

The Story of ChronoShare, or How NDN Brought Distributed File Sharing Back



Alex Afanasyev, Zhenkai Zhu, Yingdi Yu, Lijing Wang, and Lixia Zhang

IEEE MASS 2015 Workshop on Content-Centric Networking (CCN 2015)

October 19, 2015

Introduction

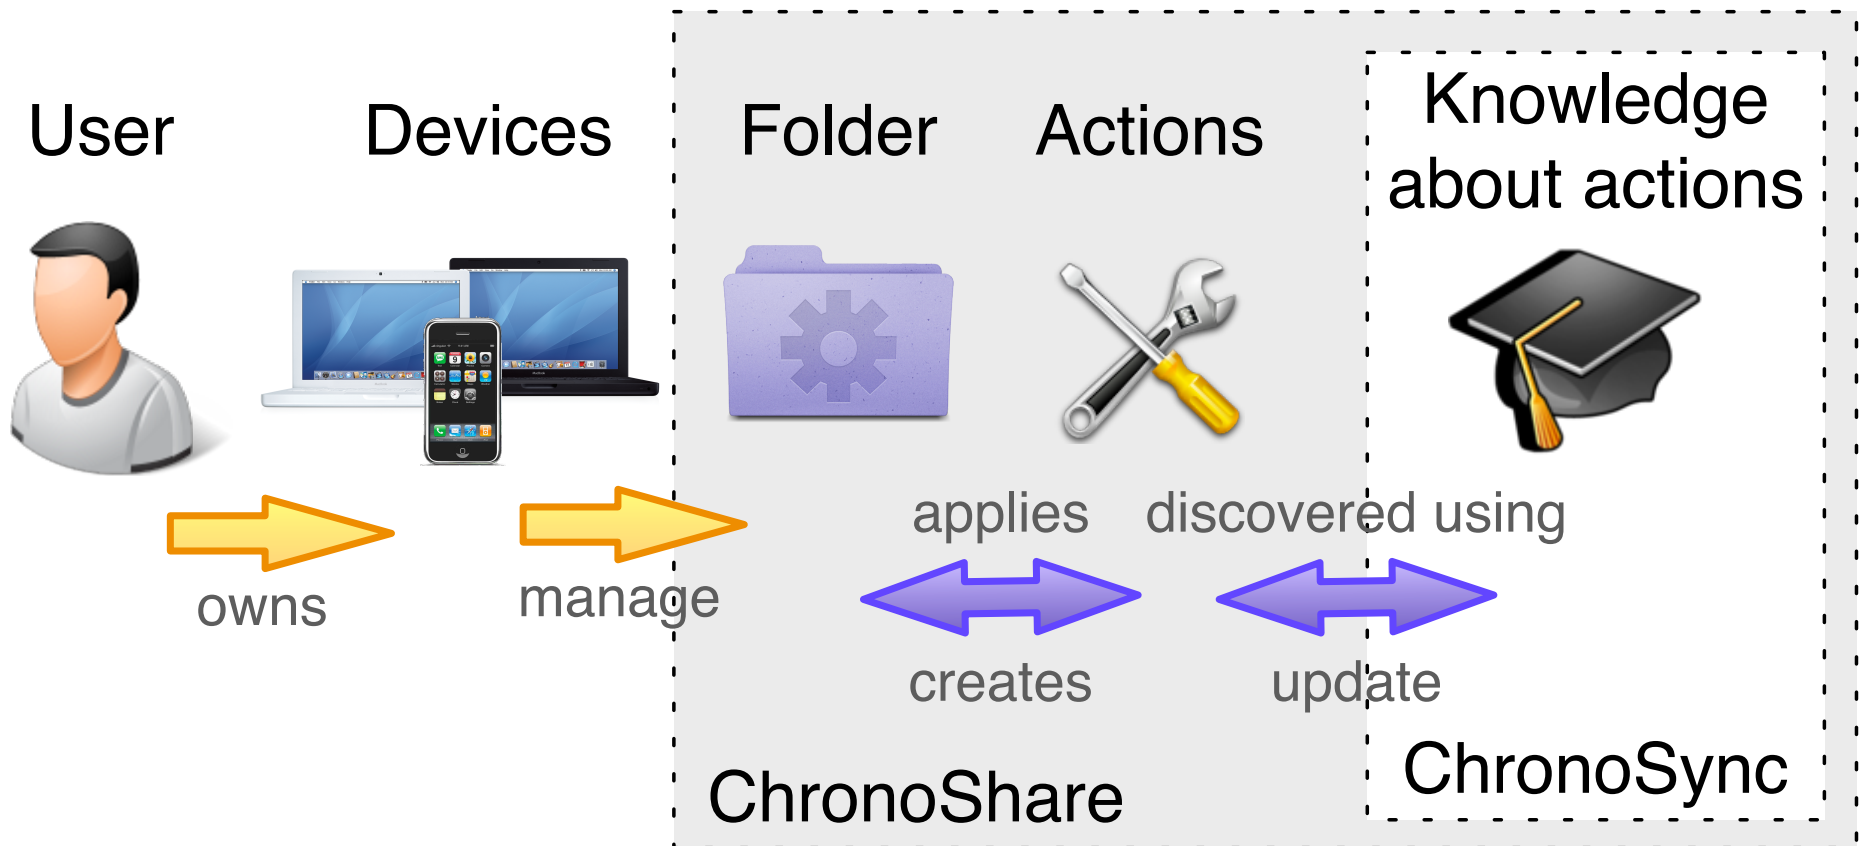
- ◇ File sharing is in great demand today
- ◇ File synchronization models
 - Centralized model
 - simple implementation, but
 - centralized control
 - single point of failure
 - delivery model mismatch
 - Peer-to-peer model
 - decentralized control
 - no single point of failure, but
 - delivery model mismatch
 - underlying and p2p topology mismatch



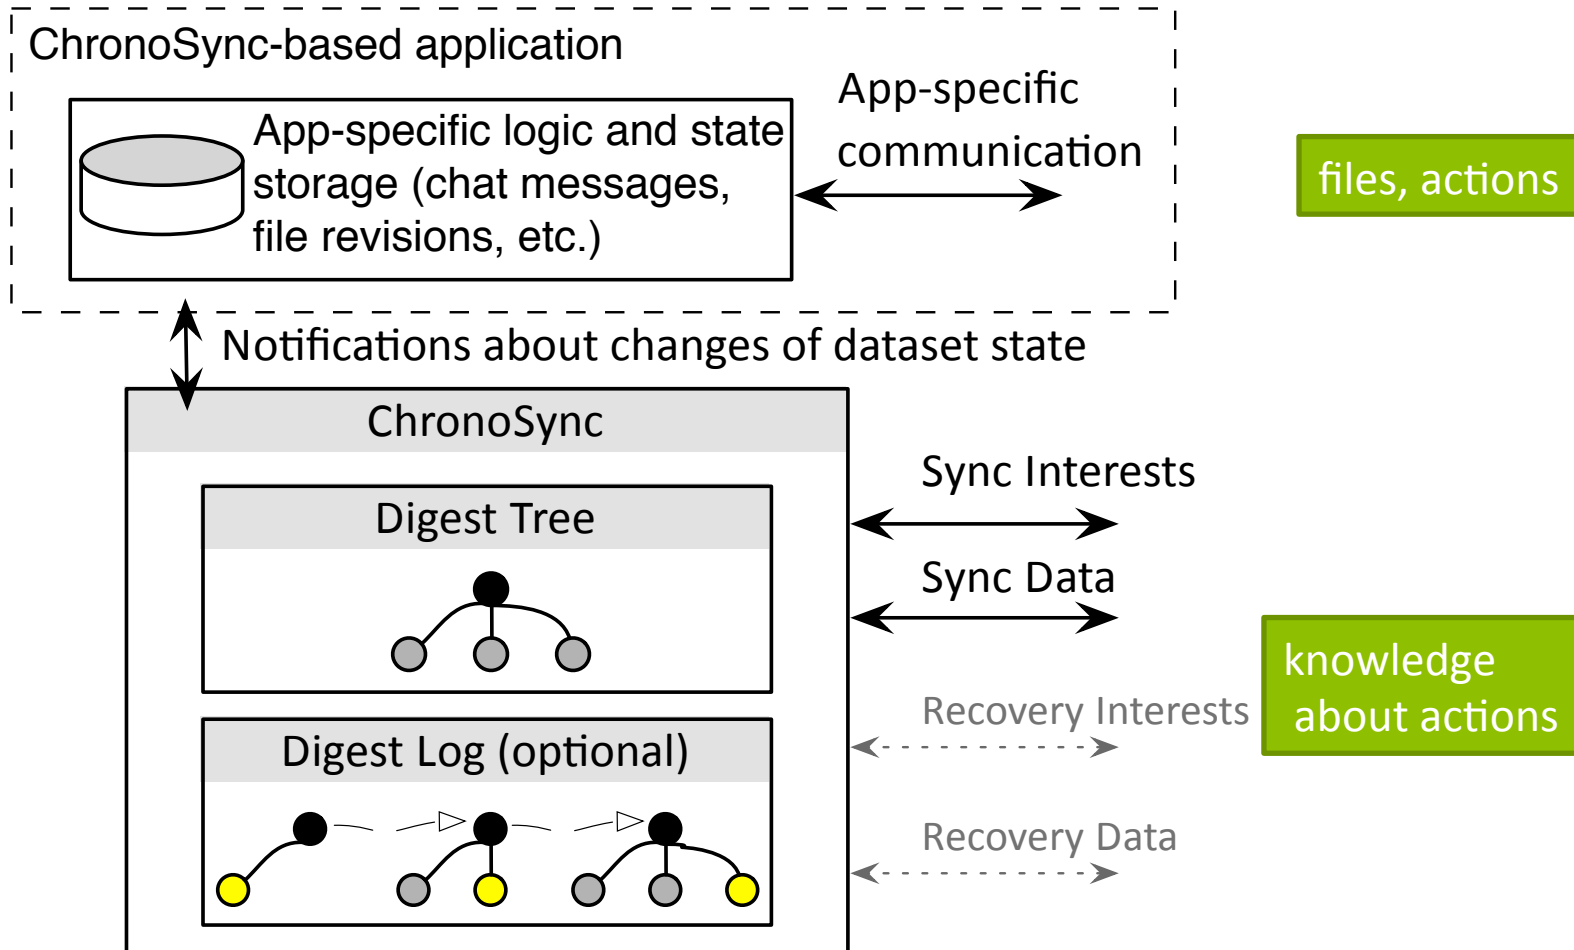
Can We Share Files in a True Distributed Way in NDN?

- ◇ Name-based data retrieval
 - expressed Interest for the can bring data
 - from the local database, local environment (WiFi, LAN)
 - from the original producer or any other user sharing the same folder
- ◇ Name-based rendezvous
 - Completely distributed, no need for central server
 - Dedicated instance can act as backup-storage/cloud
- ◇ Built-in data-oriented security
 - data-oriented provenance, integrity, group-access control
- ◇ Local communication is kept local
- ◇ Efficient data sharing
 - Multicast fashion
 - “Multi-peer-to-multi-peer” file sharing
- ◇ Support for device mobility and intermittent connectivity

ChronoShare: Distributed File Sharing in NDN



ChronoSync as a Distributed Transport Function



Z. Zhu and A. Afanasyev, "Let's ChronoSync: Decentralized Dataset State Synchronization in Named Data Networking," in Proceedings of the 21st IEEE International Conference on Network Protocols (ICNP 2013), Goettingen, Germany, October 2013.

ChronoSync: representation of knowledge about actions

◇ Actions

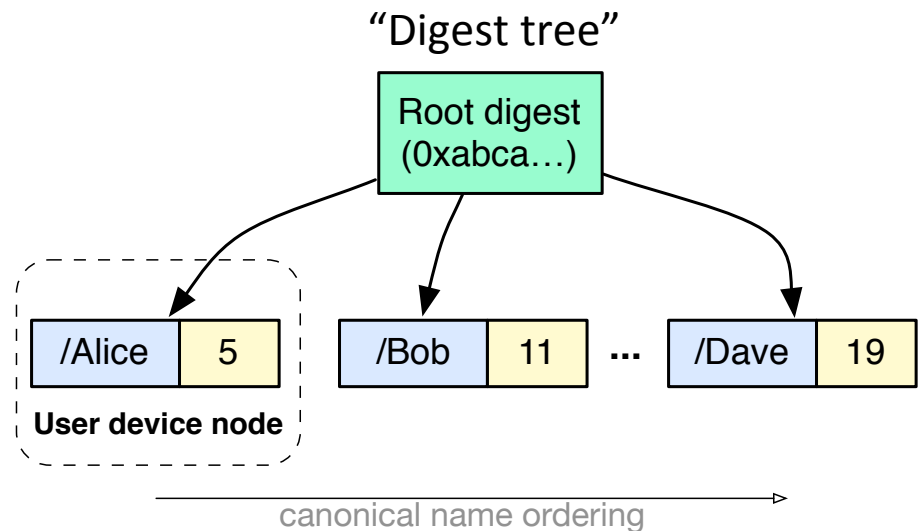
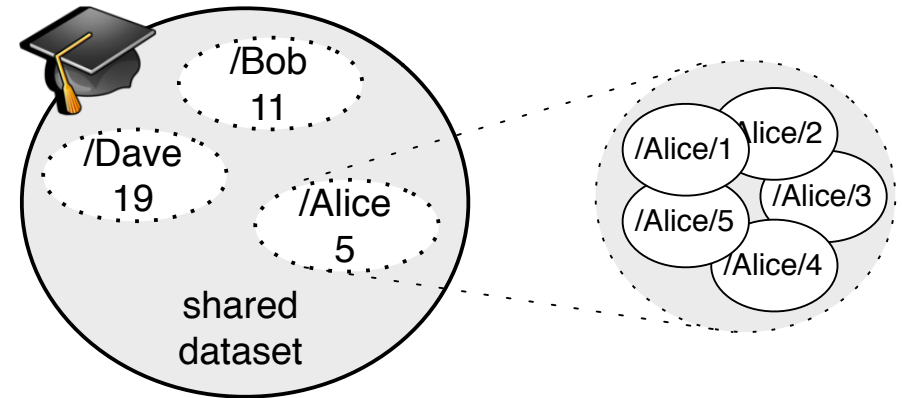
- **Operation on a shared folder**
- /Alice/1, /Alice/2, ...
- /Bob/1, /Bob/2, ...
- /Dave/1, /Dave/2

◇ User's device node

- **Latest user action -> identify all actions from the user (on a device)**
- /Alice/5 -> 5 actions from Alice
- /Bob/11 -> 11 actions from Bob
- /Dave/19 -> 19 actions from Dave

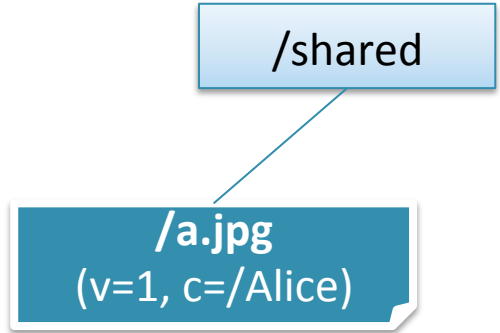
◇ State digest (root digest)

- identify state of the whole shared dataset
- **sha256(/Alice/5, /Bob/11, /Dave/19)**

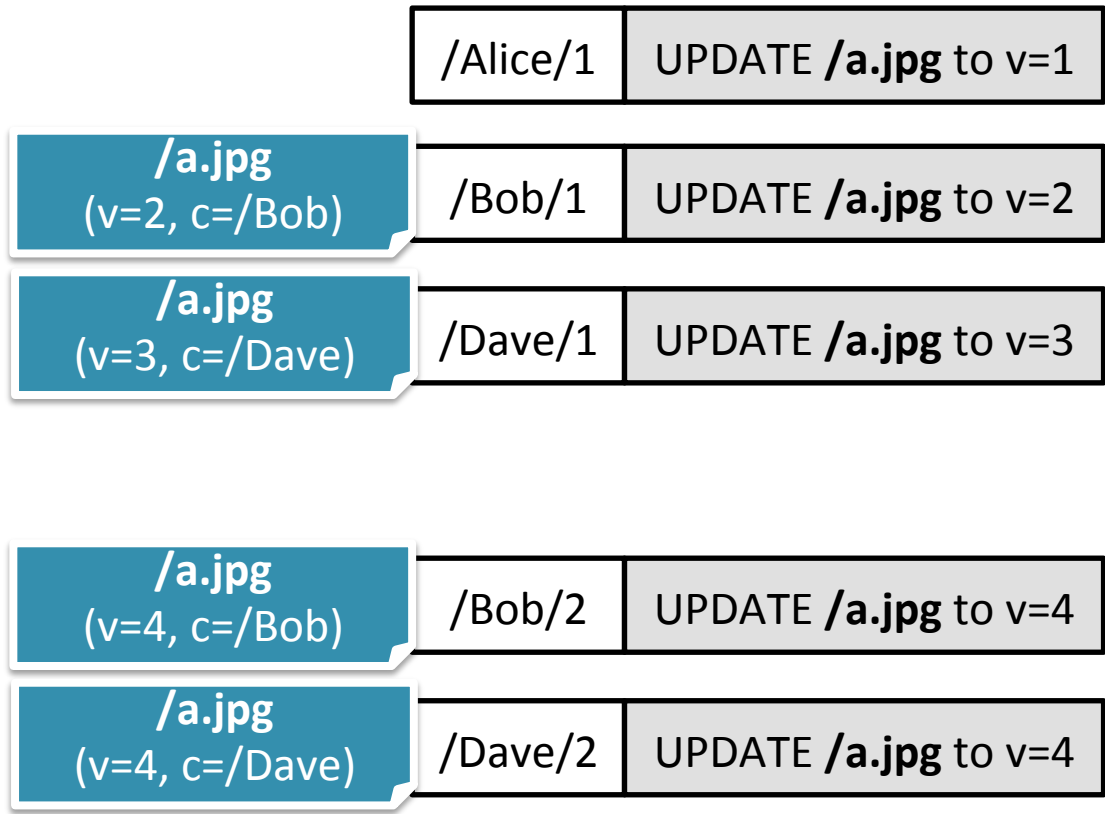


ChronoShare Operations (New/Update)

Shared Folder Tree

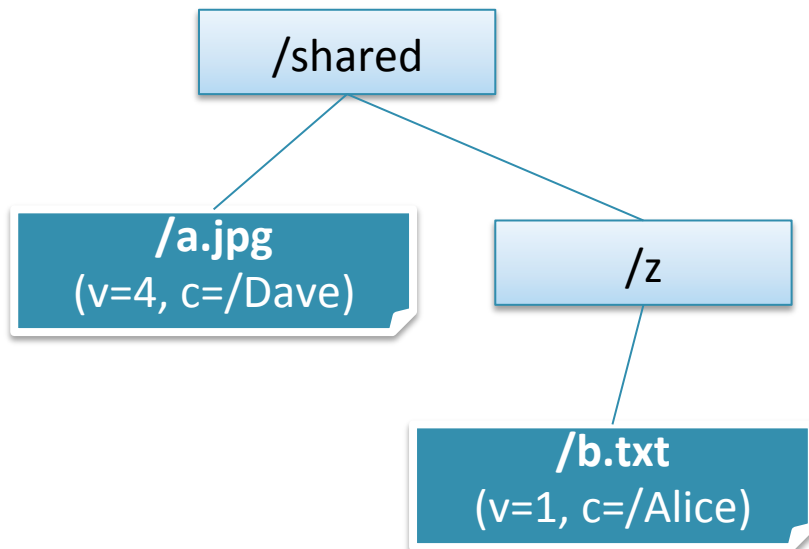


Shared Folder Actions




ChronoShare Operations (Delete)

Shared Folder Tree



Shared Folder Actions



ChronoSync Digest Tree Progress

/Alice/1	UPDATE /a.jpg to v=1
----------	----------------------

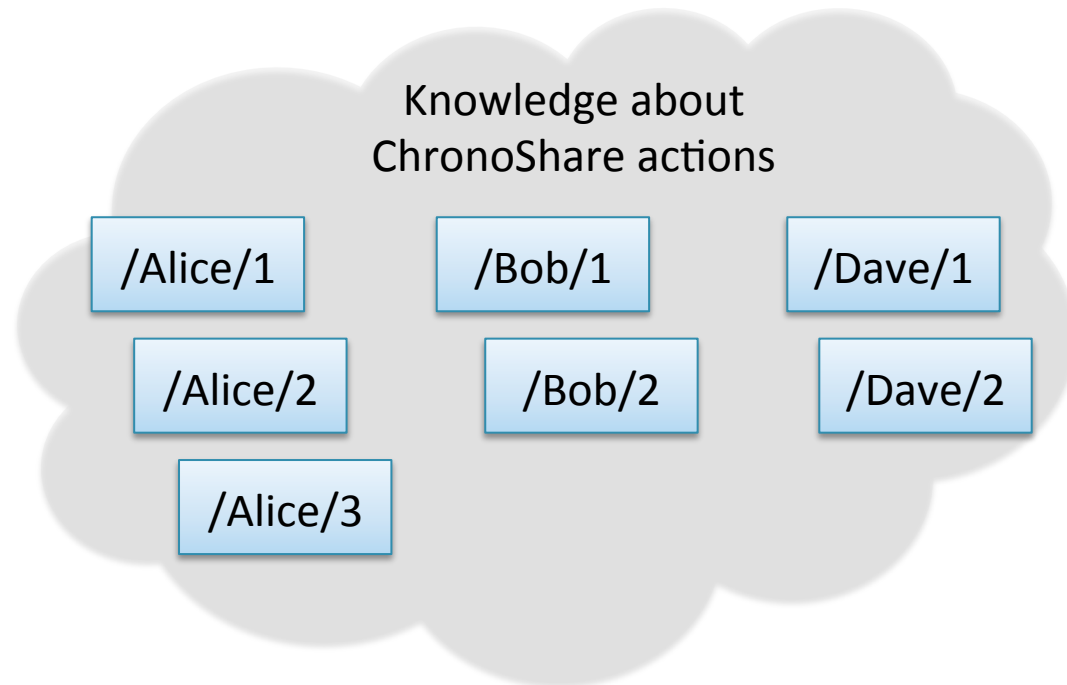

/Bob/1	UPDATE /a.jpg to v=2
--------	----------------------

/Dave/1	UPDATE /a.jpg to v=3
---------	----------------------

/Bob/2	UPDATE /a.jpg to v=4
/Dave/2	UPDATE /a.jpg to v=4

/Alice/2	UPDATE /z/b.txt to v=1
----------	------------------------

/Alice/3	DELETE /z/b.txt
----------	-----------------



ChronoShare Namespaces

- ◇ Separation of control/data/view namespaces
 - **Control namespace**
 - state (ChronoSync): who produced and how much data
 - actions: describe produced data
 - new files, file updated, file removals
 - **Data namespace**
 - files
 - **View namespace**
 - the same as file system name space
 - file names as user sees them

Control Namespace: State

- ◇ Naming pattern
 - ◇ **<name-rendezvous>/<app-name>/<folder-name>/<digest>**
- ◇ Naming components
 - ◇ **<name-rendezvous>**: NDN prefix to rendezvous ChronoShare participants (who will receive Sync Interests)
 - ◇ /ndn/multicast
 - ◇ **<app-name>**: one-component to demultiplex ChronoShare application interests
 - ◇ ChronoShare
 - ◇ **<folder-name>**: one-component to demultiplex ChronoShare interests for a specific shared folder
 - ◇ testing4
 - ◇ **<digest>**: root digest representing the shared folder state
- ◇ Example
 - ◇ **/ndn/multicast/ChronoShare/testing4/109f75ca81...**

Control Namespace: Actions

- ◇ Naming pattern
 - ◇ **<user's device name>/<app-name>/ACTION/<folder-name>/<seq-no>**
- ◇ Naming components
 - ◇ **<user's device name>**: NDN prefix under which user published data on the specific device
 - ◇ This is prefix under which user is authorized data publishing
 - ◇ Macbook of an Alex in UCLA: `/ndn/edu/ucla/alex/macbook`
 - ▷ Hierarchy in affiliation, not location
 - ▷ Requires authorization from `/ndn/edu/ucla` to publish under `/ndn/edu/ucla/alex` prefix
 - ◇ **<seq-no>**: number, representing sequence number of the action
- ◇ Example
 - ◇ `/ndn/edu/ucla/alex/macbook/ChronoShare/ACTION/testing4/_seq=42`

Data Namespace: Files

- ◇ Naming pattern
 - /<user's device name>/<app-name>/FILE/<version>/<segment>
- ◇ Naming components
 - <version>: version of the file produced by user's device
 - <segment>: segment number of the file (0 if file has only one segment)
- ◇ Example
 - /ndn/ucla.edu/zhenkai/ChornoShare/FILE/_v=1/_seg=143

Advanced Operational Concepts

- ◇ Each local action triggers ChronoSync sync'up
 - Eventual synchronization across all devices of all users
 - Remote users discover that a new action has been applied
- ◇ Devices request missing action items
 - All devices request the same piece of data
- ◇ Action items may result in fetching new/updated file data
 - Collection of action items builds a shared folder tree
- ◇ Local synchronizing policies
 - When and which file should be fetched
 - mobile devices may want to fetch only current version
 - stationary devices may fetch all the files, prioritizing the current version
 - When to facilitate heterogeneous devices and connectivity
 - Bluetooth, Wi-Fi, cellular

Handle Mobility and Disconnects

- ◇ Basic system picture
 - User/device names are in general location-independent
 - If a name prefix is in FIB, it is reachable, otherwise its location prefix should be in FIB
 - LINK object to reach one the users
- ◇ The focus is on getting desired data, instead of tracking device status or locations
 - Find device location prefix on a “need to know” basis
 - A device can report its current location prefix
- ◇ How it works: focus on file, not device
 - SYNC Interest brings back action data (=action+file name),
 - Use file-name to fetch, if no reply, send to broadcast channel, asking “who has this file?”
 - Whoever has the file replies with its location-prefix
- ◇ Handle device disconnect by fetching from cached copies

A. Afanasyev, C. Yi, L. Wang, B. Zhang, and L. Zhang, "SNAMP: Secure Namespace Mapping to Scale NDN Forwarding," in Proceedings of 18th IEEE Global Internet Symposium (GI 2015), April 2015.

Conclusions

- ◇ ChronoSync provides universal function for distributed communication
 - restrictions of ChronoSync (per-producer sequence number) does not prohibit more complex applications to be build on top of it
 - Use of an extra level of indirection
- ◇ NDN with the use of ChronoSync-based ChronoShare brings back truly distributed and secure file sharing
 - name-based multicast, mobile, and DTN-style data retrieval
 - name-based rendezvous
 - name-based communication scoping
 - data-oriented security

Thanks

◇ Questions?

◇ Email: aa@cs.ucla.edu