# A Long-Term User-Centric Analysis of Deduplication Patterns

## 32nd International Conference on Massive Storage Systems and Technology (MSST 2016)

Zhen Sun,[1,2] Geoff Kuenning,[3] Sonam Mandal,[2] Philip Shilane,[4]
Vasily Tarasov,[5] Nong Xiao,[1,6] Erez Zadok[2]

*[1]HPCL, NUDT, China; [2]Stony Brook University; [3]Harvey Mudd College;*
*[4]EMC Corporation; [5]IBM Research – Almaden; [6]SYSU, China*


# Outline

- Introduction
- Data-set description
- Deduplication-ratio & File-based Analysis
- User-based Analysis
- Conclusion and Future Work

# Introduction

- Deduplication has been widely deployed in both backup and primary storage.

- Data sets analysis plays an important role in deduplication study.
  - Backup Storage (FAST'13, MSST'14).
  - Primary Storage (ATC'15, SYSTOR'09, SYSTOR'12, FAST'11).
  - Archival Storage (ICIVC'12).
  - HPC centers (SC'12).
  - And more……

# Motivation

- More data-set studies are needed:
  - Data-set characteristics vary significantly.
    - Whole file chunking (WFC) efficiency varies from 20%~87% (ATC'12, SC'12, FAST'12).
  - Most previous works study static data-set or cover a short period.
  - New findings can help us make better design decisions.

- What makes our work special:
  - Long-term backup study.
    - Covering > 4,000 snapshots from > 21 months.
  - User-Centric:
    - Study from users' perspective produces surprising results.

# Data Set: *FSL-Homes*

| Data Set | FSL-Homes |
|---|---|
| Organization | 1 snapshot per user per day |
| Total Size | 456TB |
| Start and end time | 03/09/2012 – 11/23/2014 |
| Number of users | 33 |
| Number of Snapshots | 4,181 dailies (about 21 months) |
| Chunking methods | Content-defined Chunking, Whole File Chunking |
| Average Chunking Size | 2, 4, 6, 8, 16, 32, 64 and128 KB |
| Hashing Method | 48 bit MD5 hash. (Hash collision rate < 0.004% using 2KB chunking) |
| Number of files | 130 million |
| Meta-data included | File pathname, size, atime, mtime, ctime, UID, GID, permission bits, device ID, inode number |

# Data Set: *FSL-Homes*

- Limitations:
  - File content is not stored.
    - Time/Space consuming to store all the data.
    - Not suitable for content-based analysis.
  - Some periods were not collected.
    - Data-collection is hard for many reasons.
    - Long breaks when data-set remained unchanged.

- Link: *http://tracer.filesystems.org*
  - Contains both tools and data-set.
  - Has been used in a number of papers.
  - Data set will be periodically updated.

# Deduplication Ratio Analysis

- Simulated 3 backup methods:
  - Daily-Full backup.
  - Incremental backup.
  - Weekly-full backup.
- Due to high redundancy:
  - Meta-data consumes large fraction of total space.
  - Small chunking size is not always better.
  - Different backup methods have their own best chunking size.

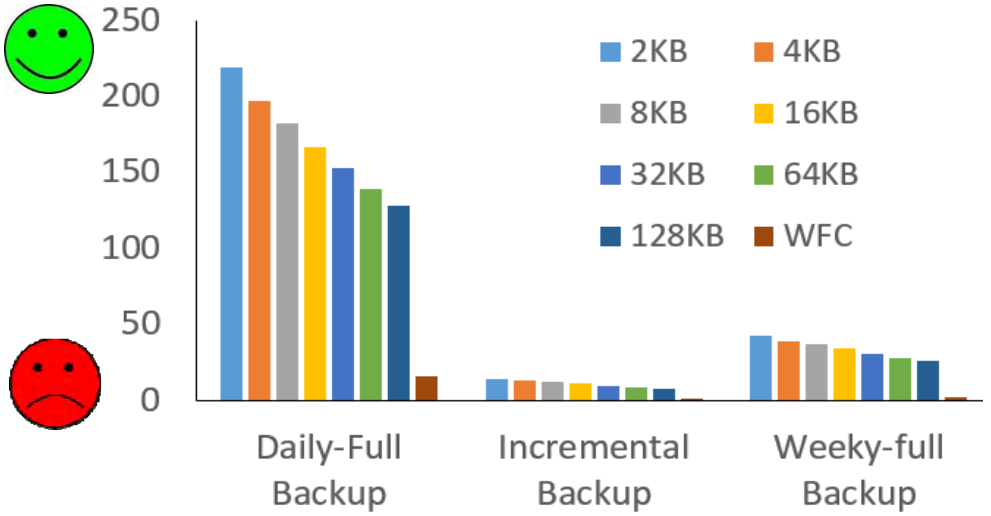


Raw Deduplication Ratio

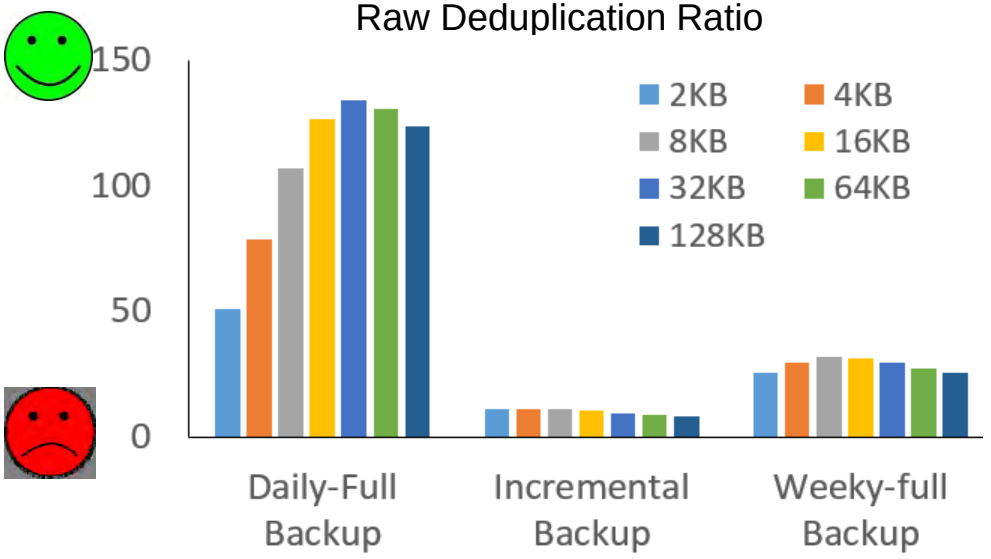


Effective Deduplication Ratio

# Whole File Chunking

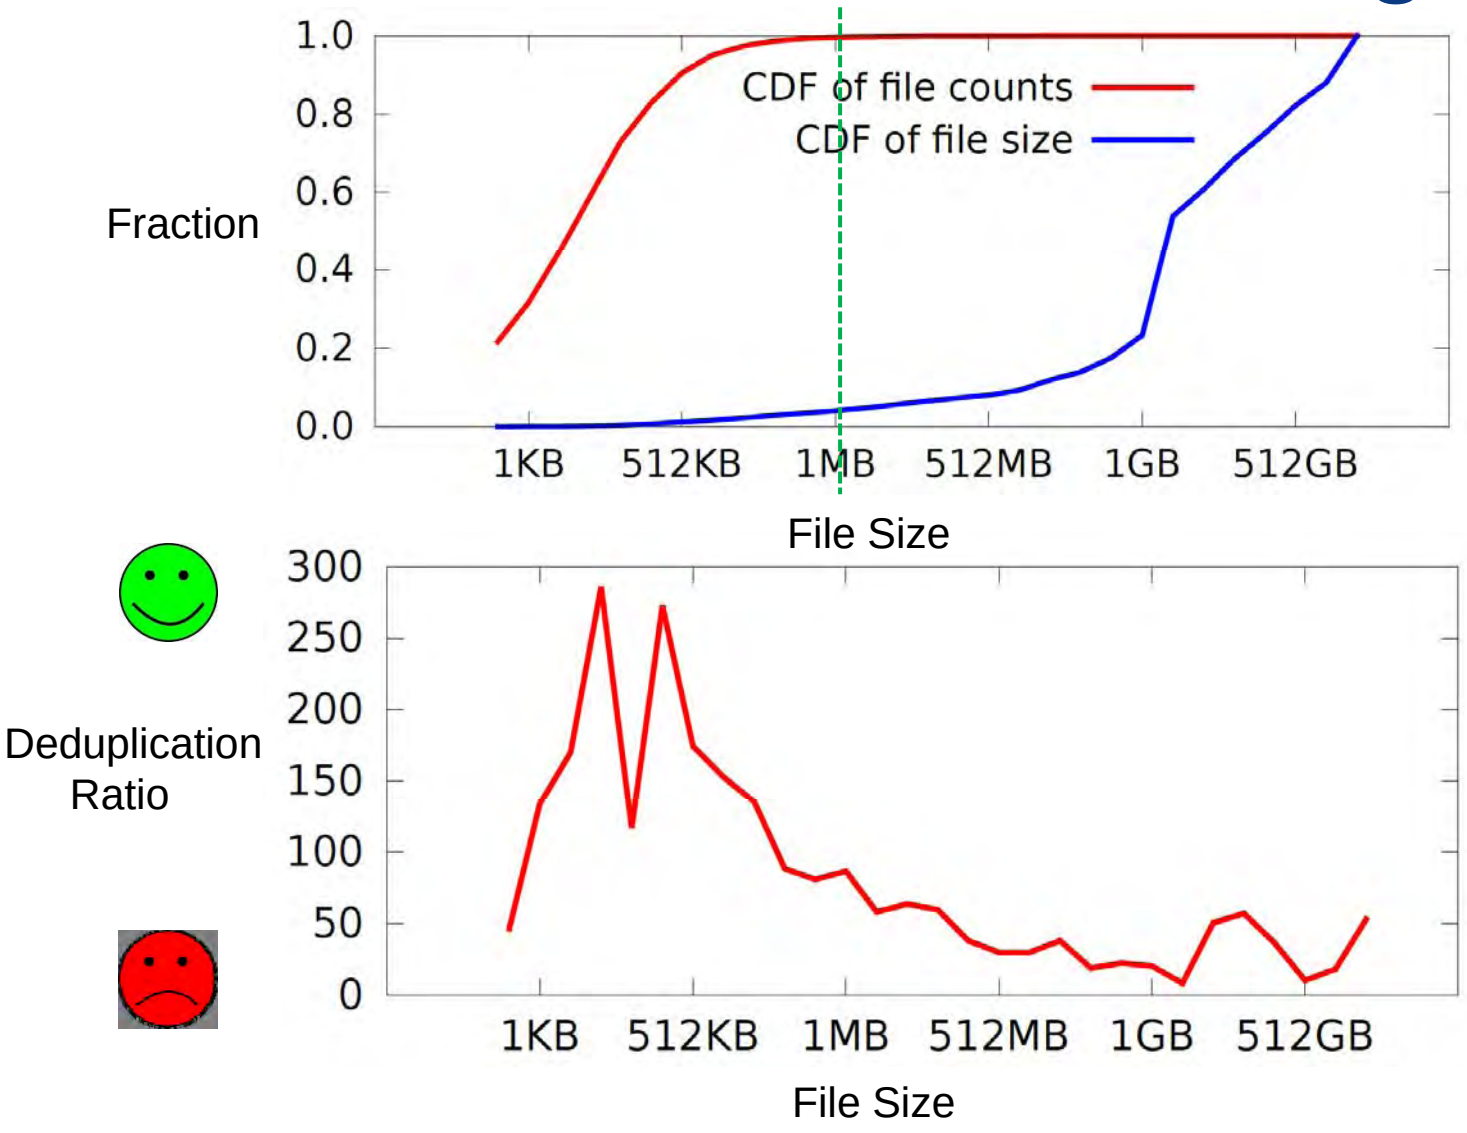

# File Analysis

- VMDK files take ~60% of total space .
- Different file types have hugely different deduplication ratio and sensitivity to chunking

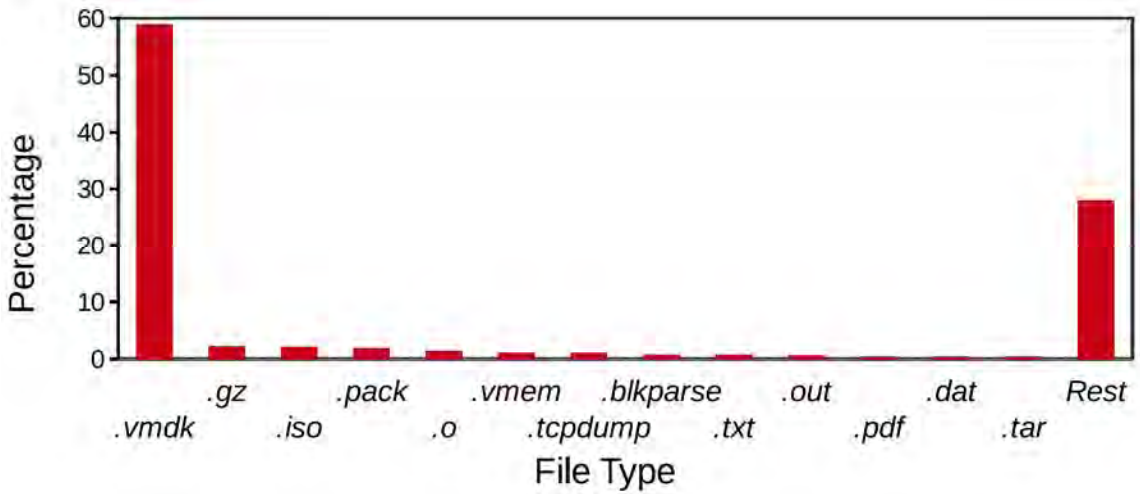


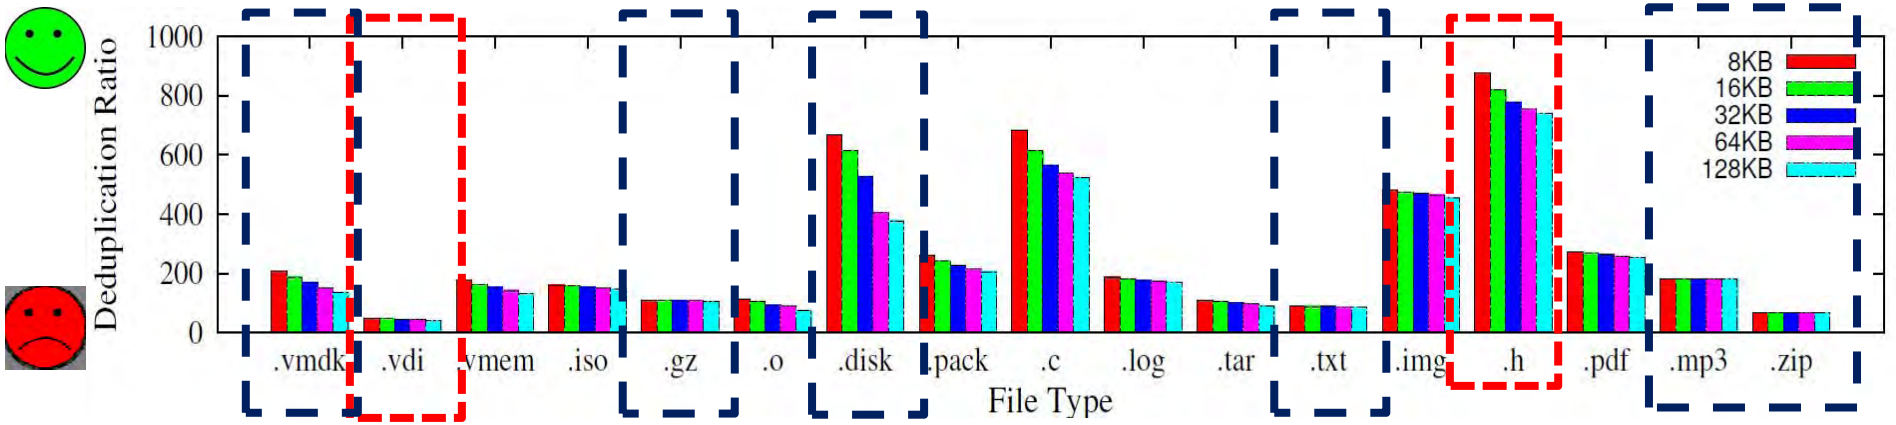

# Per-User Analysis 1/2

- All representative users are carefully chosen.
  - We selected users that covered different characteristics.
- Users' deduplication ratio differs a lot.
- Users' sensitivity to chunking size is also different.

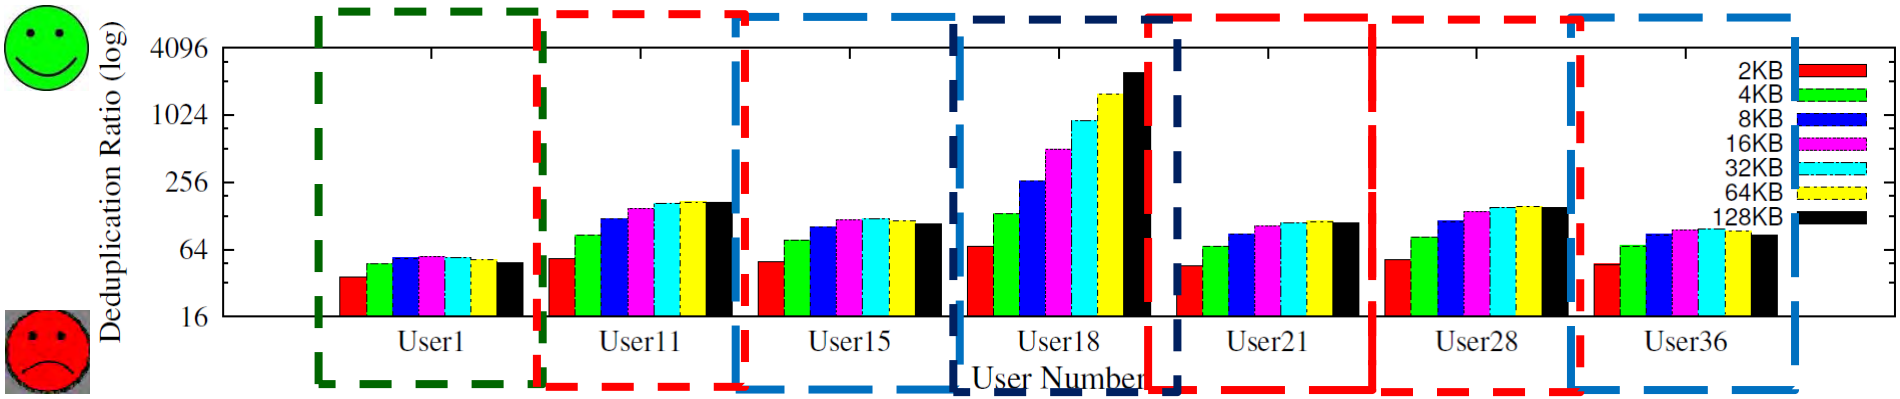

# Per-User Analysis 2/2

- Why users' deduplication ratio differ so much?
  - Users' lifetime?
  - Users' file types?
  - Users' own characteristics:
    - Internal deduplication ratio.
    - Activity level.

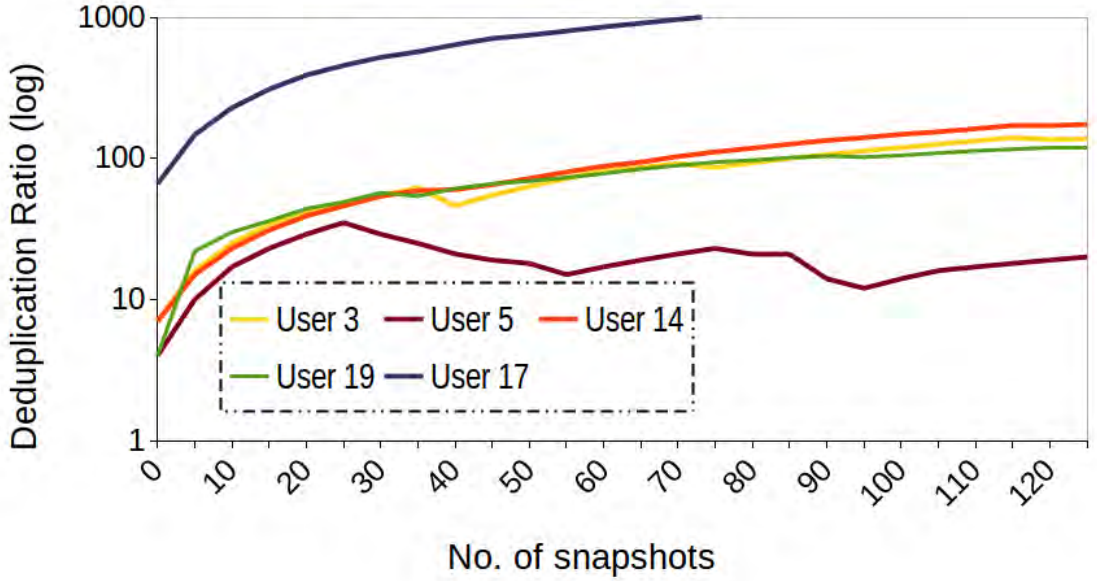


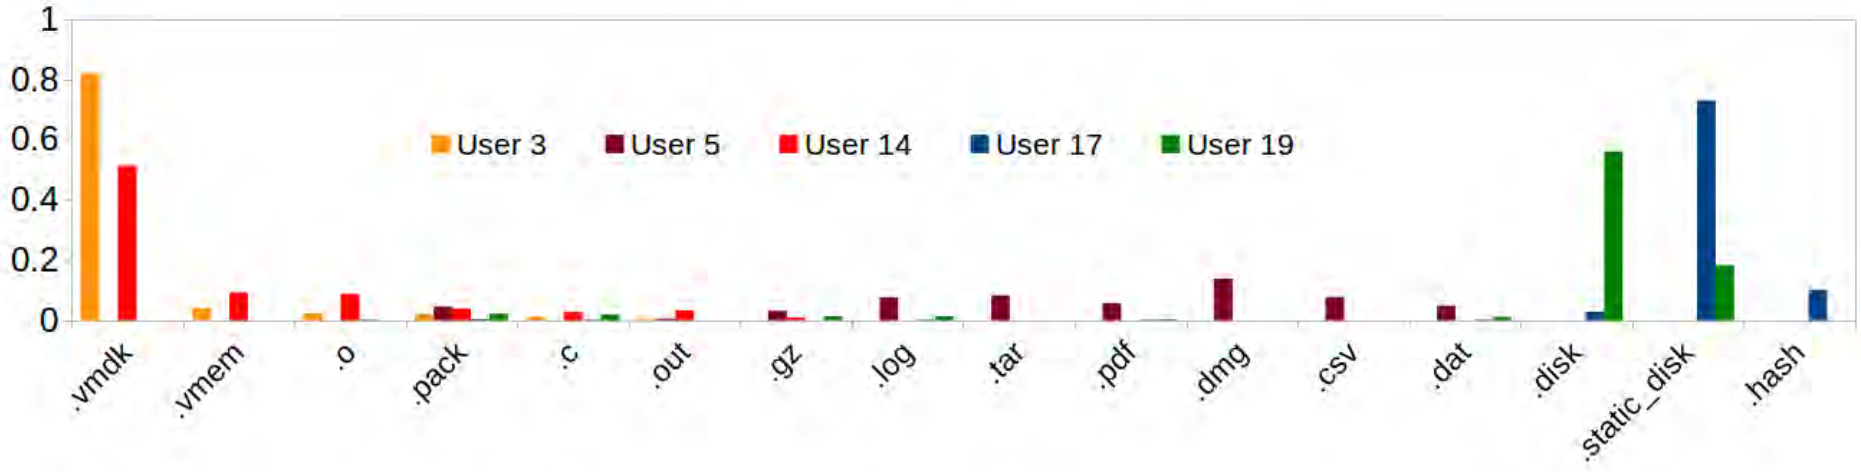

# User-Groups Analysis

- Redundancies among users vary significantly.
- Users can be divided into groups.

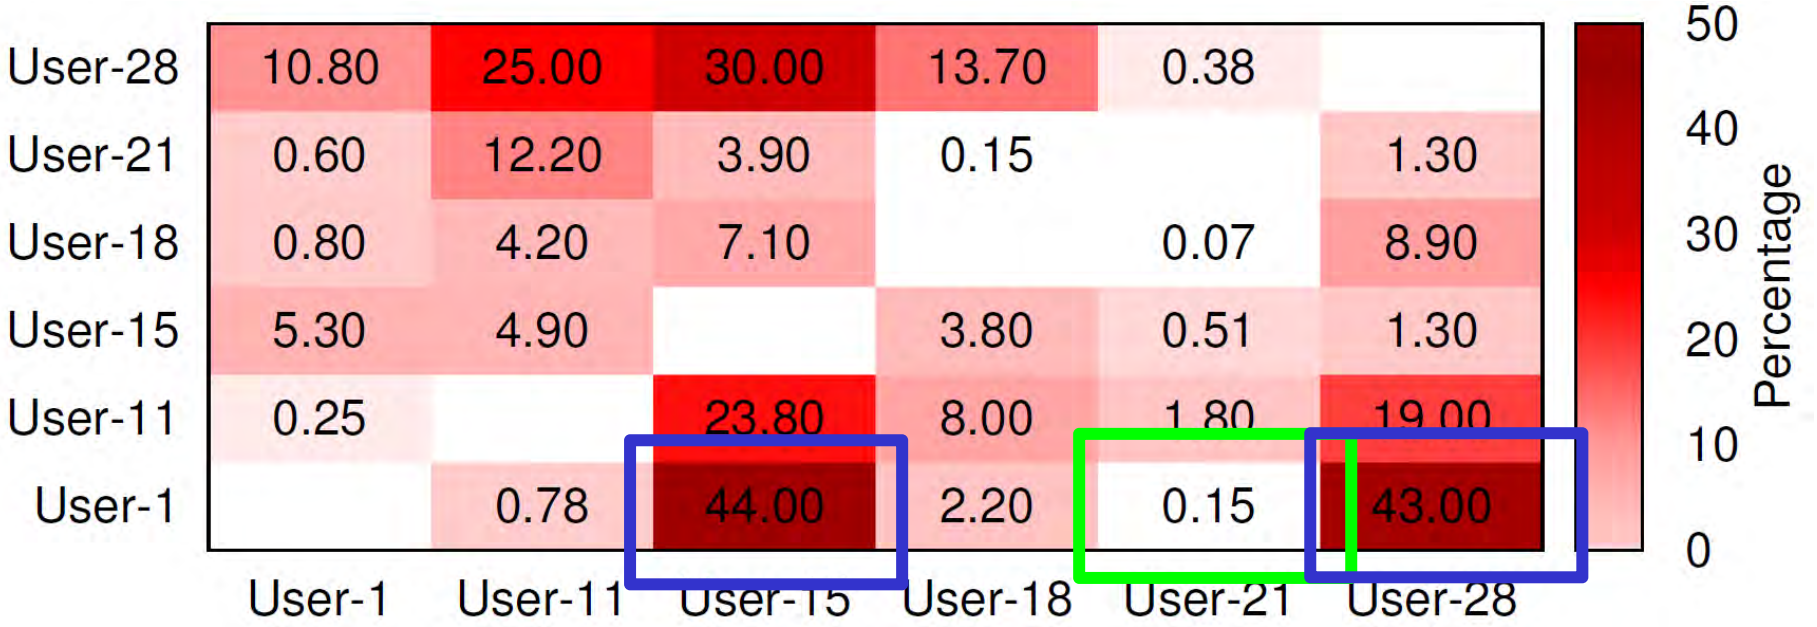

# Conclusion and Future Work

- Conclusion:
  - A long-term large-scale data-set collected and published online.
  - Data-set analyzed from whole data-set and users' perspective.
    - Large chunking size may performs better in deduplication ratio.
    - WFC is not suitable for our data-set.
    - File types have different deduplication ratio and chunk size sensitivity.
    - Data in different users vary in deduplication ratio and chunk sensitivity.
    - User shared data have much higher popularity than average.
  - Future work:
    - Cluster-deduplication.
    - Fragmentation in deduplication backup system.

# A Long-Term User-Centric Analysis of Deduplication Patterns


Zhen Sun,[1,2] Geoff Kuenning,[3] Sonam Mandal,[2] Philip Shilane,[4]
Vasily Tarasov,[5] Nong Xiao,[1,6] Erez Zadok[2]

*[1]HPCL, NUDT, China; [2]Stony Brook University; [3]Harvey Mudd College;*
*[4]EMC Corporation; [5]IBM Research – Almaden; [6]SYSU, China*

*Link for our data-set and tools: tracer.filesystems.org*

# Tools

- *Fs-hasher* : Collect snapshots
  - Scans a file-system everyday.
  - Collect file's meta-data and chunk's information.
  - Supports multiple chunking strategies, chunking size and hash functions.
- *Hf-state*: Parse snapshots
  - Prints snapshots in human-readable manner.
  - Multiple options to control it's output.

- Link: *tracer.filesystems.org*

# Data-set: FSL-*Homes*

- FSL-Homes: A long-term user-based backup data-set:
  - One snapshot per user per day.
  - Covered 33 users, >4000 snapshots, > 21months.
  - 7 variable chunking sizes + whole file chunking (WFC).
  - Rich meta-data which makes it suitable for multiple purpose studies.
  - 48 bit MD5 hash. (Hash collision rate < 0.004%)
- Limitation:
  - Real data is not stored.
    - Time/Space consuming to store all the data.
    - Unable for content-based analysis.
  - Some periods were not collected.
    - Data-collection is hard for many reasons.
- Link: http://tracer.filesystems.org/traces/fslhomes/
  - Data set will be periodically updated.

# Data-set: FSL-*Homes*

| Data Set | Homes |
|---|---|
| Organization | 1 snapshot per user per day |
| Total Size | 456TB |
| Start and end time | 03/09/2012 – 11/23/2014 |
| Number of users | 33 |
| Number of Snapshots | 4181 dailies (about 21 months) |
| Chunking methods | Content-defined Chunking, Whole File Chunking |
| Average Chunking Size | 2, 4, 6, 8, 16, 32, 64 and128KB |
| Hashing Method | 48 bit MD5 hash. (Hash collision rate < 0.004%) |
| Number of files | 130 million |
| Meta-data included | File pathname, size, atime, mtime, ctime, UID, GID, permission bits, device ID, inode number |

# User-groups Analysis (2)

- Redundant data shared by users in a group are largely similar.
- Chunks shared among users have much higher popularity than average.