Argonne NATIONAL LABORATORY

Towards Optimizing Large-Scale Data Transfers with End-to-End Integrity Verification

Raj Kettimuthu

Argonne National Laboratory & University of Chicago

Si Liu, Xian-He Sun, Illinois Institute of Technology

Eun-Sung Jung, Jongik University

U.S. DEPARTMENT OF ENERGY

Exploding data volumes

Astronomy

MACHO et al.: 1 TB
Palomar: 3 TB
2MASS: 10 TB
GALEX: 30 TB
Sloan: 40 TB
Pan-STARRS: 40,000 TB

Climate

2004: 36 TB
2014: 3,300 TB
2020: 100+ EB

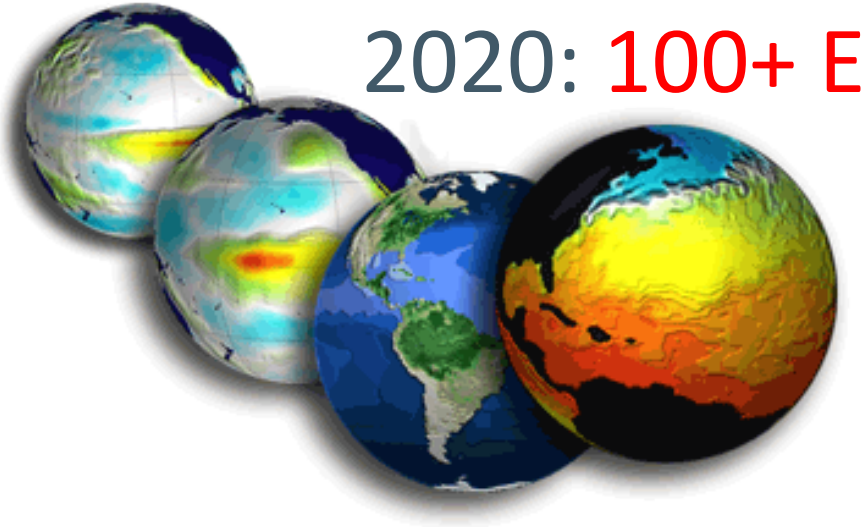

Genomics

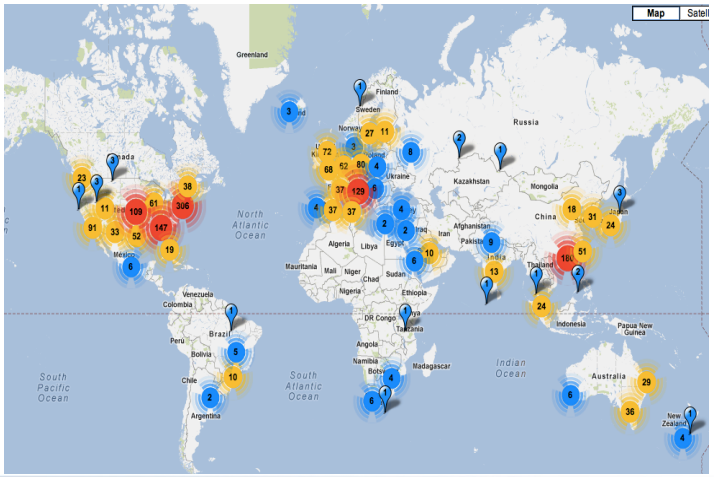

10^5 increase in data volumes in 6 years

End-to-end wide-area data transfers

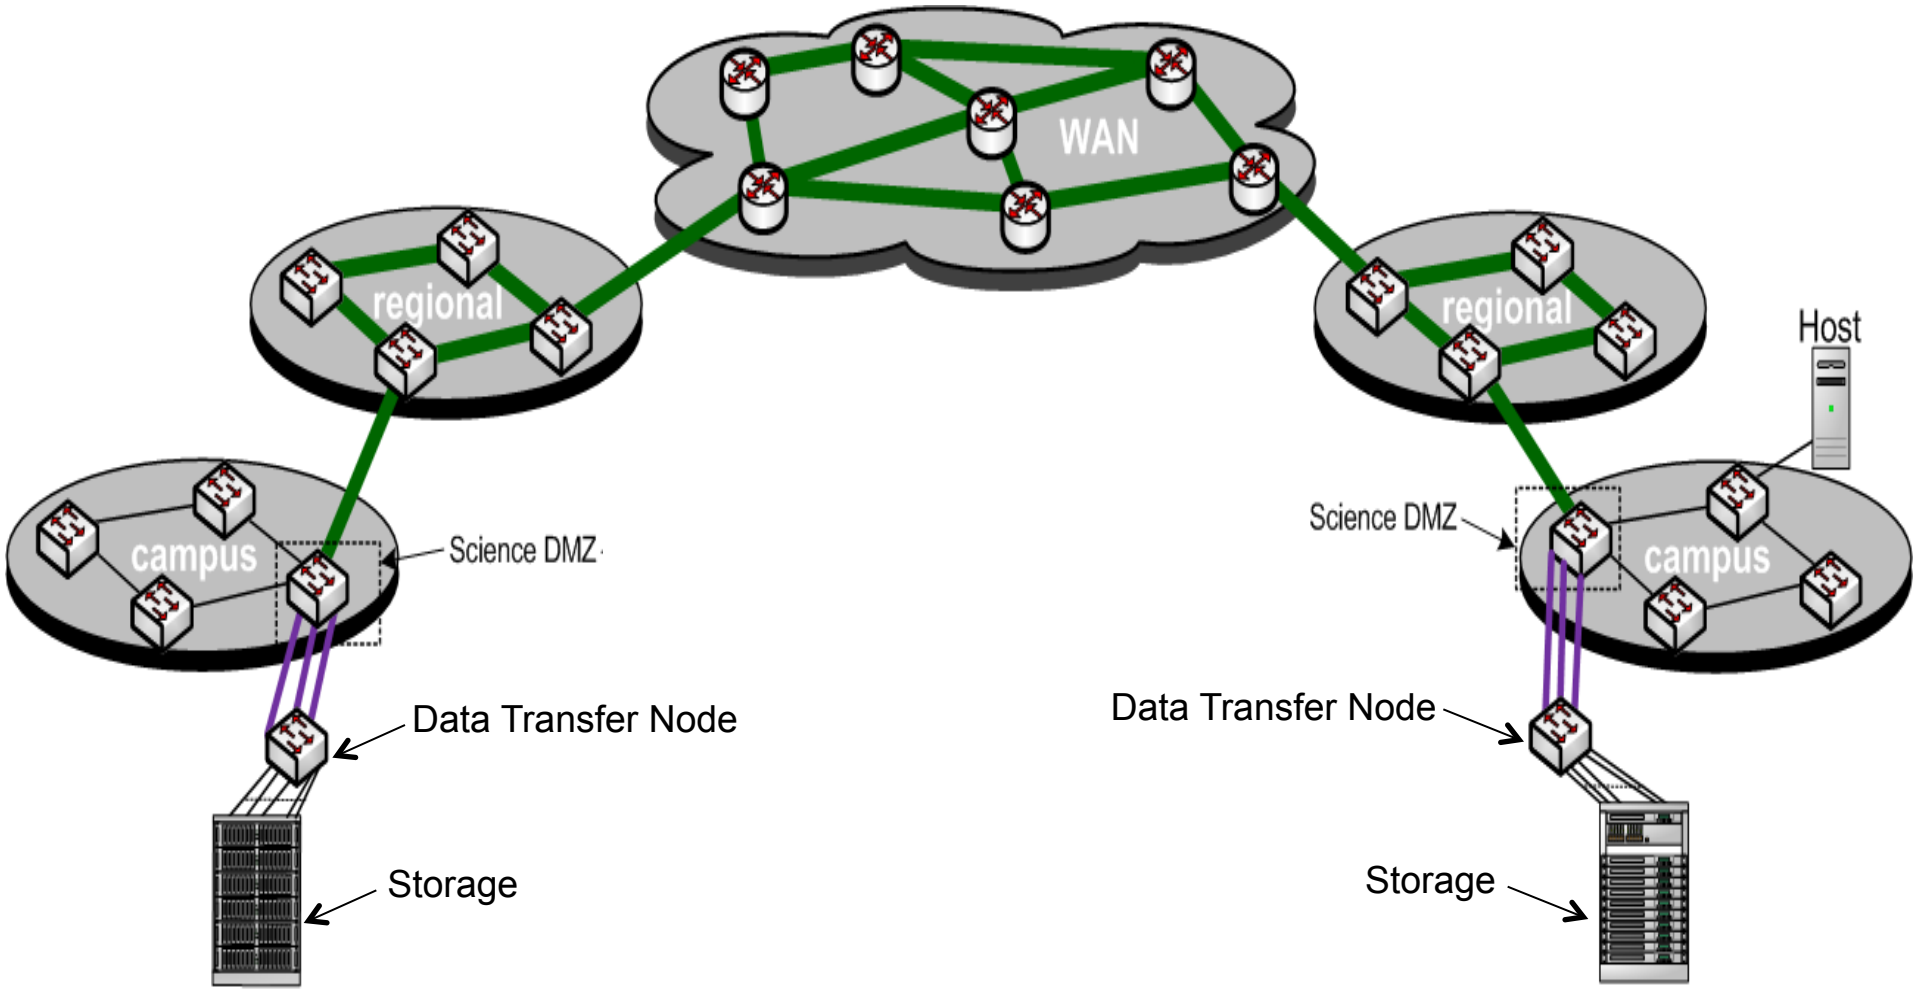

Pipeline Transfer and Checksum

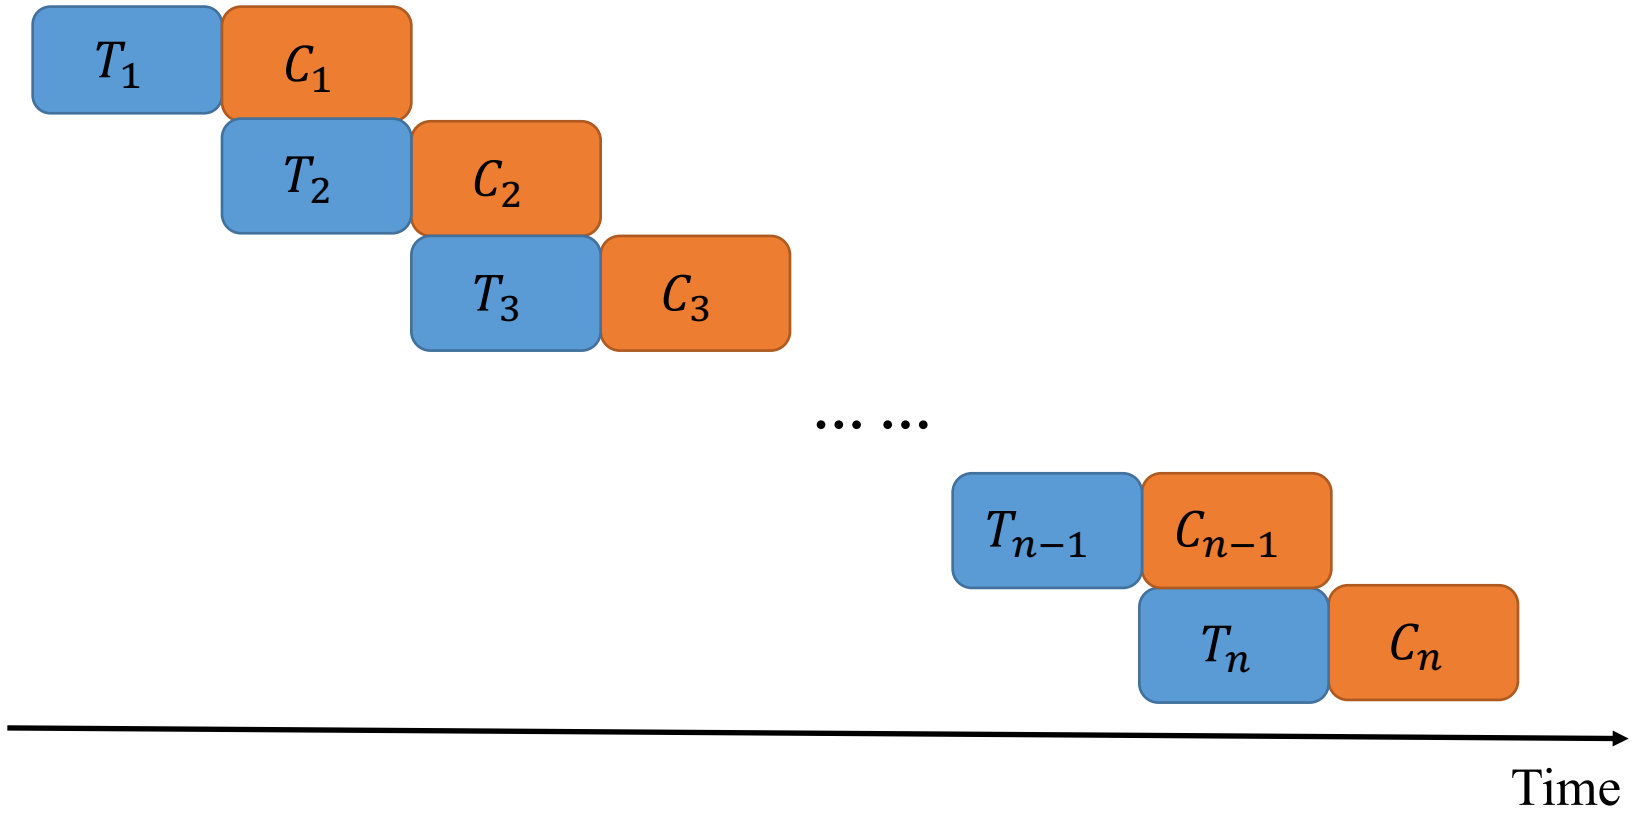

Pipelining Data Transfer and End-to-End Data Integrity Check

- Pipelining
  - File-level pipelining: overlap a file transfer and a file integrity check
  - Block-level pipelining: overlap a block transfer and block data integrity check
    - Block size is less than the average file size in a dataset
- Analytical Modeling
  - **t**: Transfer time of 500MB data **c**: Checksum time of 500MB data

Case	Dataset	Block-level Pipeline				File-level Pipeline	File Sequential
		100MB	500MB	1GB	2GB		
Transfer-Dominant	20-10GB	$400 \times t + 1/5 \times c$	$400 \times t + c$	$400 \times t + 2 \times c$	$400 \times t + \|4 \times c$	$400 \times t + 20 \times c$	$400 \times (t + c)$
$t > c$	10GB-500MB	$210 \times t + c$	$210 \times t + c$	$210 \times t + c$	$210 \times t + c$	$200 \times t + 201 \times c$	$210 \times (t + c)$
Checksum-Dominant	20-10GB	$400 \times c + 1/5 \times t$	$400 \times c + t$	$400 \times c + 2 \times t$	$400 \times c + 4 \times t$	$400 \times c + 20 \times t$	$400 \times (c + t)$
$t < c$	10GB-500MB	$208 \times c + 51/5 \times t$	$210 \times c + t$	$210 \times c + 2 \times t$	$201 \times c + 40 \times t$	$200 \times t + 201 \times c$	$210 \times (c + t)$

- Enhancing Block-level Pipelining
  - Based on the analysis, the best performance can be achieved when the data transfer time is close to the data checksum time
  - Checksum-Dominant case: reduce the data checksum time (Current Work)
  - Transfer-Dominant case: reduce the transfer time (Future Work)

Block-level pipelining -- Results

- Results on Cooley

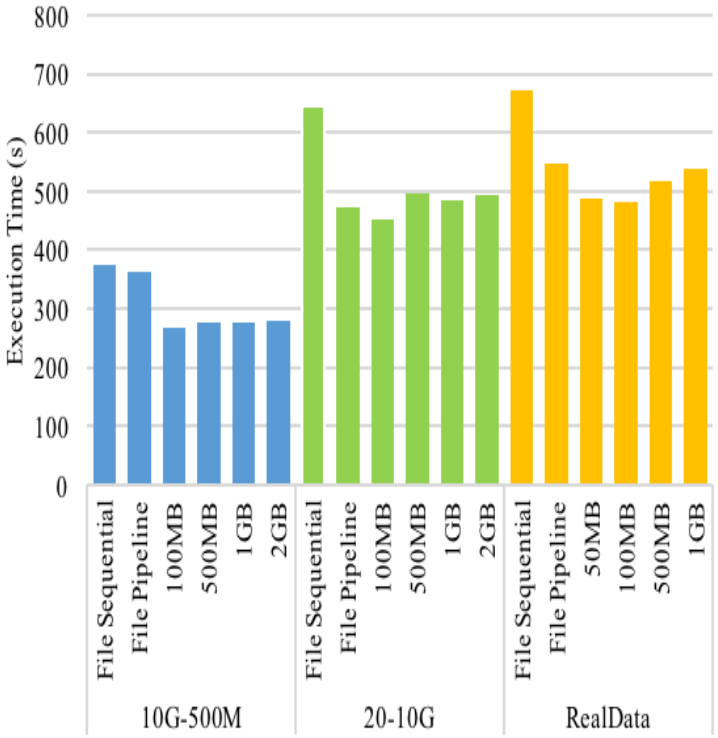

- Results on Rain

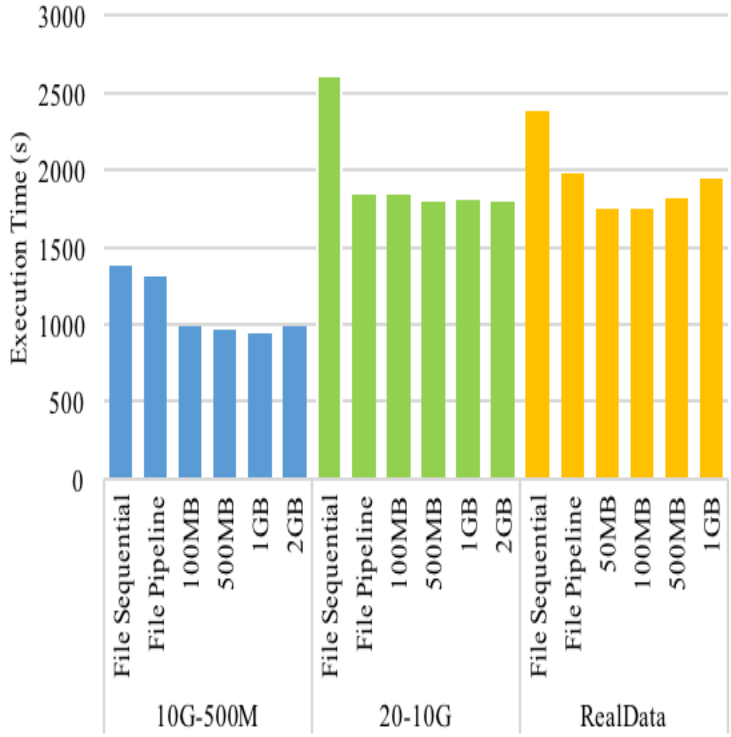

Block-level Pipelining - Perfect Pipeline

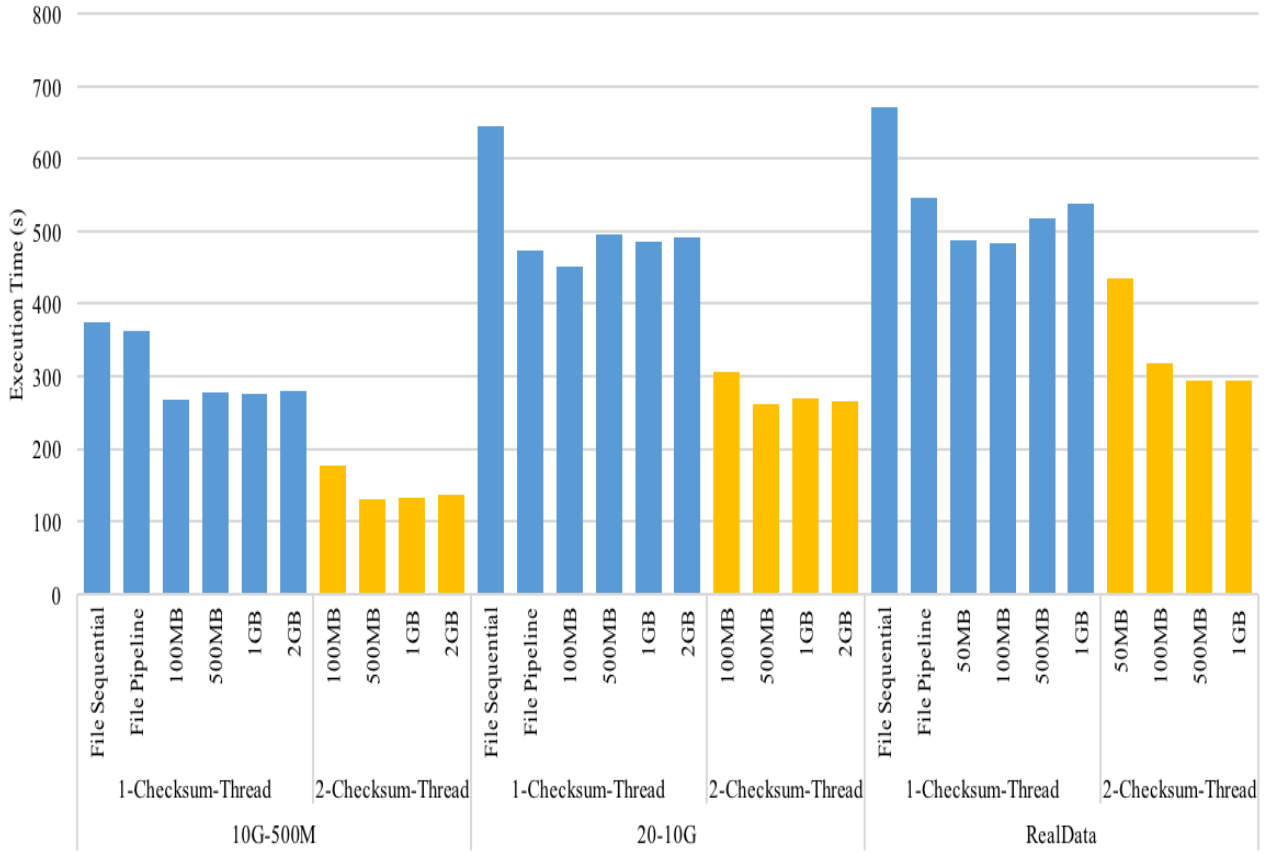

Comparison of the performance of 1-Checksum-Thread and 2-Checksum-Thread on Cooley

Questions