Explaining Wide Area Data Transfer Performance

By: Zhengchun Liu, Prasanna Balaprakash, Rajkumar Kettimuthu and Ian Foster

Presented by: Zhengchun Liu

April 11, 2017, Chicago

Motivation

Armed with a large collection of Globus transfer records, and experiments performed in the ESnet testbed environment, we want to:

- Extract factors that affect the transfer performance based on domain knowledge, and study their importance (***explanation***);
- Predict transfer performance by using data-driven model (***prediction***);
- Model based performance optimization (***optimization***, future work).

Outline

- Background & Motivation;
- Which factors are affecting the transfer performance?
- Deriving features from log to explain transfer performance.
- Make prediction by using derived features.
- Conclusion and future work.

What affect transfer performance?

4 kinds (3 known and 1 unknown):

For a given endpoint pair:

1) Tunable transfer parameters, e.g., concurrency, parallelism and pipeline;
2) Transfer file characteristic, e.g., file size;
3) Contentions from other simultaneous Globus transfers (known to us) and,
4) Contentions from other programs (unknown to us), e.g., sharing PFS with SC, network.

What affect transfer performance?

Tunable parameters:

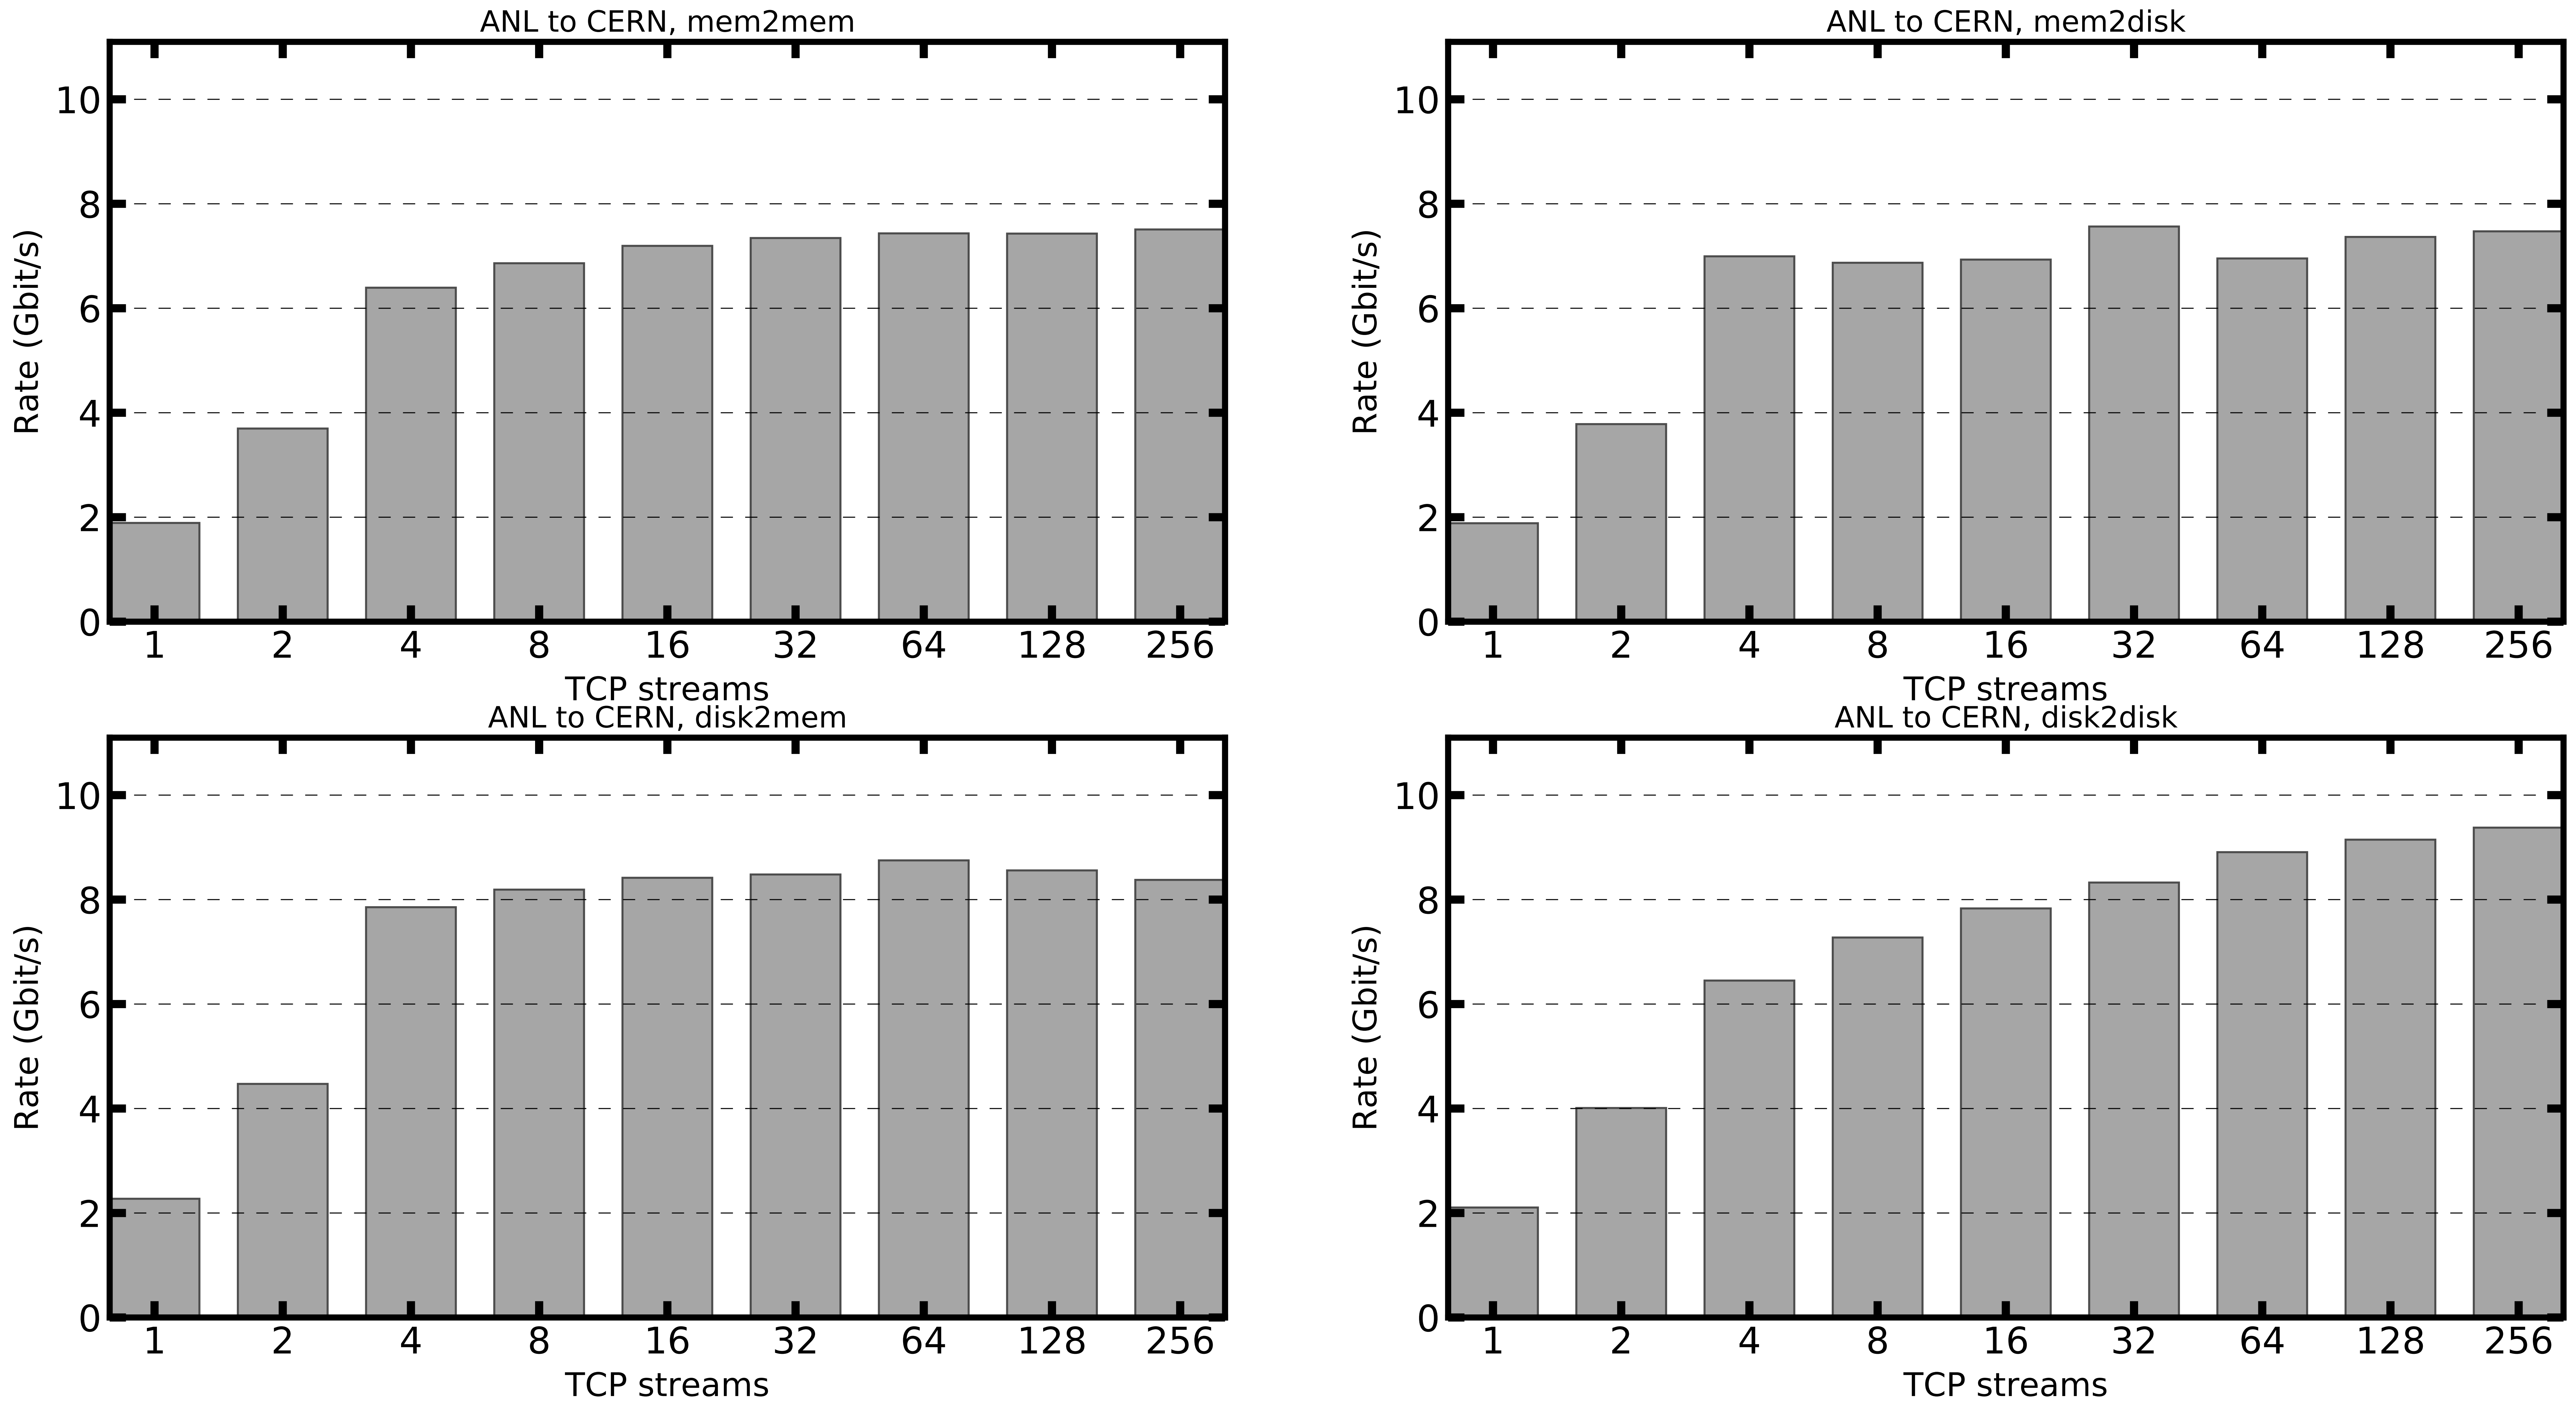

What affect transfer performance?

File characteristics:

File characteristics versus transfer performance.

Large transfers with big average file size are more likely to have better performance.

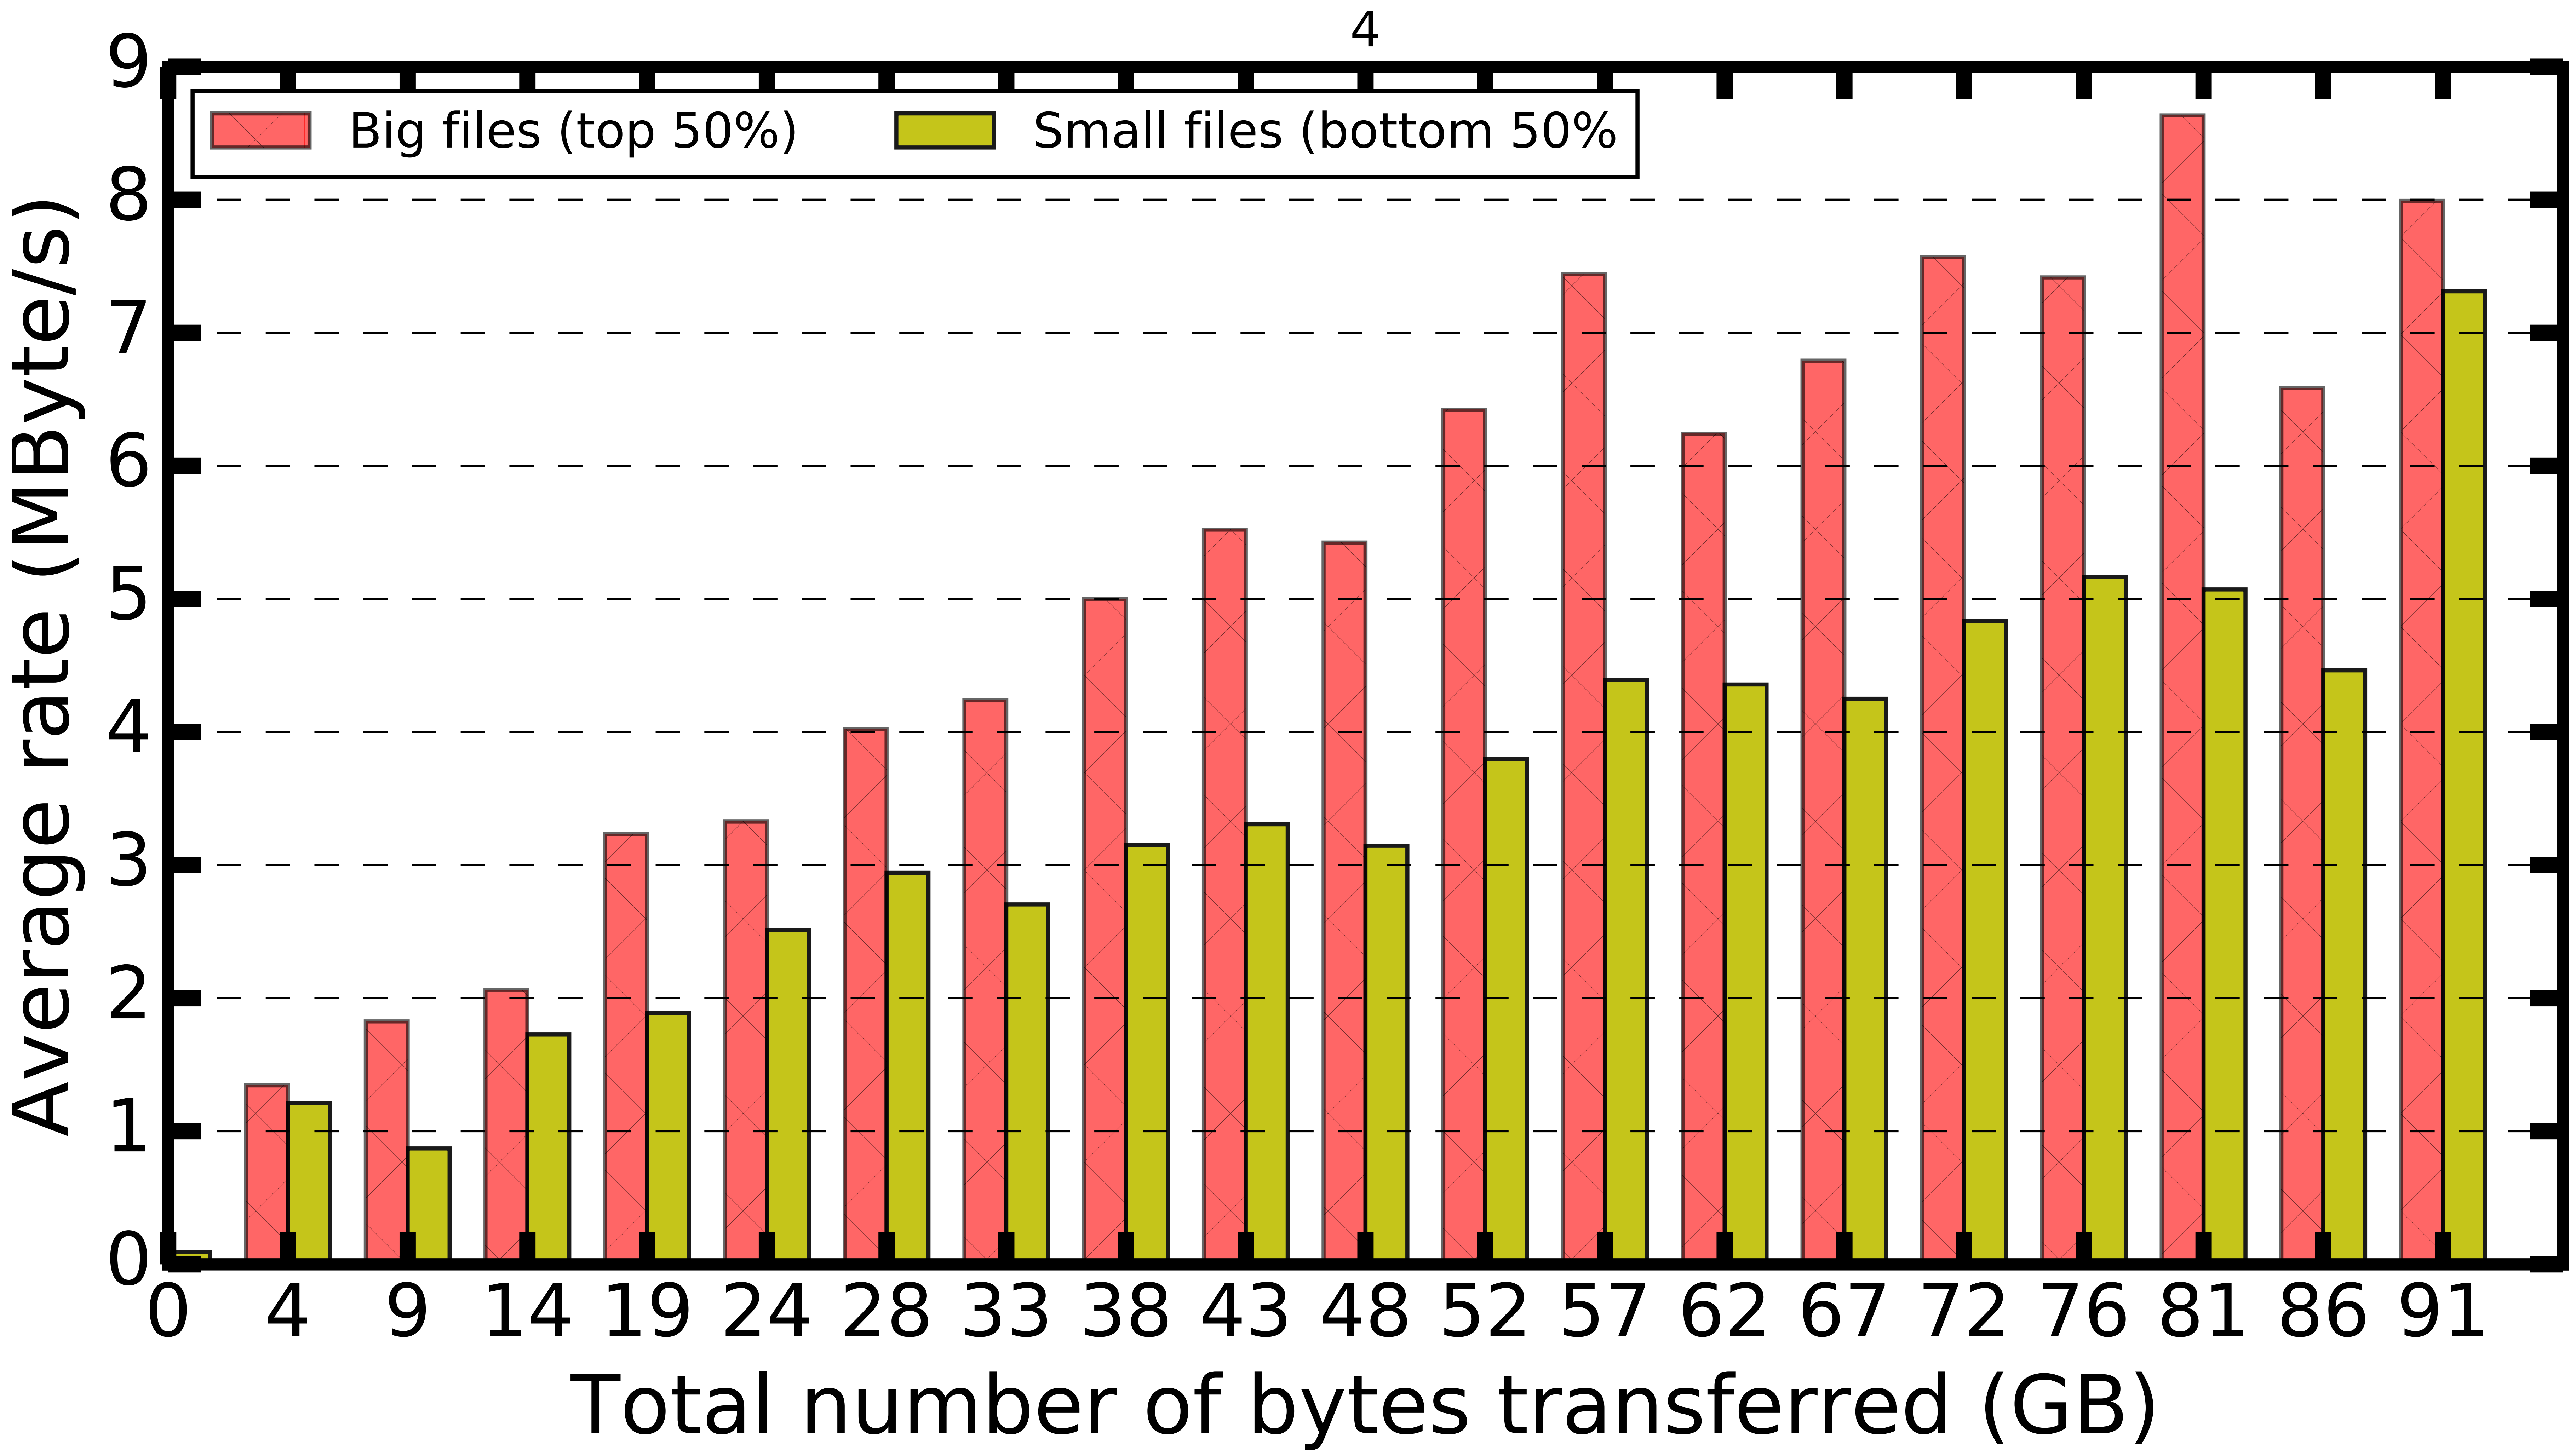

What affect transfer performance?

- ☑ *Tunable transfer parameter, e.g., concurrency, parallelism and pipeline;*
- ☑ *Transfer file characteristic, e.g., file size;*
- ☑ **Contentions from other Globus transfers (known to us).**
- ☑ **Contentions from non-globus programs (unknown to us), e.g., sharing storage, network.**

What affect transfer performance?

Contention from other non-globus program:

Transfer over ESnet testbed

(less likely to have non-globus load on endpoints)

Transfer over production DTN

(more likely to have non-globus load on endpoints)

What affect transfer performance?

Contention from simultaneous globus transfers (I/O, NIC, CPU & RAM):

What affect transfer performance?

Contention from simultaneous globus transfers (I/O, NIC, CPU & RAM):

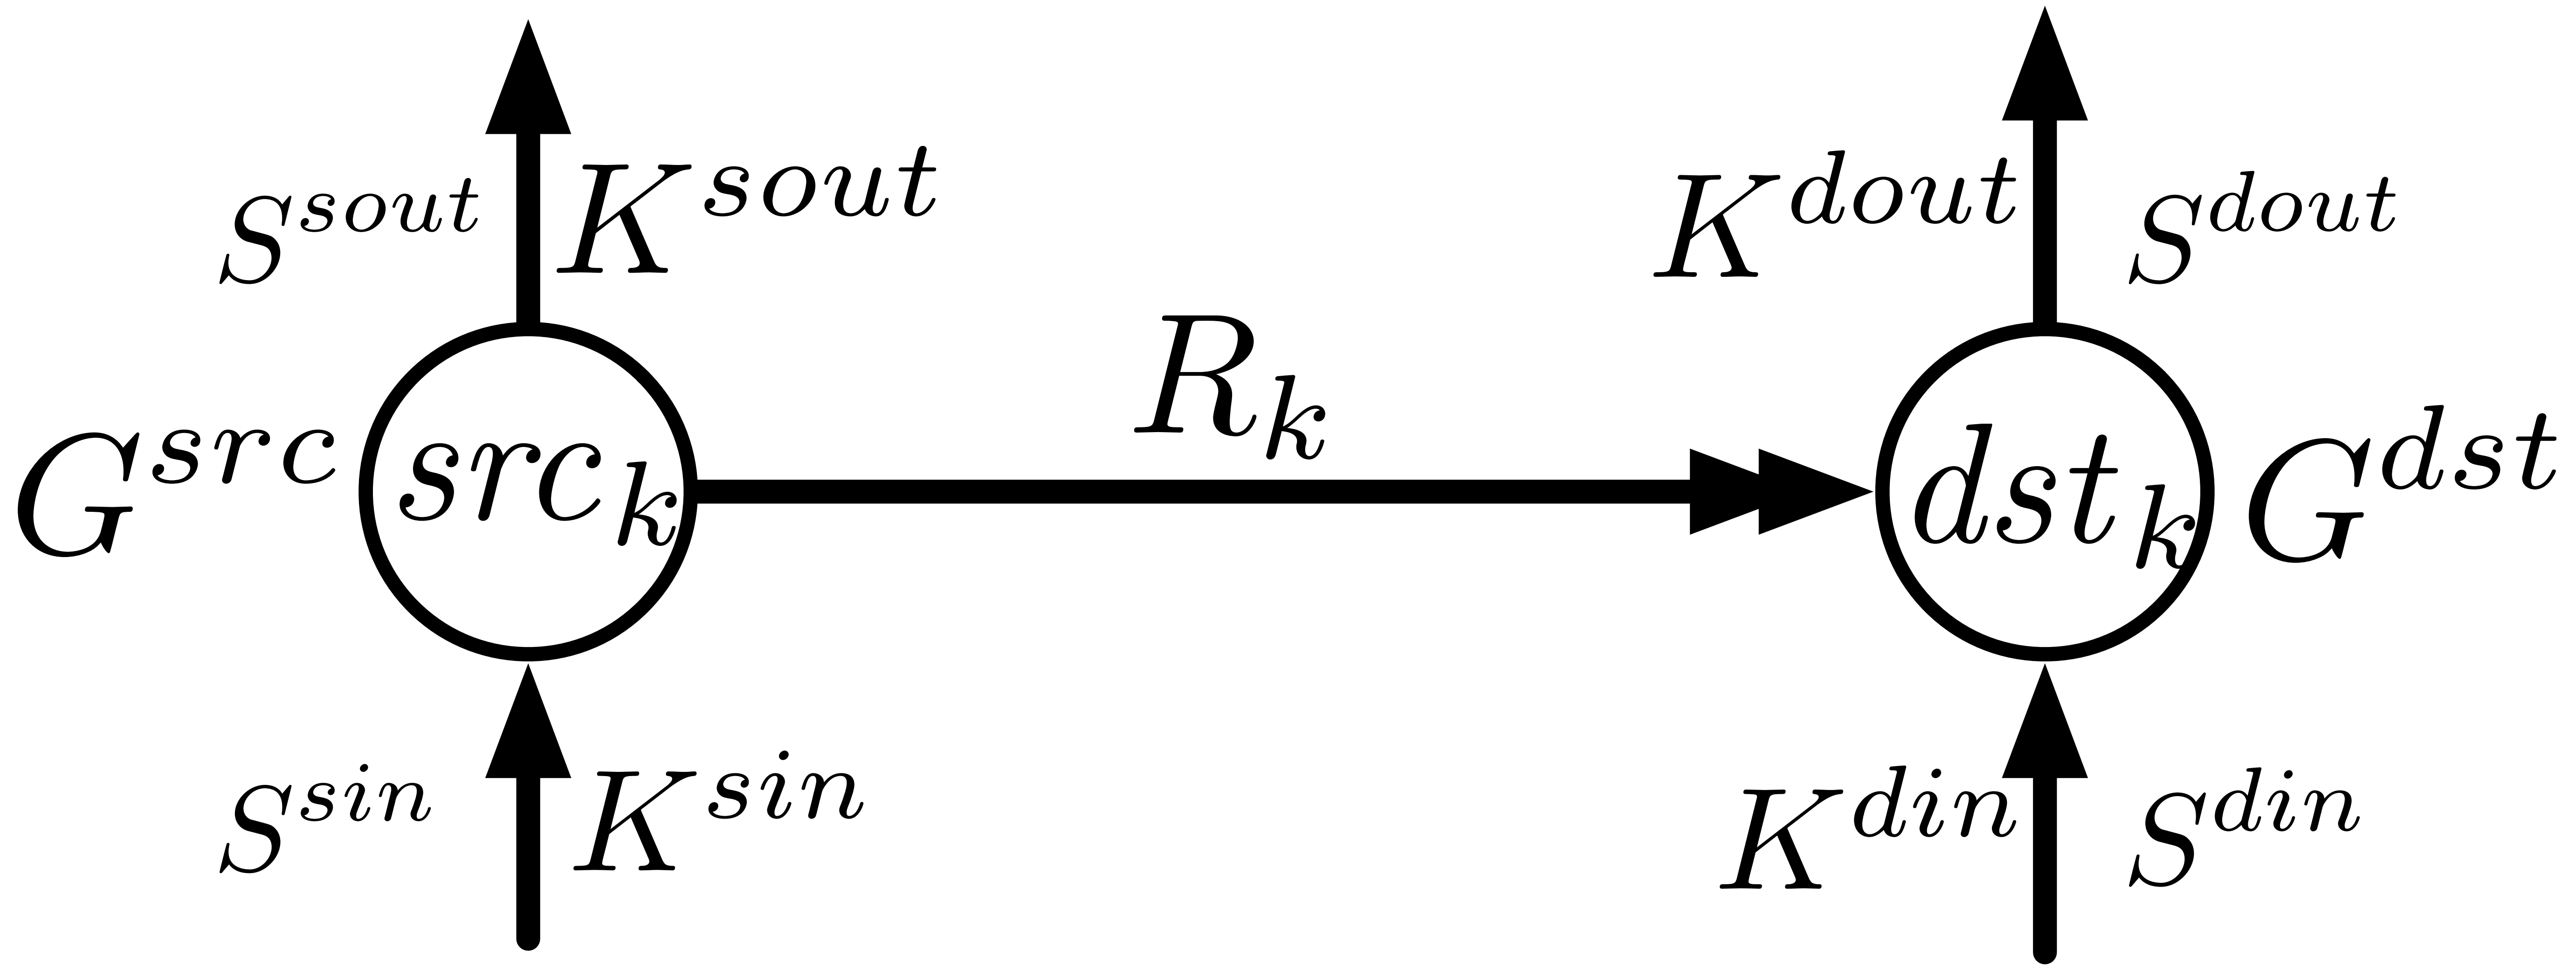

What affect transfer performance?

Contention from simultaneous globus transfers (I/O, NIC, CPU & RAM):

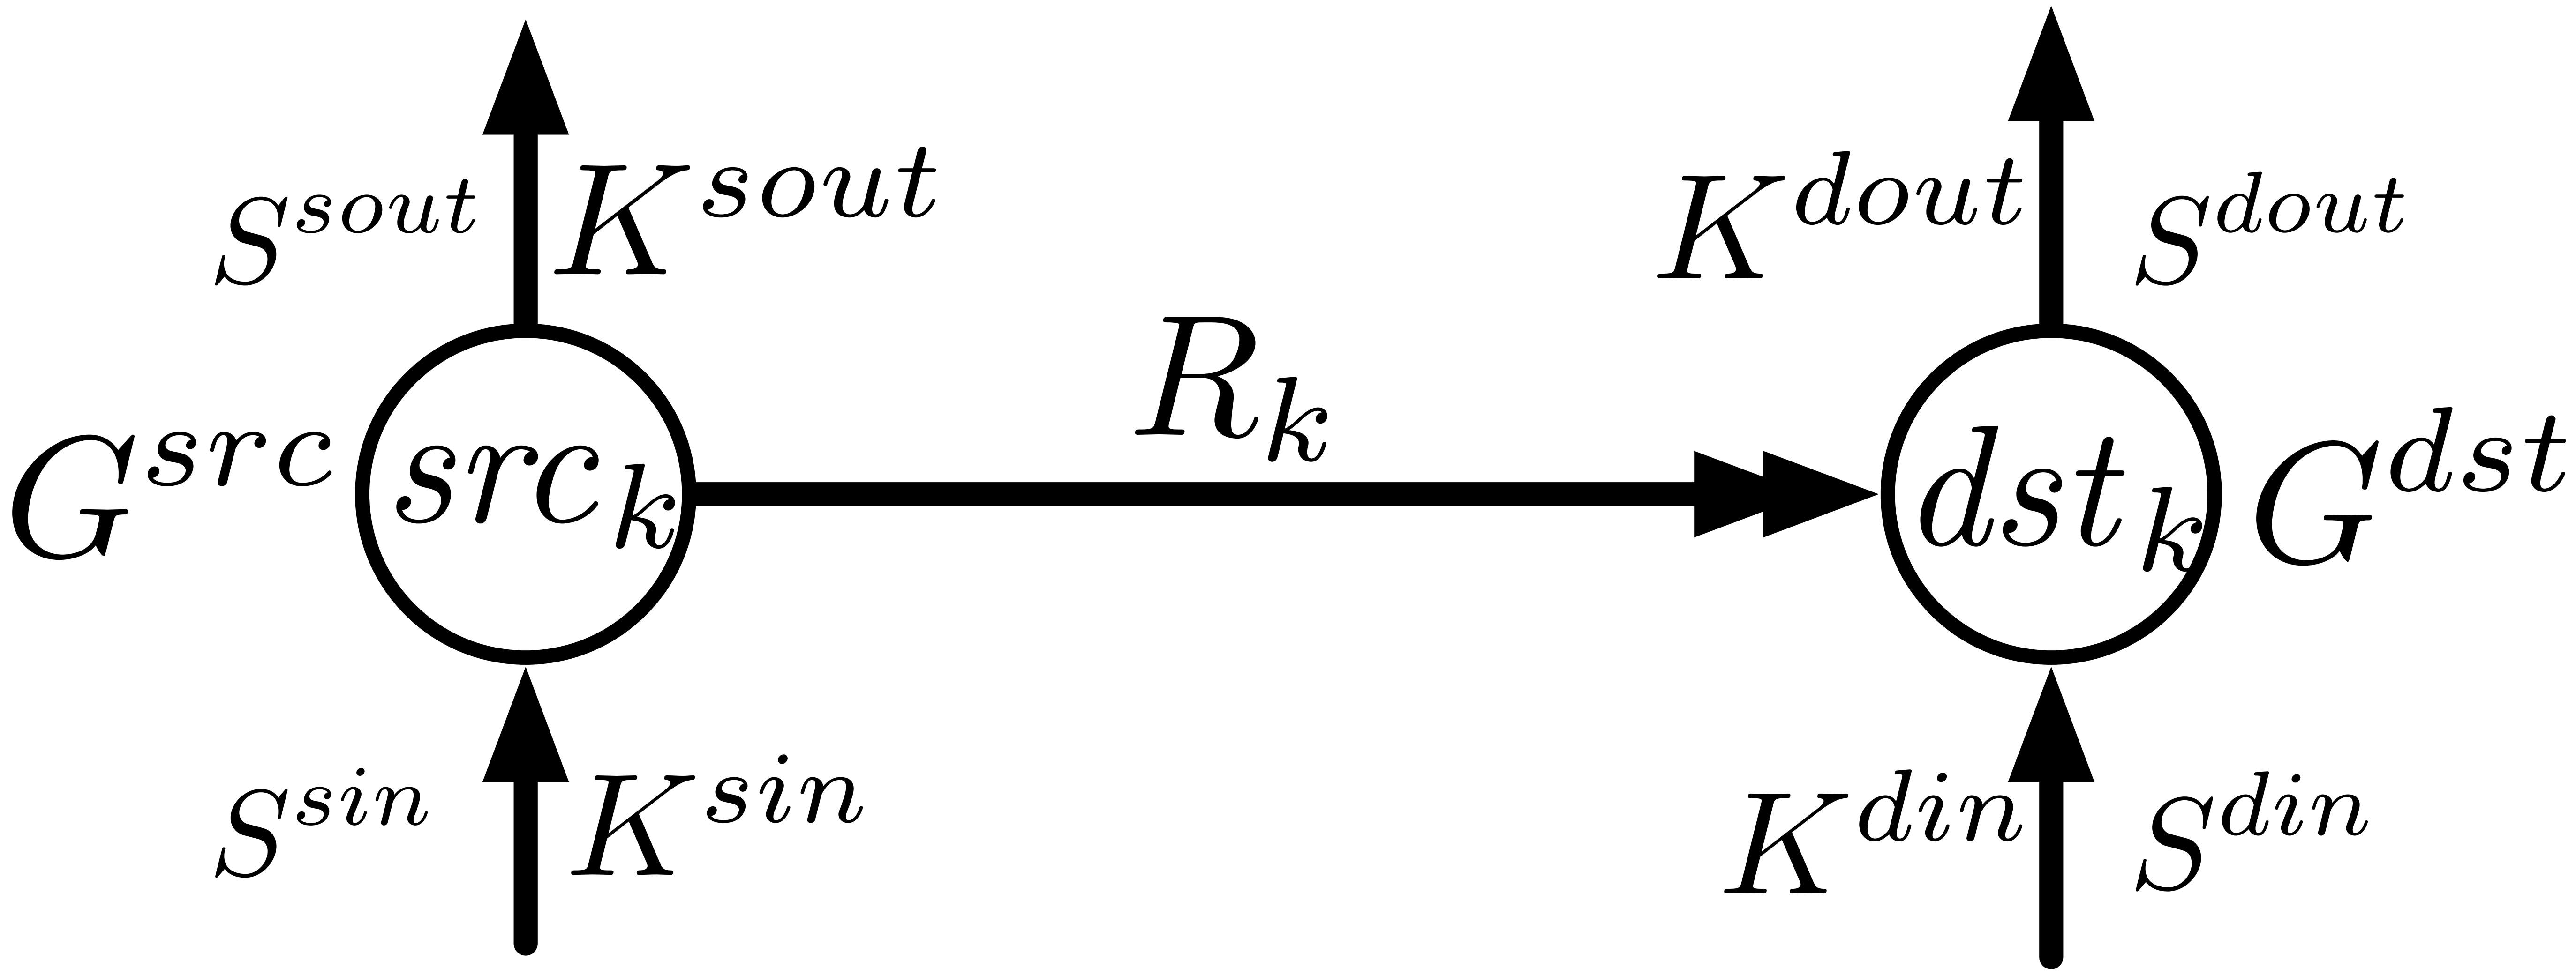

The Globus **contending transfer rate** for a transfer k at its source (src_k) and destination (dst_k) endpoints (demonstrated in Figure 7) is as follows:

$$K^{x \in \{sout, sin, dout, din\}}(k) = \sum_{i \in A_x} \frac{O(i,k)}{Te_k - Ts_k} R_i, \qquad (2)$$

where A_x is the set of transfers (excluding k) with src_k as source, when $x = sout$; src_k as destination, when $x = sin$; dst_k as source, when $x = dout$; and dst_k as destination when $x = din$; and $O(i,k)$ is the overlap time for the two transfers:

$$O(i,k) = \max\left(0, \quad \min(Te_i, Te_k) - \max(Ts_i, Ts_k)\right).$$

What affect transfer performance?

Contention from simultaneous globus transfers (I/O, NIC, CPU & RAM):

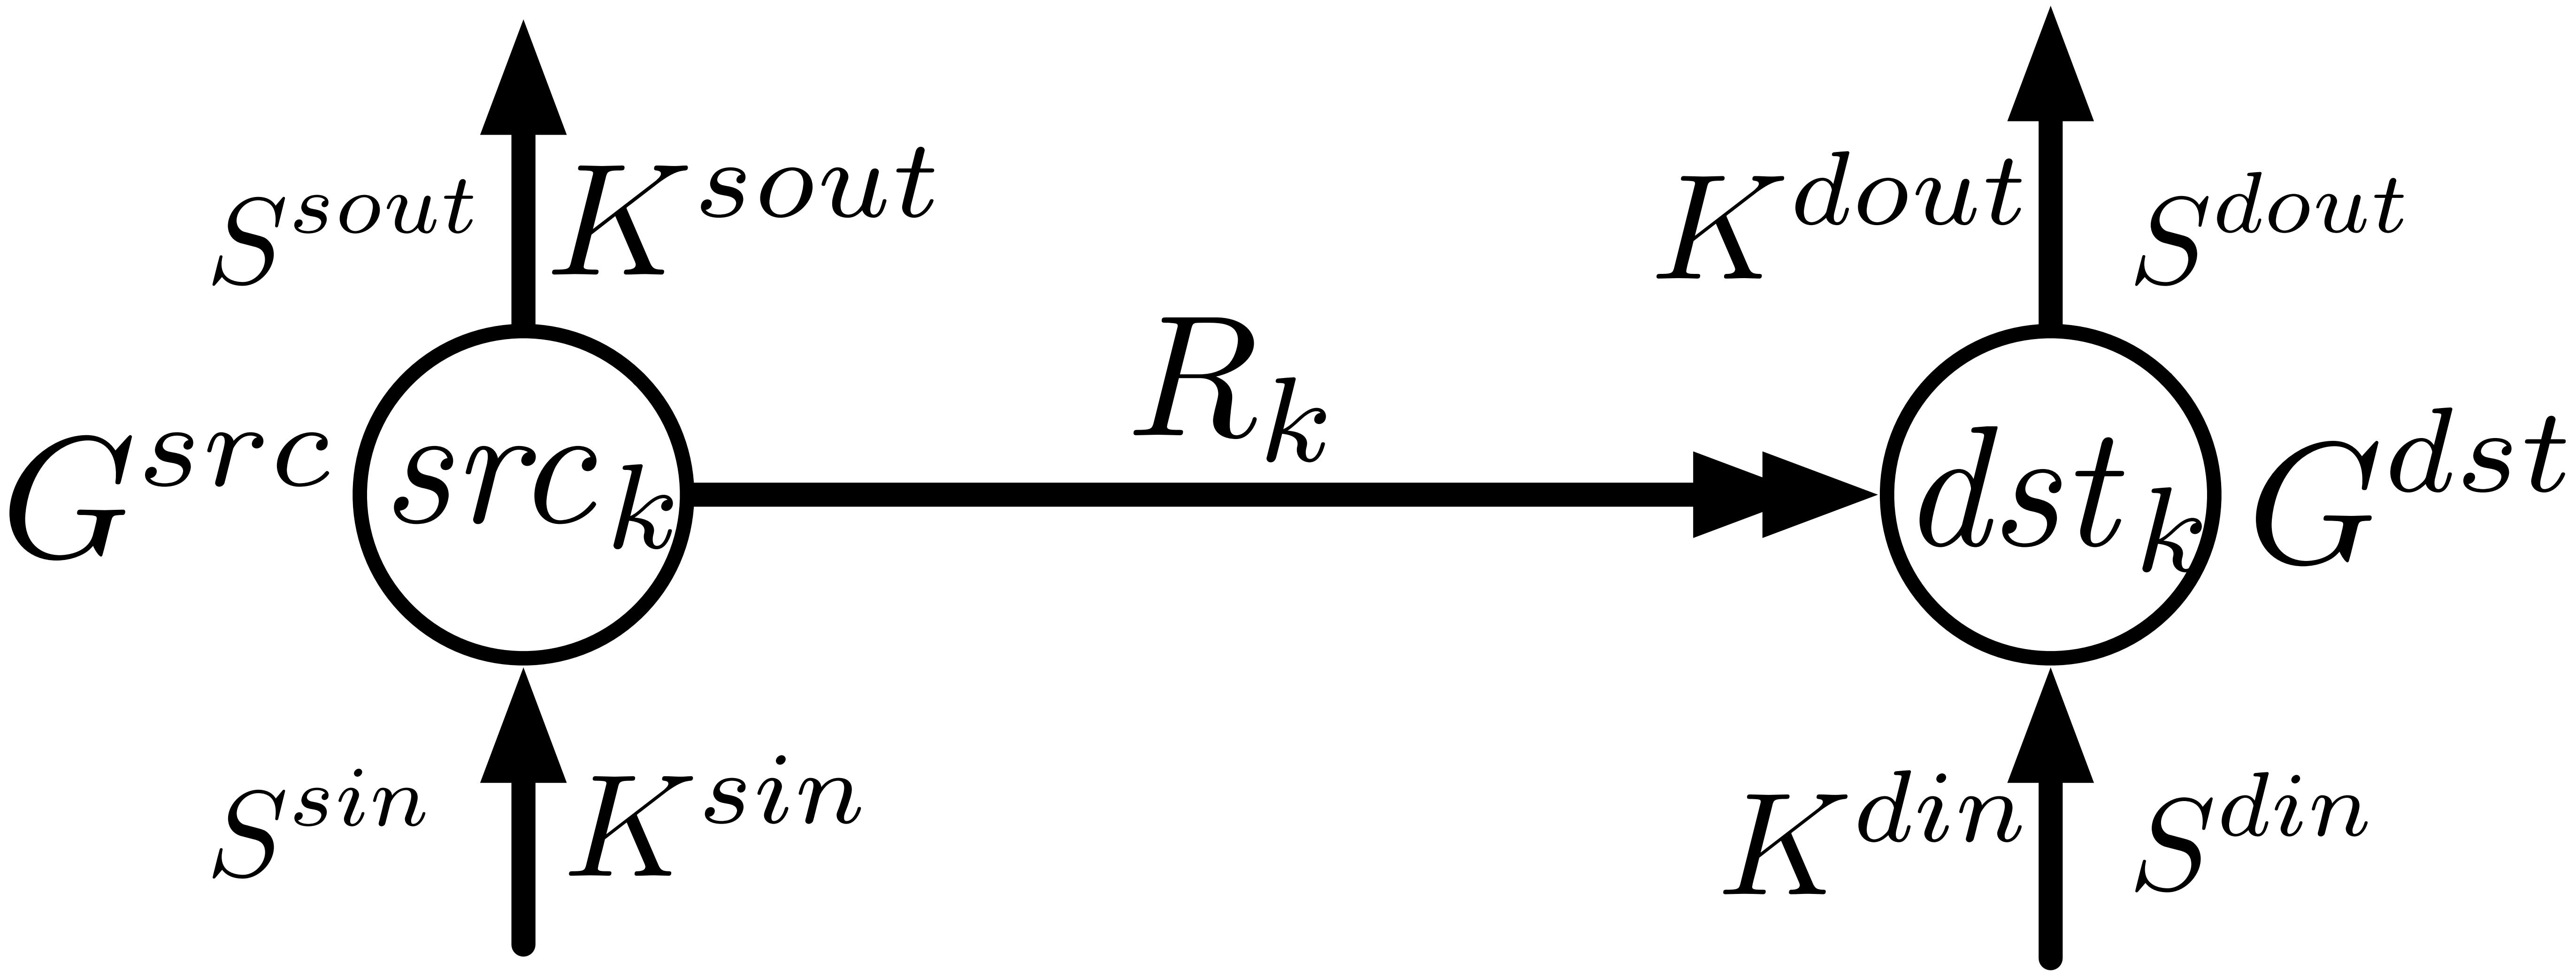

The Globus **contending transfer rate** for a transfer k at its source (src_k) and destination (dst_k) endpoints (demonstrated in Figure 7) is as follows:

$$K^{x\in\{sout, sin, dout, din\}}(k) = \sum_{i\in A_x} \frac{O(i,k)}{Te_k - Ts_k} R_i, \qquad (2)$$

where A_x is the set of transfers (excluding k) with src_k as source, when $x = sout$; src_k as destination, when $x = sin$; dst_k as source, when $x = dout$; and dst_k as destination when $x = din$; and $O(i,k)$ is the overlap time for the two transfers:

$$O(i,k) = \max\left(0, \quad \min(Te_i, Te_k) - \max(Ts_i, Ts_k)\right).$$

K^{sin}	Contending incoming transfer rate on src_k.
K^{sout}	Contending outgoing transfer rate on src_k.
K^{din}	Contending incoming transfer rate on dst_k.
K^{dout}	Contending outgoing transfer rate on dst_k.
C	Concurrency: Number of GridFTP processes.
P	Parallelism: Number of TCP channels per process.
S^{sin}	Number of incoming TCP streams on src_k.
S^{sout}	Number of outgoing TCP streams on src_k.
S^{din}	Number of incoming TCP streams on dst_k.
S^{dout}	Number of outgoing TCP streams on dst_k.
G^{src}	GridFTP instance count on src_k.
G^{dst}	GridFTP instance count on dst_k.
Nf	Number of files transferred.
Nd	Number of directories transferred.
Nb	Total number of bytes transferred.

Feature importance based on linear model

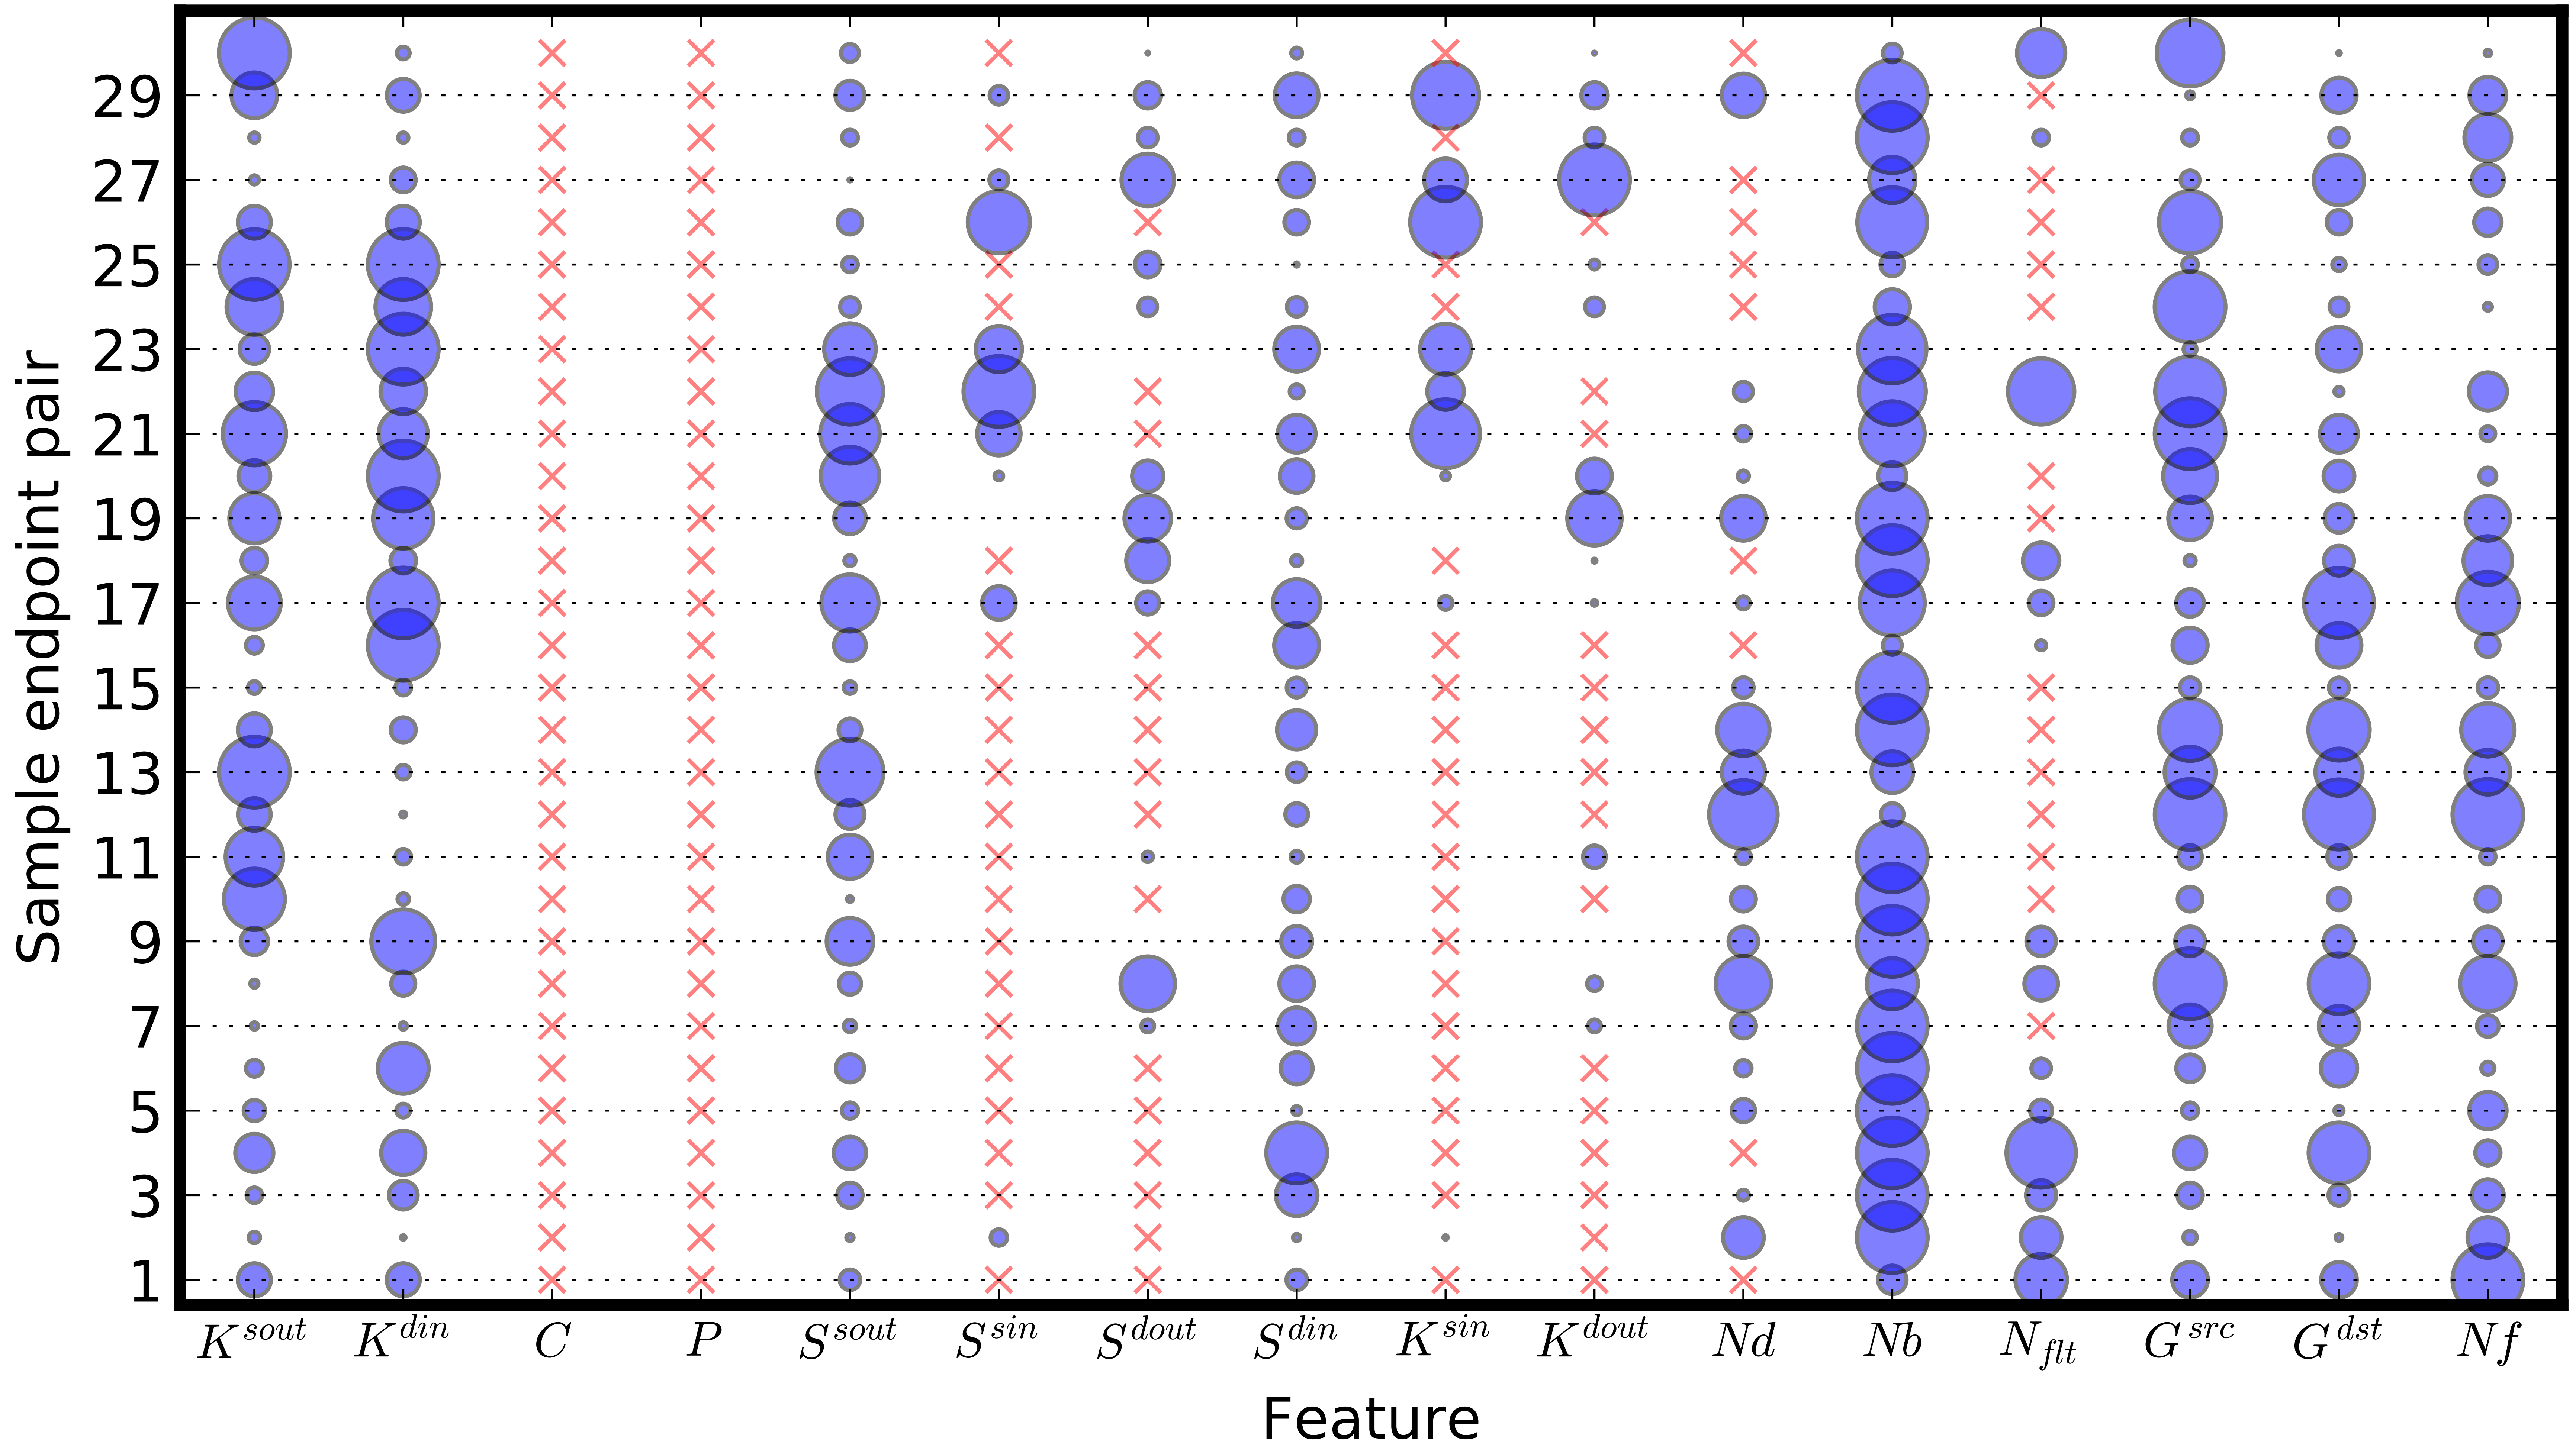

Circle size indicates the relative importance of features in the linear model, for each of 30 heavily used endpoint pairs. A red cross point means that feature was eliminated because of low variance.

Feature importance based on nonlinear model

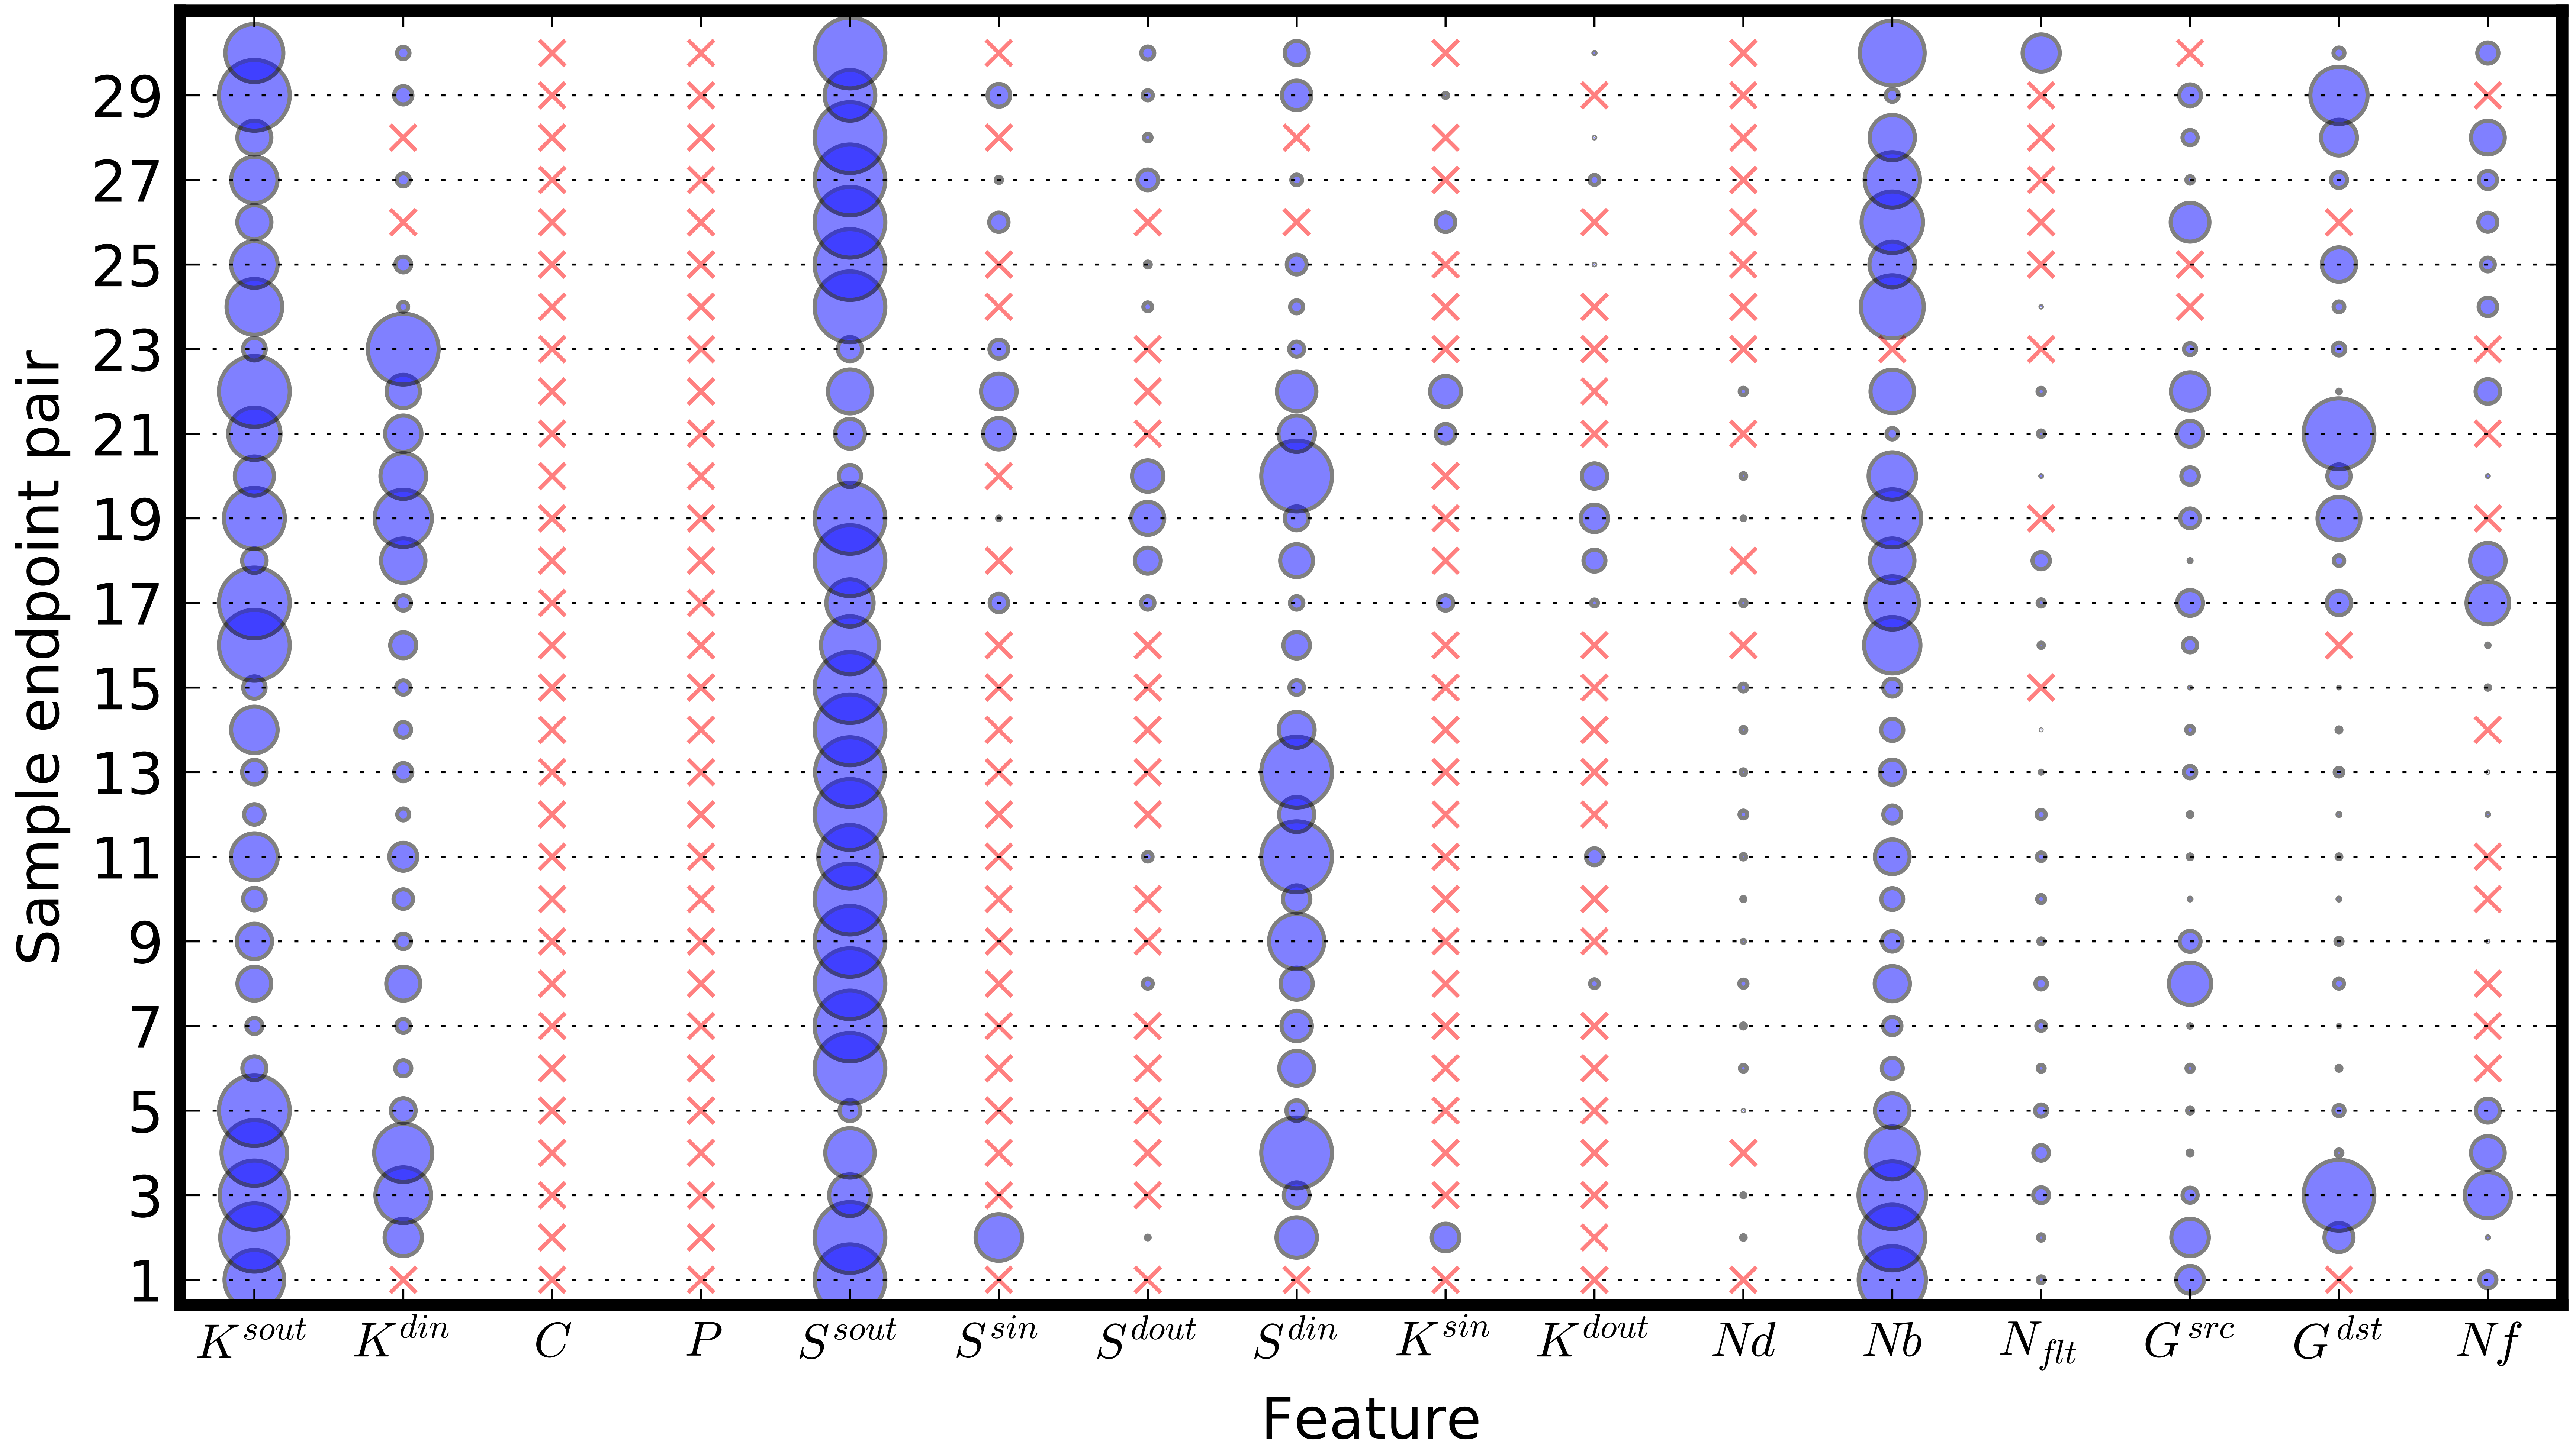

Circle size indicates the relative importance of features in the linear model, for each of 30 heavily used endpoint pairs. A red cross point means that feature was eliminated because of low variance.

Prediction

Linear versus nonlinear:

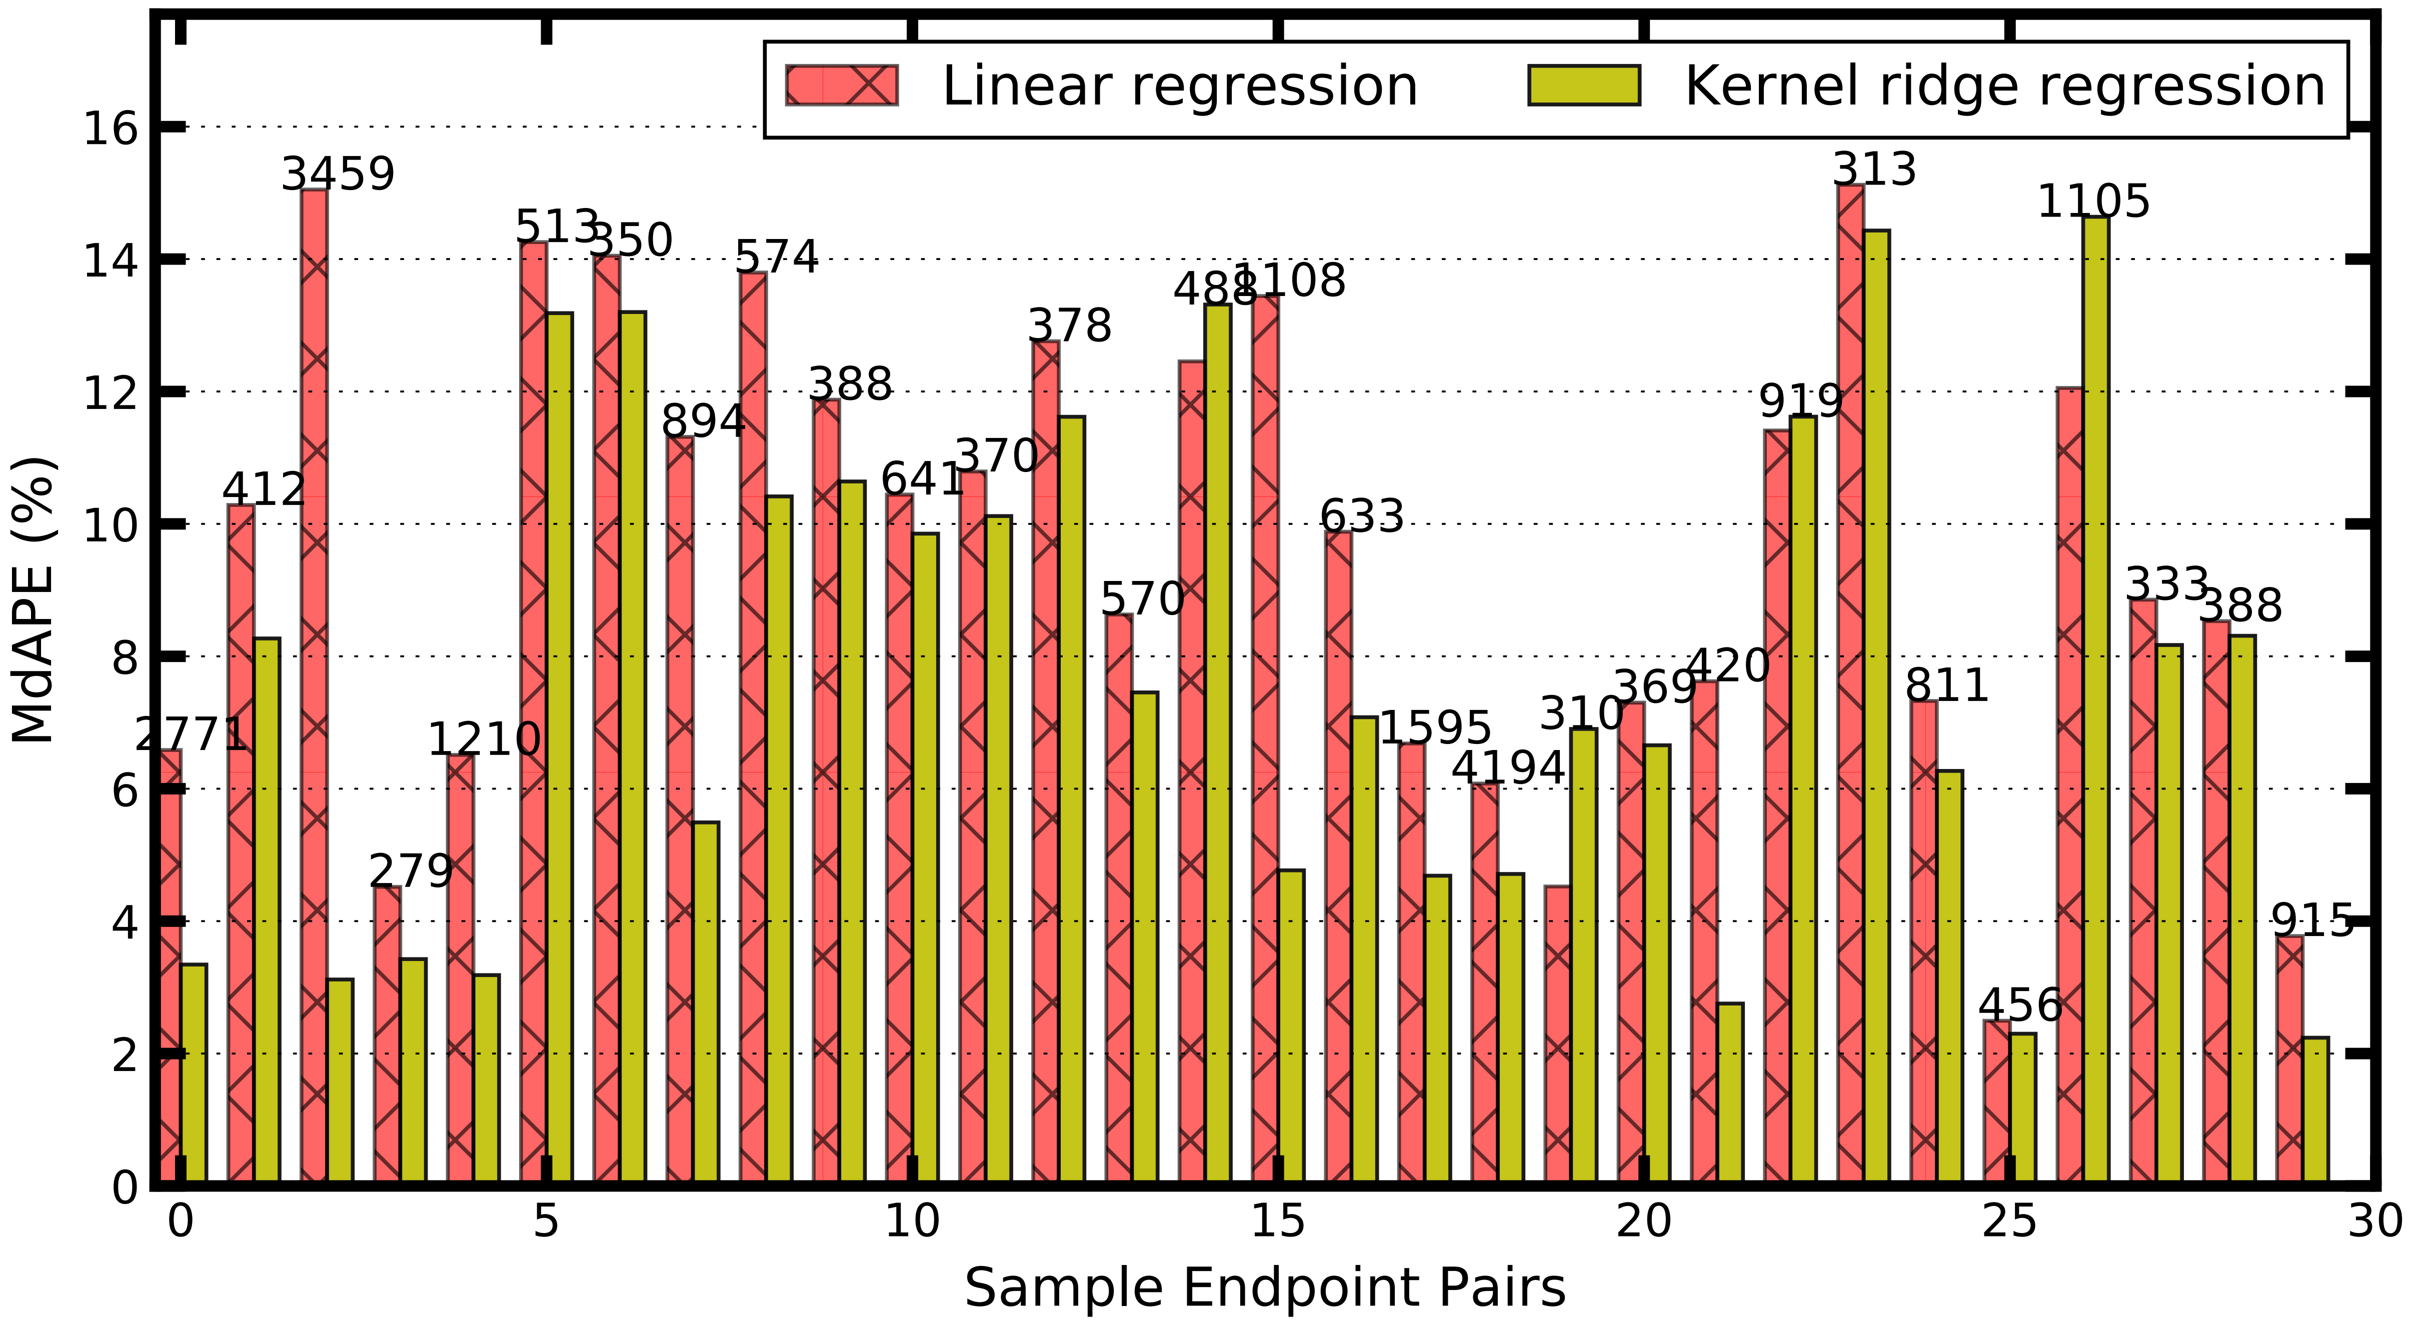

Linear regression versus kernel ridge regression[1]

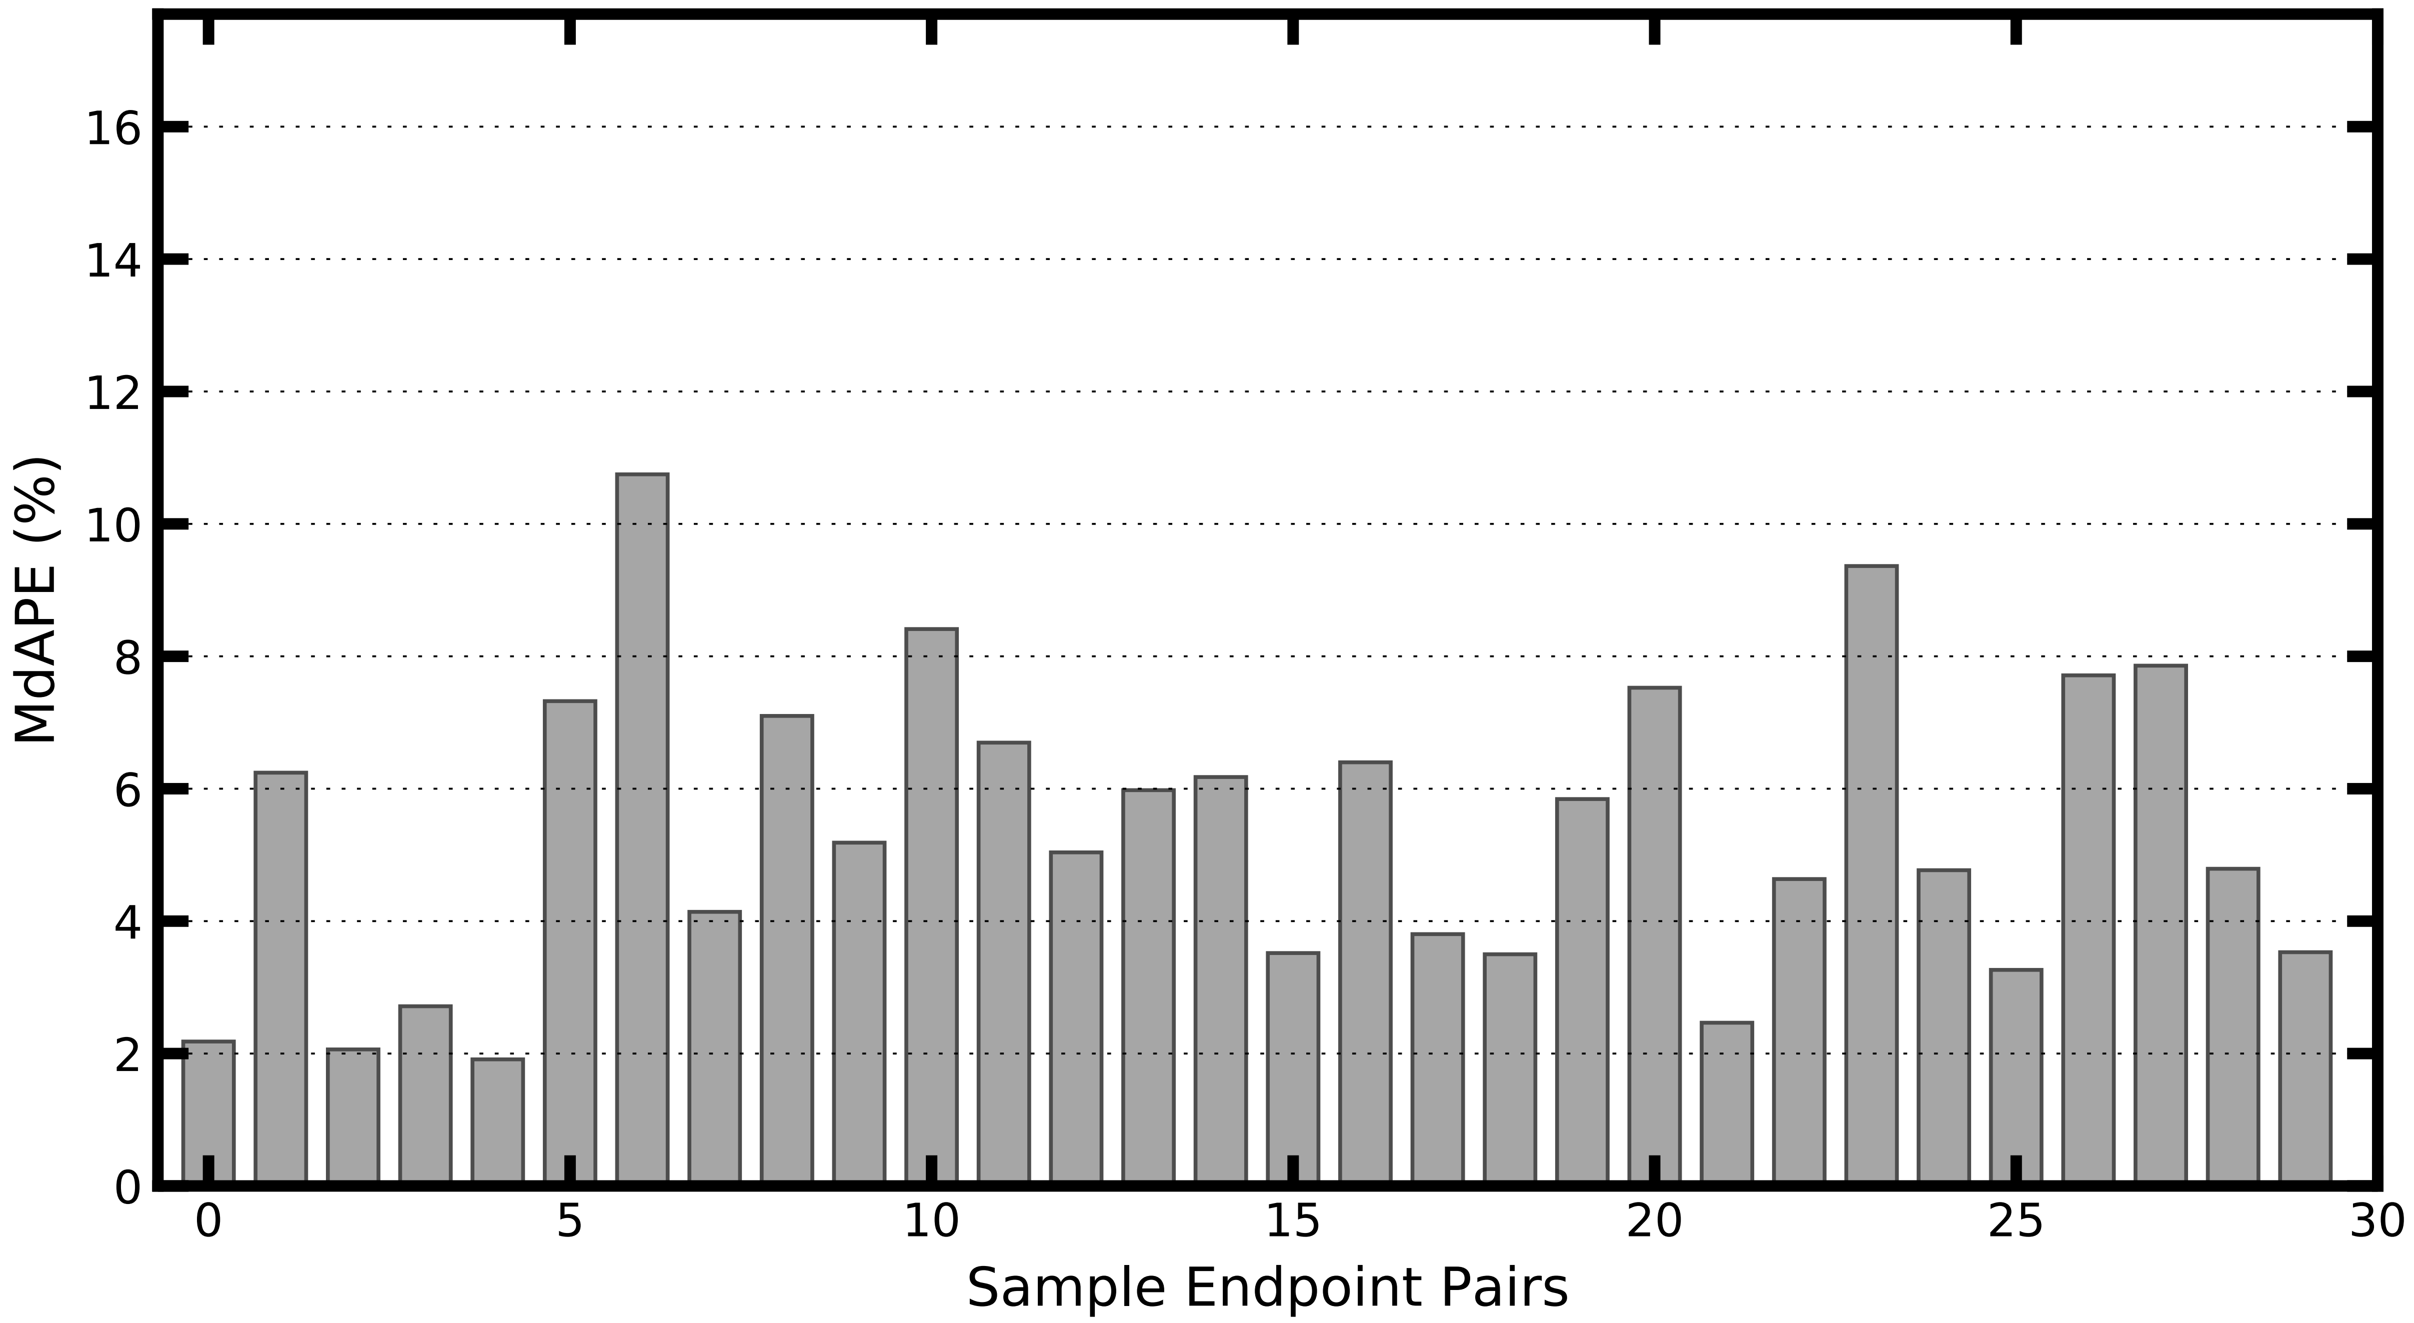

XGBoost[2]

[1] Murphy, K. P., Machine Learning: A Probabilistic Perspective,- chapter 14.4.3, pp. 492-493, The MIT Press, 2012
[2] https://xgboost.readthedocs.io/en/latest/

Argonne NATIONAL LABORATORY

Conclude and Future work

- Gain insights into the behavior of wide area data transfers.
- We derived features from Globus transfer log and studied their importance.
- We tried to make prediction based on the features we derived.
- Our models achieve good accuracy when there is less unknown load.
- This work has been accepted by HPDC'17, more details are available in the paper.

- Unknown load coming from non-globus load is troublesome;
- Can the cutting edge methods, like deep learning, help?