2	Wired/Wireless	2,528	1,902,824	1,814	6,609
Home 1	FTTH/ADSL	18,785	1,438,369	1,153	3,350
Home 2	ADSL	13,723	693,086	506	1,313
Total			4,204,666	3,624	11,561

- 42 consecutive days in March and April 2012
 - 4 vantage points in Europe

	Type	IP Addr.	Dropbox		
			Flows	Vol. (GB)	Devices
Campus 1	Wired	400	167,189	146	283
Campus 2	Wired/Wireless	2,528	1,902,824	1,814	6,609
Home 1	FTTH/ADSL	18,785	1,438,369	1,153	3,350
Home 2	ADSL	13,723	693,086	506	1,313
Total			4,204,666	3,624	11,561

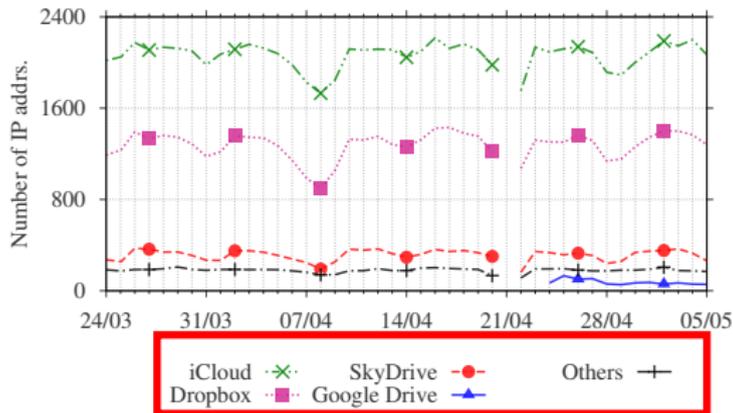
- 42 consecutive days in March and April 2012
 - 4 vantage points in Europe
 - Number of IP addresses in home probes \approx installations

	Type	IP Addr.	Dropbox		
			Flows	Vol. (GB)	Devices
Campus 1	Wired	400	167,189	146	283
Campus 2	Wired/Wireless	2,528	1,902,824	1,814	6,609
Home 1	FTTH/ADSL	18,785	1,438,369	1,153	3,350
Home 2	ADSL	13,723	693,086	506	1,313
Total			4,204,666	3,624	11,561

- 42 consecutive days in March and April 2012
 - 4 vantage points in Europe
 - Number of IP addresses in home probes \approx installations
 - **11,561 unique devices**
- 2nd capture in Campus 1 in June 2012

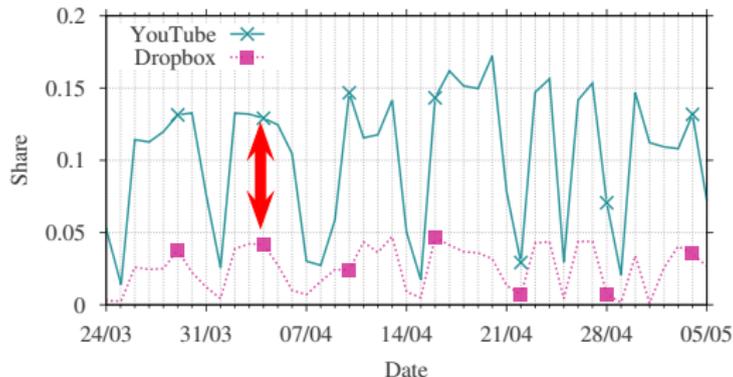
How much traffic to personal cloud storage?

7



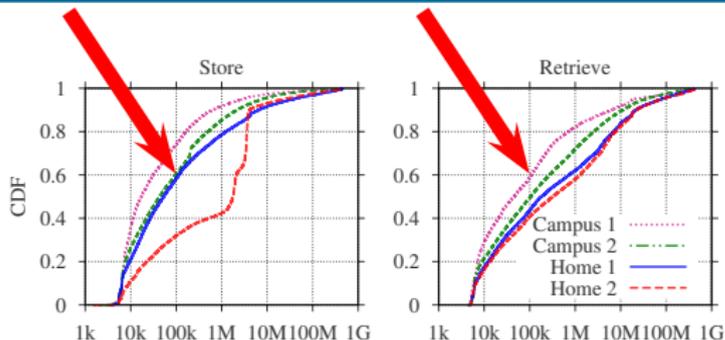
- **Server names** to check popularity (DN-Hunter)
- 6 – 12 % adoption in home networks
- iCloud tops in terms of devices

How much traffic to personal cloud storage?



- Equivalent to **1/3 of YouTube volume at Campus 2**
- 90 % of the Dropbox traffic is from the native client

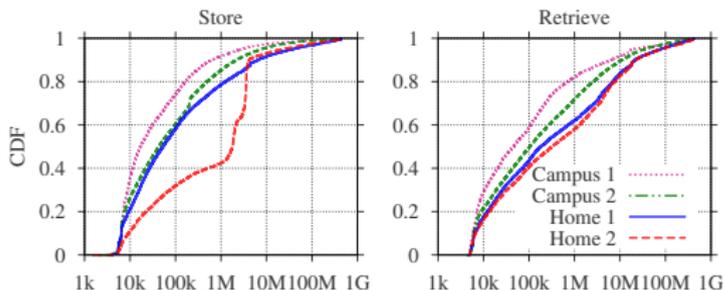
How does the storage traffic look like?



■ Flow size

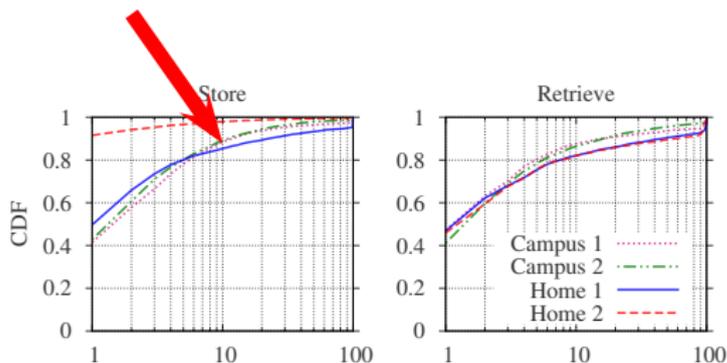
- **Store: 40 % – 80 % < 100 kB**
- Small files and deltas
- **Larger retrieve flows**

How does the storage traffic look like?



■ Flow size

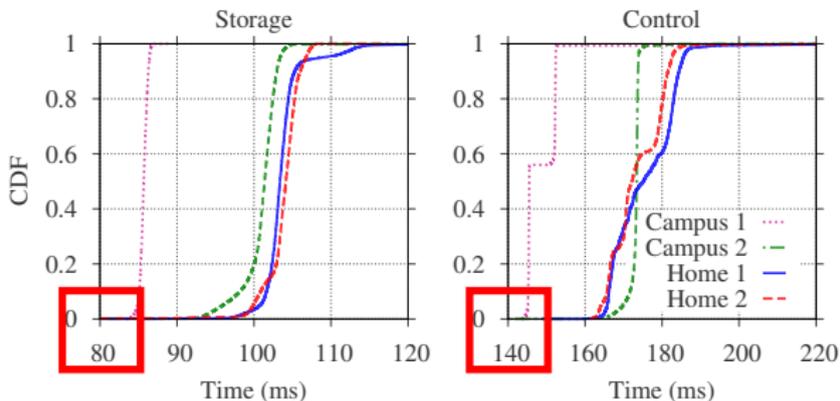
- **Store: 40 % – 80 % < 100 kB**
- Small files and deltas
- **Larger retrieve flows**



■ Chunks per flow

- **80 % ≤ 10 chunks**
- **Remaining: up to 100**
- Limited by the client

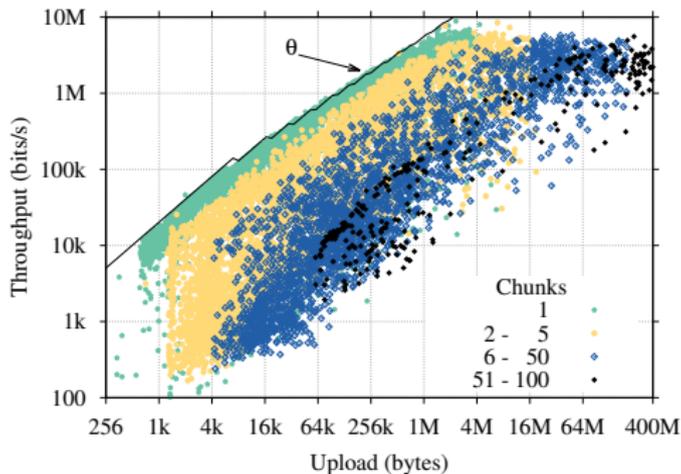
Where are the servers located?



- Minimum RTT per flow → stable over 42 days
- PlanetLab experiments → **the same U.S. data centers worldwide**
- *“less than 35 % of our users are from the USA”*

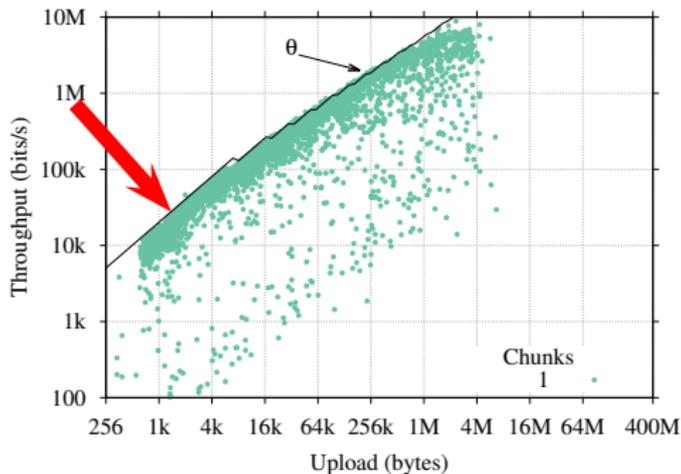
How is the performance far from the data centers?

11



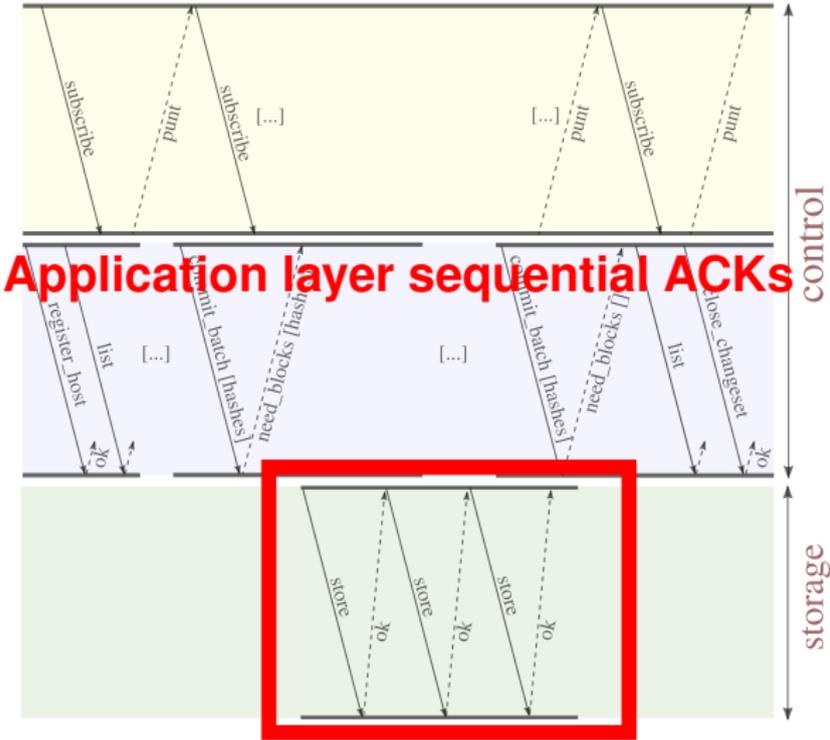
- Storage throughput in campuses
- **Most flows experience a low throughput**

How is the performance far from the data centers?

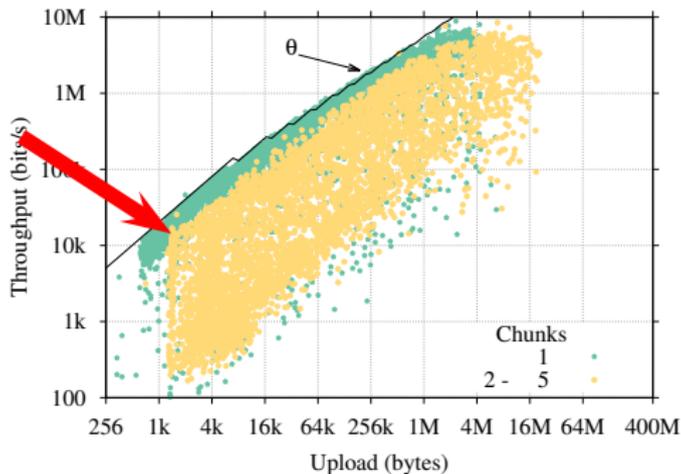


- Flows carrying 1 chunk
 - Size \leq 4 MB, RTT \approx 100 ms
 - Most of them finish in **TCP slow-start**

How is the performance far from the data centers?



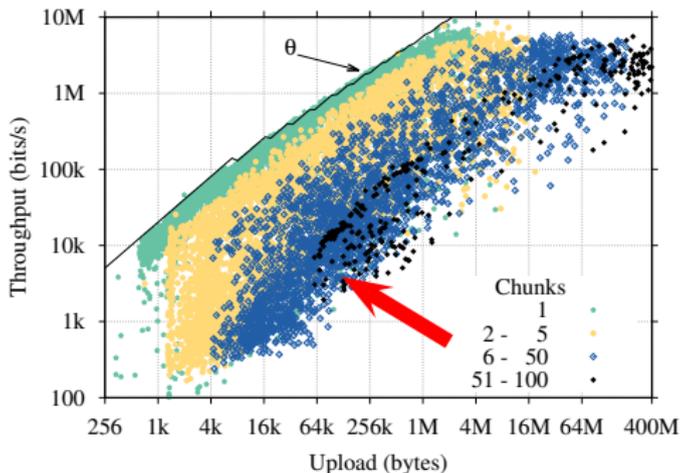
How is the performance far from the data centers?



- Flows carrying several chunks
 - **Pause between chunks** → RTT and client/server reaction

How is the performance far from the data centers?

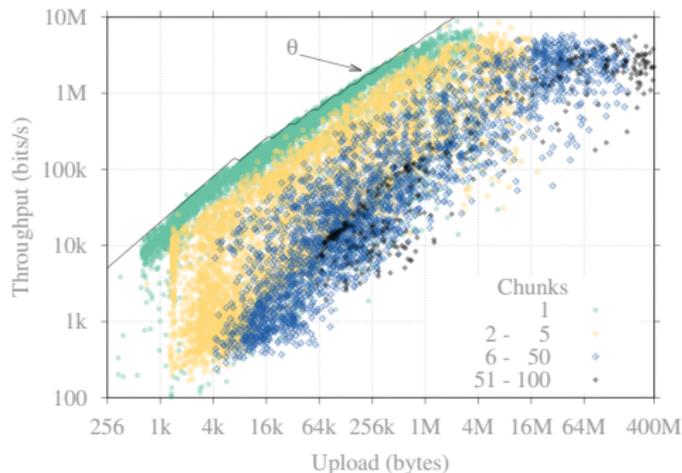
11



- Flows carrying several chunks
 - Transferring **100 chunks takes more than 30 s**
 - RTTs → 10 s of inactivity

How is the performance far from the data centers?

11



- **Delaying acknowledgments**
- **Bundling chunk** → deployed after our 1st capture
- **Distributing servers** → storage traffic is heavy!

How much improvement from chunk bundling?

- New protocol released on Apr 2012 (v 1.4.0)
 - Small chunks are bundled together

	Mar/Apr		Jun/Jul	
	Median	Average	Median	Average
Flow size				
<i>Store</i>	16.28 kB	3.91 MB	42.36 kB	4.35 MB
<i>Retrieve</i>	42.20 kB	8.57 MB	70.69 kB	9.36 MB
Throughput (kbits/s)				
<i>Store</i>	31.59	358.17	81.82	552.92
<i>Retrieve</i>	57.72	782.99	109.92	1293.72

- **Less small flows** → less TCP slow-start effects
- Average **throughput** is up to **65 % higher**

How are other storage systems implemented?

	Dropbox	SkyDrive	Wuala	Google Drive	Cloud Drive
Chunking	4 MB	variable	variable	8 MB	X
Bundling	✓	X	X	X	X
Deduplication	✓	X	✓	X	X
Delta encoding	✓	X	X	X	X
Compression	always	never	never	smart	never

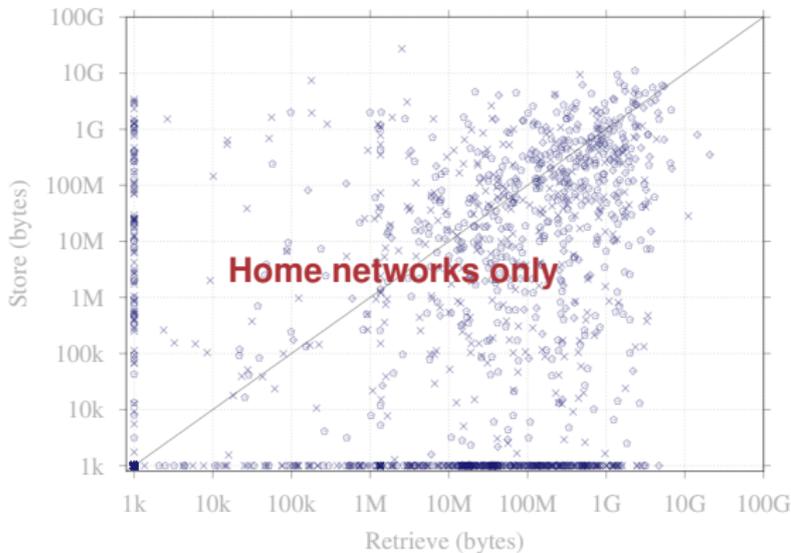
- Comparison of design choices of different providers
- **Benchmarking Personal Cloud Storage** – IMC 2013

How are other storage systems implemented?

	Dropbox	SkyDrive	Wuala	Google Drive	Cloud Drive
Chunking	4 MB	variable	variable	8 MB	X
Bundling	✓	X	X	X	X
Deduplication	✓	X	✓	X	X
Delta encoding	✓	X	X	X	X
Compression	always	never	never	smart	never

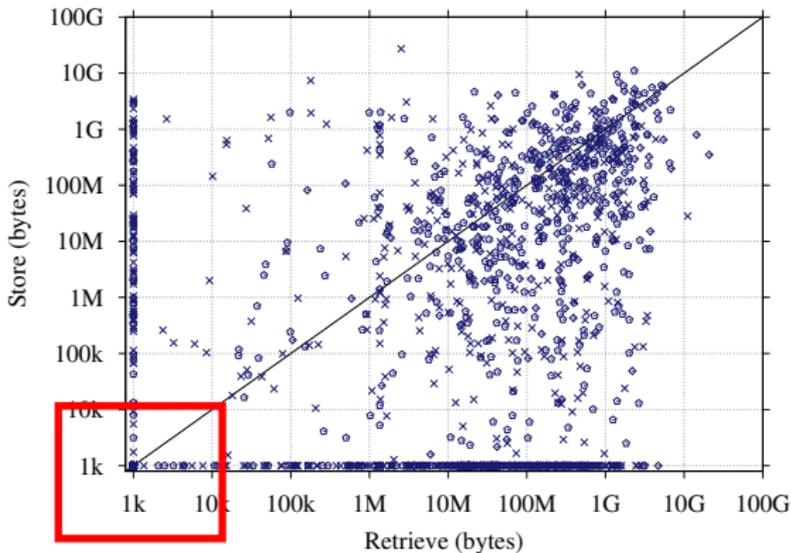
- Are **batches of files** exchanged in a single transaction?
 - Cloud Drive and Google Drive open several TCP connections per file

Are there typical usage scenarios?



- **More downloads** → download/upload ratio up to **2.4**
- What about download/upload per user?

Are there typical usage scenarios?



Occasional:

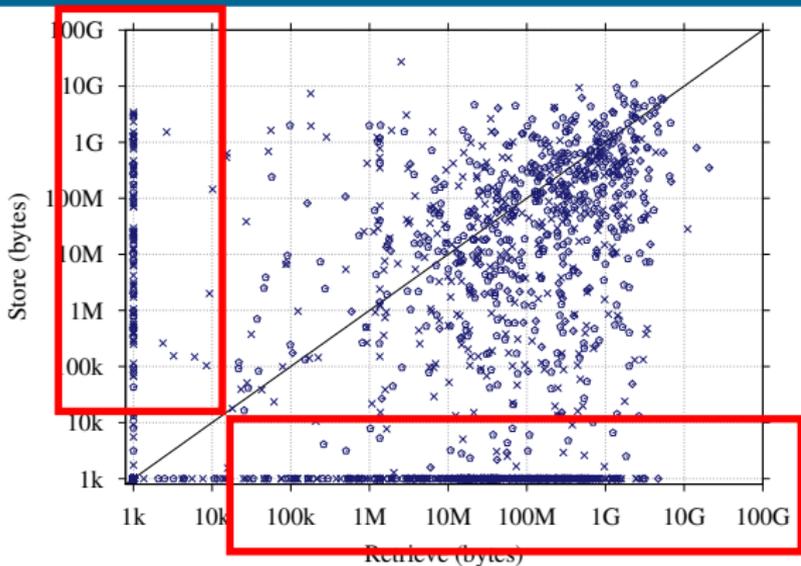
□ Users: 31 %

□ Devices per user: 1.22

■ Abandoned Dropbox clients

■ No storage activity for 42 days

Are there typical usage scenarios?



- Backup and content sharing
- Geographically dispersed devices

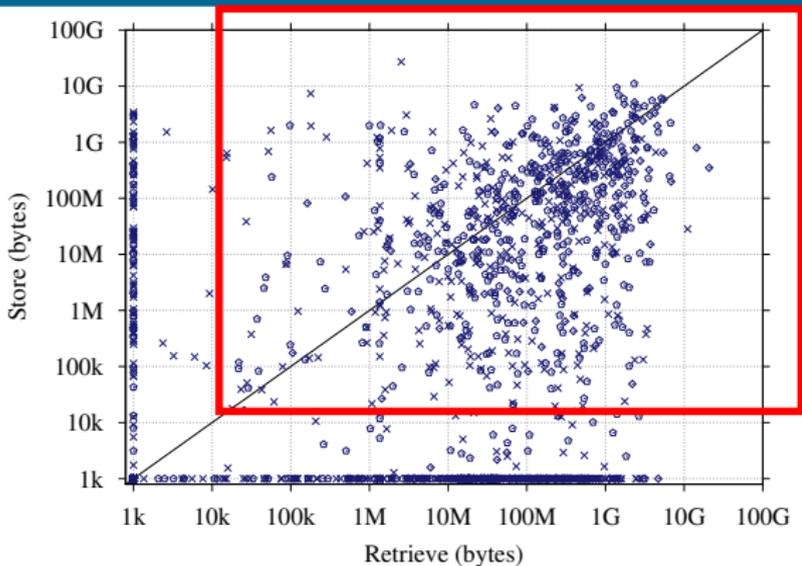
■ Upload-only:

- Users: 6 %
- Uploads: 11 – 21 %
- Devices per user: **1.36**

■ Download-only:

- Users: 26 %
- Downloads: 25 – 28 %
- Devices per user: **1.69**

Are there typical usage scenarios?



■ Heavy:

- Users: 37 %
- Uploads: 79 – 89 %
- Downloads: 72 – 75 %
- Devices per user: 2.65

■ Synchronization of content in a household

- **1st to analyze Dropbox usage on the Internet**
- Cloud storage is a new data-intensive application
 - Adoption above 6 % in our datasets
- Architecture and performance
 - **Bottlenecks from system design choices**
- **Extensive characterization of workload and usage**
 - User groups, number of devices, daily activity etc.

- **Thank You.**

- Anonymized traces and scripts

- `http://traces.simpleweb.org/dropbox/`