Analyzing scientific data-sharing patterns for in-network data caching

Elizabeth Copps[1], Alex Sim (advisor)[2], John Wu (advisor)[2]
[1]Middlebury College, [2]Lawrence Berkeley National Laboratory

BERKELEY LAB — WORKFORCE DEVELOPMENT & EDUCATION

U.S. DEPARTMENT OF ENERGY — Office of Science

ABSTRACT

The volume of data generated by new scientific projects is exponentially increasing, driving up network traffic and data delivery requirements. In-network regional data caching systems are being used to minimize the negative effects of the increased data volume. These systems reduce the number of data transfers from the source by storing files locally. When multiple users need to access the same file, it can be transferred once and then stored in the cache. Regional caches have been shown to save network traffic volume, reduce data latency, and improve overall application performance. In this work, an in-network XCache system is analyzed to determine how much traffic volume is saved over time. The researchers also studied the system to predict how new caches can be distributed to further decrease traffic volume and increase data availability. The cache system was shown to reduce network traffic by an average factor of 2.91.

BACKGROUND

- Study uses data from Southern California regional cache of 14 XCache installations
- Southern California regional cache handles data from Large Hadron Collider (LHC), which is expected to produce 30x more data in 2028 than 2018
- Study period: June - August 2020

What happens when users submit a request to the regional cache?

- File request goes to regional redirector
- Requests sent to XCache nodes based on available disk space on each node
- If file is already cached, immediately sent to the user (share/cache hit)
- If not, file transferred to the user from the data origin (transfer/cache miss)

	UCSD	Caltech	ESnet
Number of Nodes	11	2	1
Node names	Cache 00 - Cache 10	Xrd1, Xrd2	ESnet
Disk Capacity	24 TB each	180 TB each	40 TB

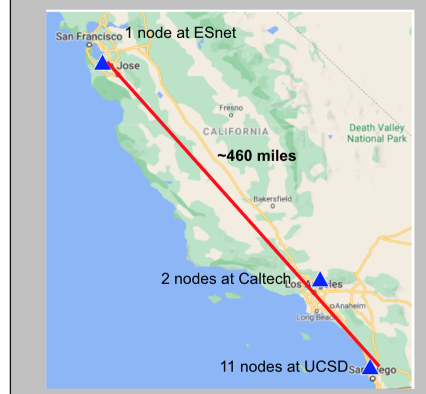

Table 1 (above): Summary of nodes in the regional cache
Figure 1 (left): Map of regional cache

RESEARCH QUESTION

How does an in-network regional caching system save network traffic volume? How can caches be distributed to increase data availability?

DATA

	Number of accesses	Data transfer size (GB)	Shared data size (GB)	Share Percentage
June 2020	1,804,697	532,037.7	818,956.9	60.62%
Jul. 2020	1,426,585	354,452.8	764,351.3	68.32%
Aug. 2020	995,324	249,583.5	586,188.8	70.14%
Total	4,226,606	1,136,074.0	2,169,497.0	65.63%

NODE DOWNTIMES

- Share count decreased by avg. of 264.59 shares/hr, while share size decreased by avg. of 149 GB/hr.
- Transfer count increased by avg. of 67.81 transfers/hr, while transfer size increased by avg. of 6.76 GB/hr.
- Remaining nodes evenly split the loads of nodes that had gone down

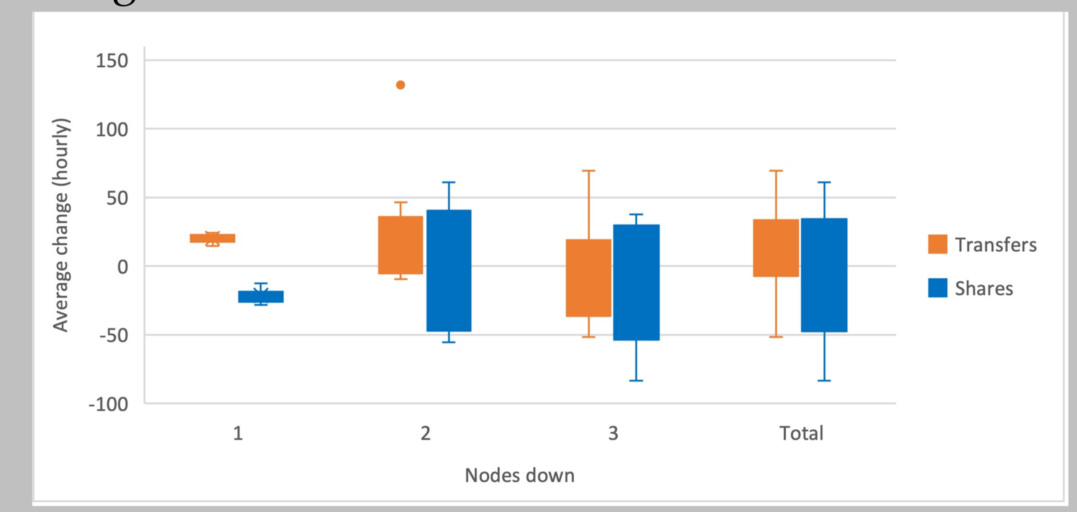

Figure 10: Box and whisker plot of the hourly transfer and share count changes after node downtimes. 'Nodes down' refers to how many nodes were down during the studied period.

RESULTS

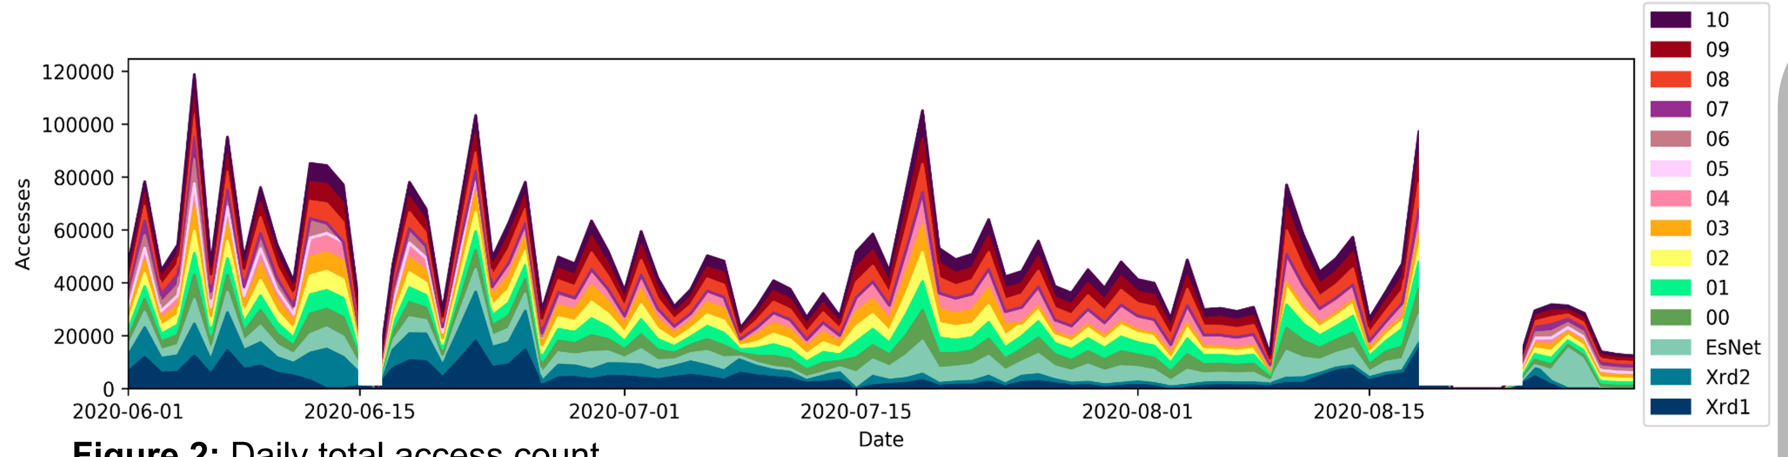

Figure 2: Daily total access count

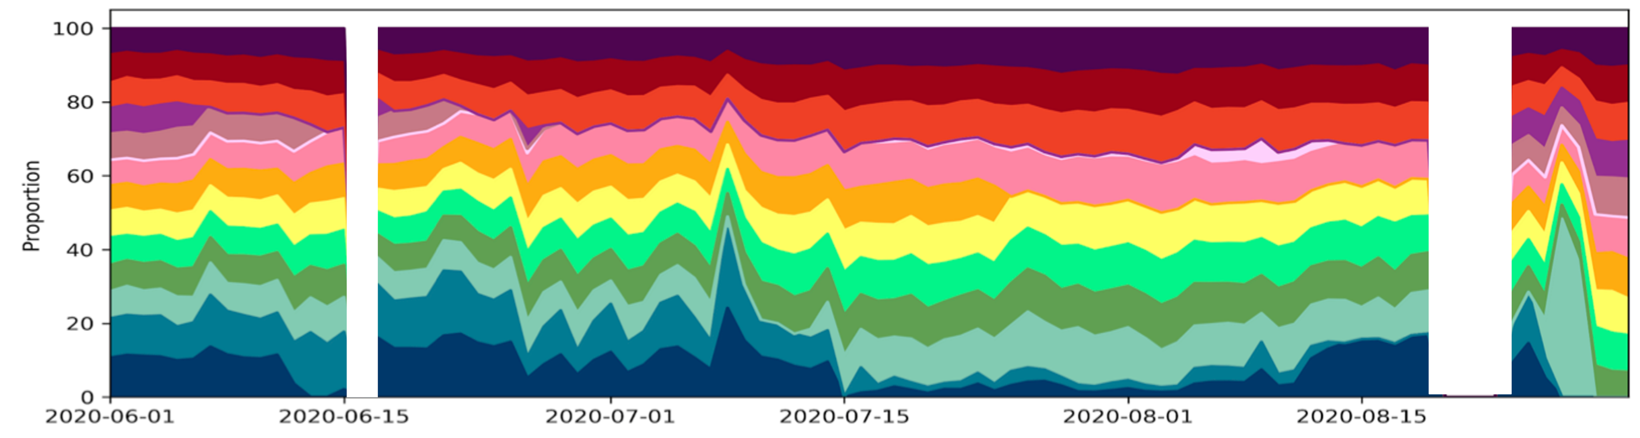

Figure 3: Daily proportion of cache hits on each node

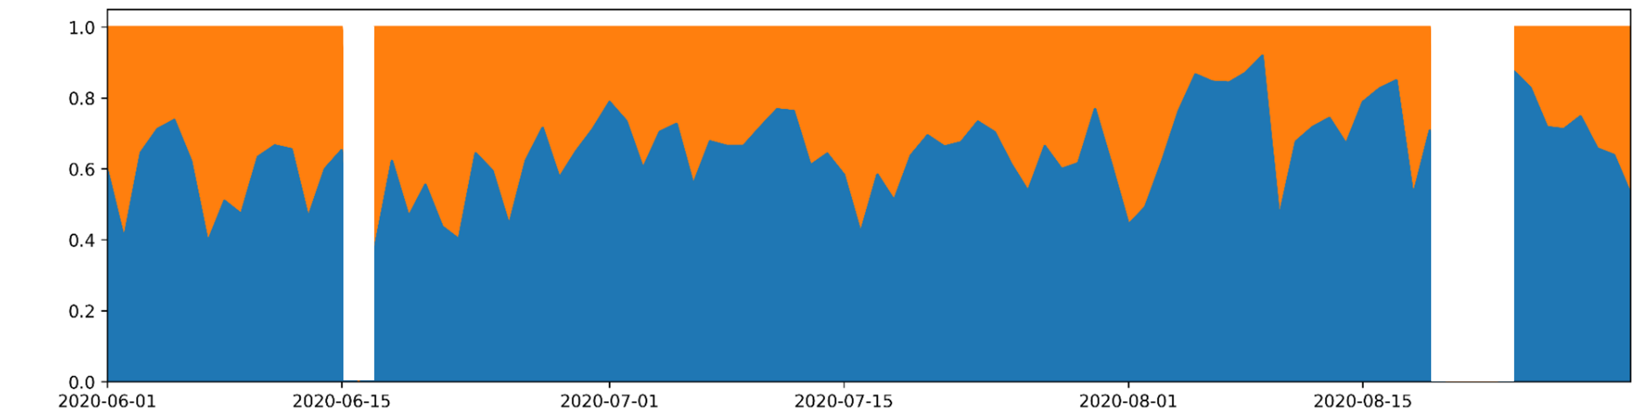

Figure 4: Daily proportion of total accesses that are shares (blue) and transfers (orange)

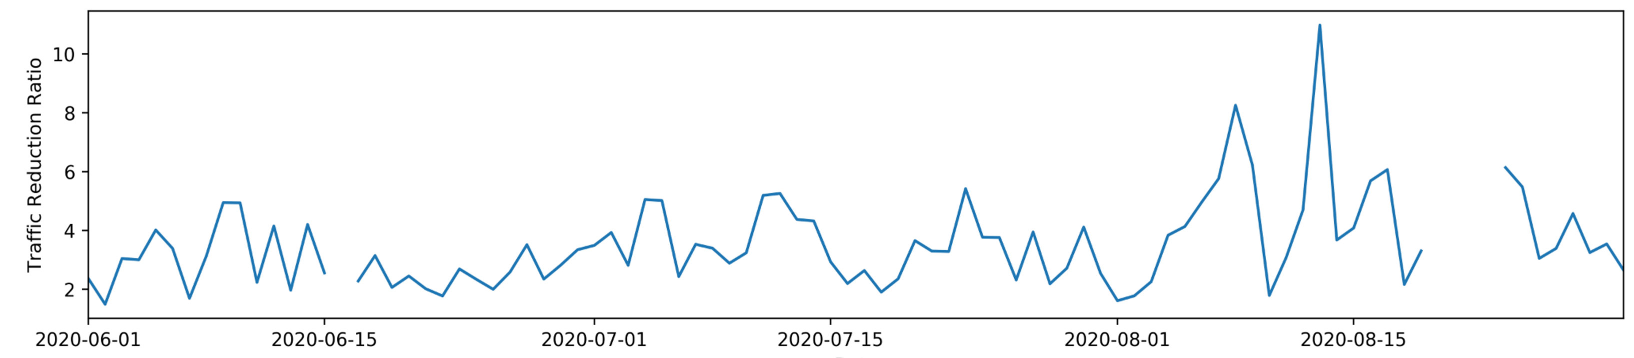

Figure 5: Network traffic reduction ratio = (Total access size / Total transfer size)

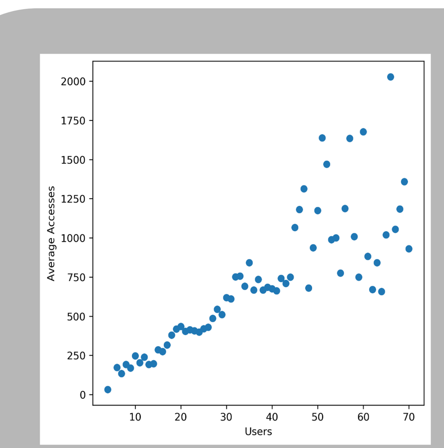

Figure 6: Hourly users vs. average hourly accesses

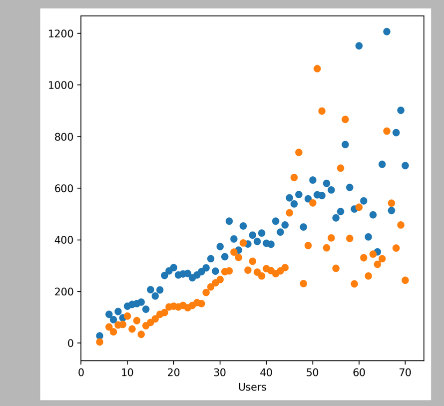

Figure 9: Hourly users vs. number of shares (blue) and transfers (orange)

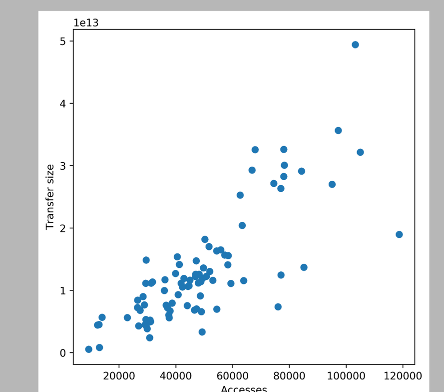

Figure 7: Hourly accesses vs transfer size

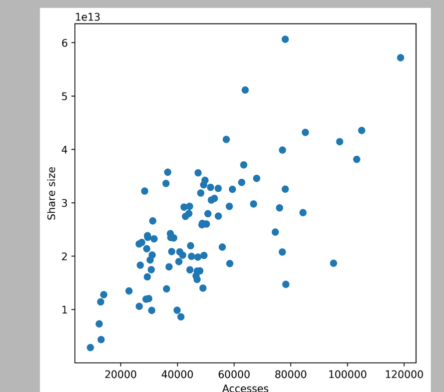

Figure 8: Hourly accesses vs share size

DISCUSSION

- Larger nodes have <1.5x as many accesses as other nodes at their peak, despite having 7x the disk capacity (Fig. 2)
- Accesses and users have linear relationship: as users double, accesses double (Fig. 6)
- Total daily transfer size is less than doubled when accesses double (Fig. 7)
- Total daily share size more than doubles when daily access count doubles (Fig. 8)
- Percentage of shared data and traffic reduction ratio gradually increases throughout study period, peaks in mid August (Figs. 4 & 5)
- Average traffic reduction ratio = **2.91** (Fig. 5)

CONCLUSIONS

- Nodes with larger capacity do not take proportionally larger loads off of the network
- Larger nodes are able to cache files for longer, increasing the traffic reduction rate over time
- Shared size increases more than transfer size as network load increases. This is because users are more likely to hit cached files as the cache grows
- Node downtimes cause an increase in transfers and a decrease in shares, evenly affecting the remaining nodes
- Adding nodes expected to reduce loads on other nodes evenly and to increase the proportion of requests that are already cached and do not need to be transferred

FURTHER READING

ACKNOWLEDGMENTS

Thank you to my mentors at Lawrence Berkeley National Laboratory, Alex Sim and John Wu. Thank you also to the Scientific Data Management (SDM) Group at Berkeley Lab. This work was supported in part by the U.S. Department of Energy, Office of Science, Office of Workforce Development for Teachers and Scientists (WDTS) under the Science Undergraduate Laboratory Internship (SULI) program.