

# Investigating Hardware Offloading for Reed-Solomon

## Encoding

Gustavo De Leon  
gdeleon@berkeley.edu  
UC Berkeley

John Dermer  
John.Dermer@enmu.edu  
Eastern New Mexico University

Tyler Rau  
tylersrau@gmail.com  
Dakota State University

### Introduction

Reed-Solomon (RS) encoding is a storage scheme which offers better scalability, but requires heavier computation, compared to other models. This presents a problem as it requires users to purchase browner CPUs than would be otherwise necessitated. However, Mellanox's ConnectX-4&5 Infiniband cards have the capability to perform RS encoding on the HCA hardware; removing the need for powerful CPUs.

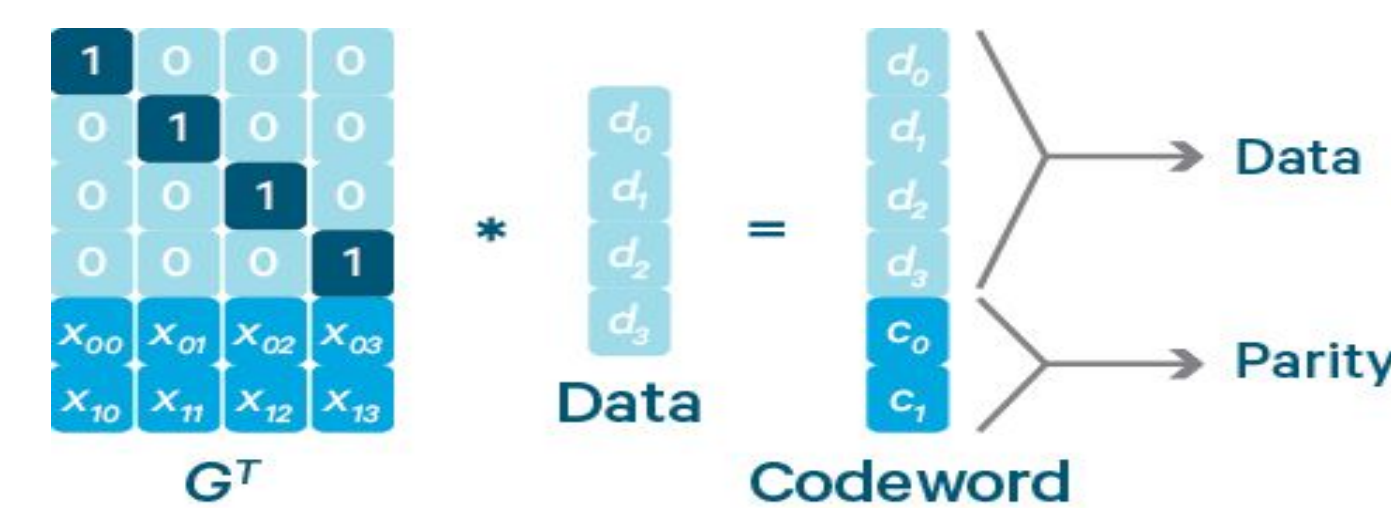
### Specifications

- 2 Intel(R) Xeon(R) CPU E5-2620 v4 @ 2.10GHz
- Mellanox Connectx-4/5 IB Cards
- 64 GB RAM
- CentOS 6.9
- GCC 6.2.1

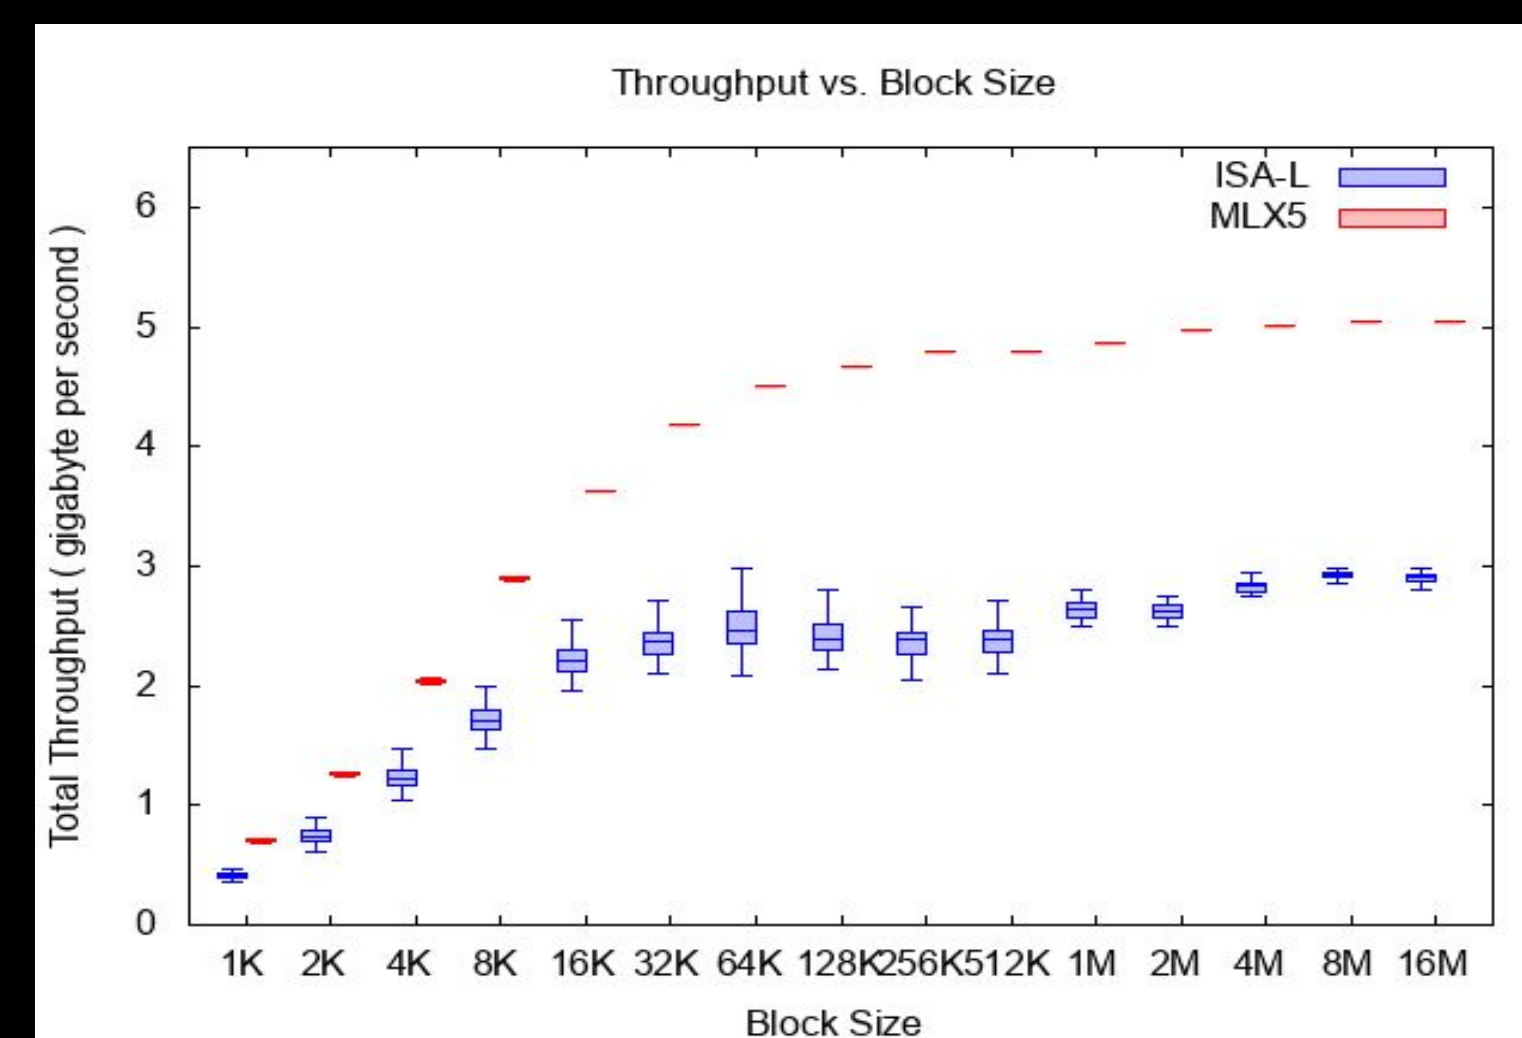
### Methodology

- The first two were run on both ISA-L and card offloading to identify their effect on throughput
  - Block size test sizes covered  $1024 \cdot 2^x$  bytes where X goes from 0 to 15 on a single thread
  - Concurrency testing was done with a constant block size of 1MB and involved between 2 and 20 threads
- We tested concurrency on a single machine running 6 ConnectX-4 cards going from 1 to 20 processes per card, and then ramping up from 1 to 6 cards, 1 to 4 processes on each

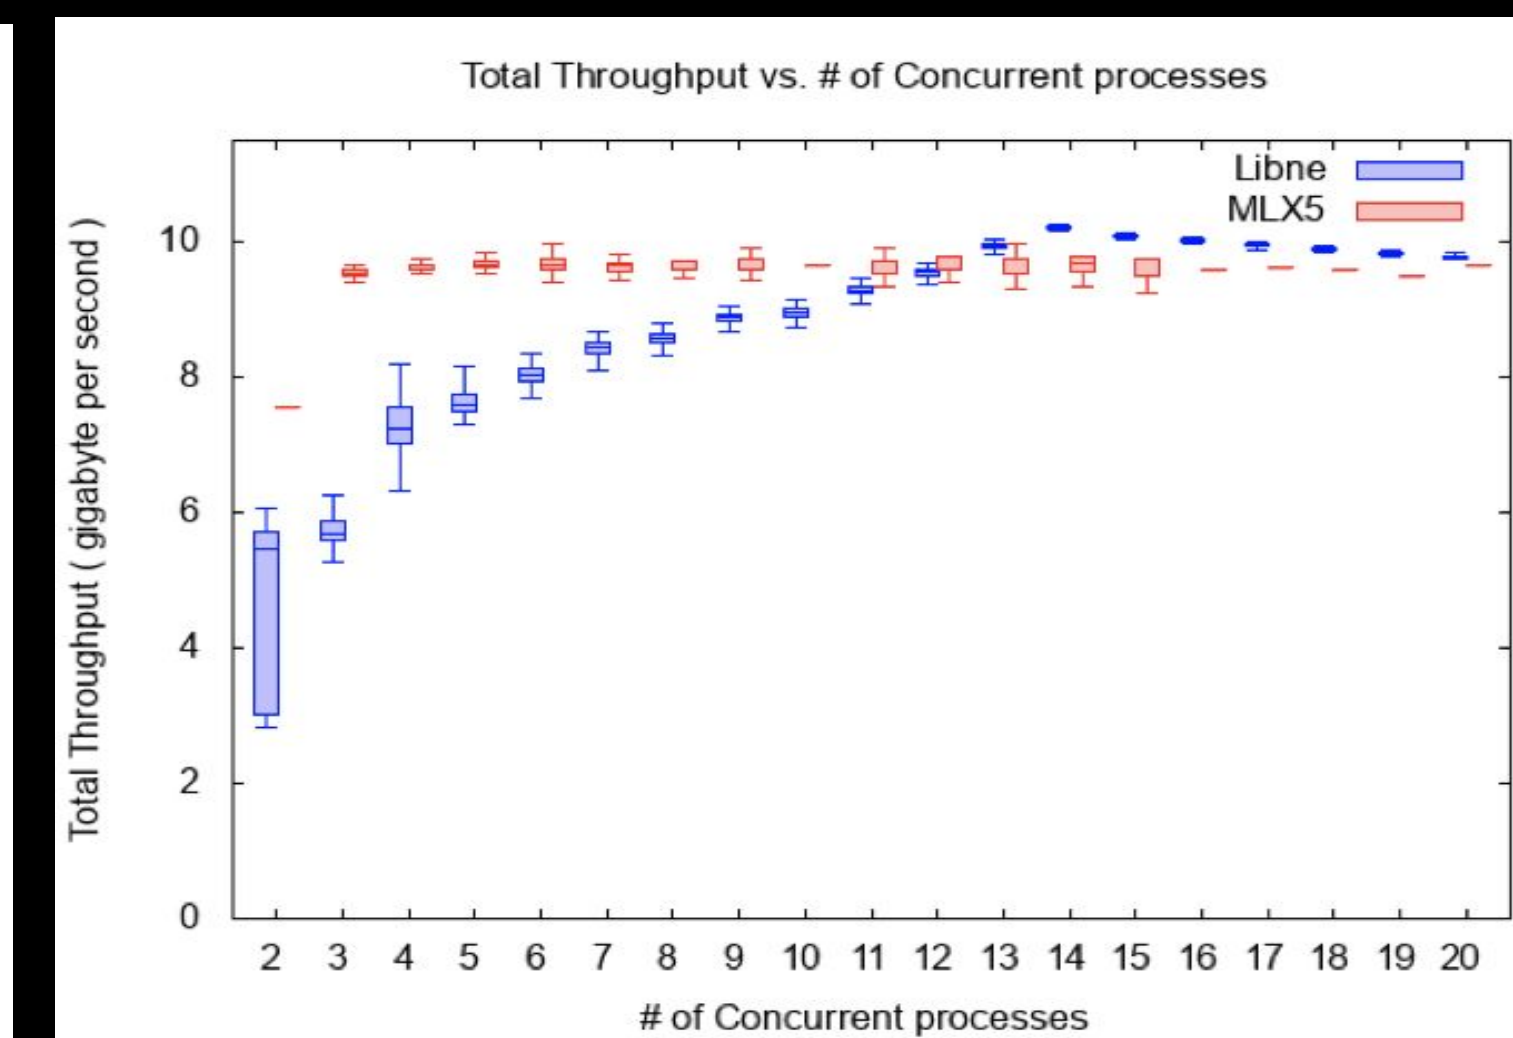
### Erasure Encoding



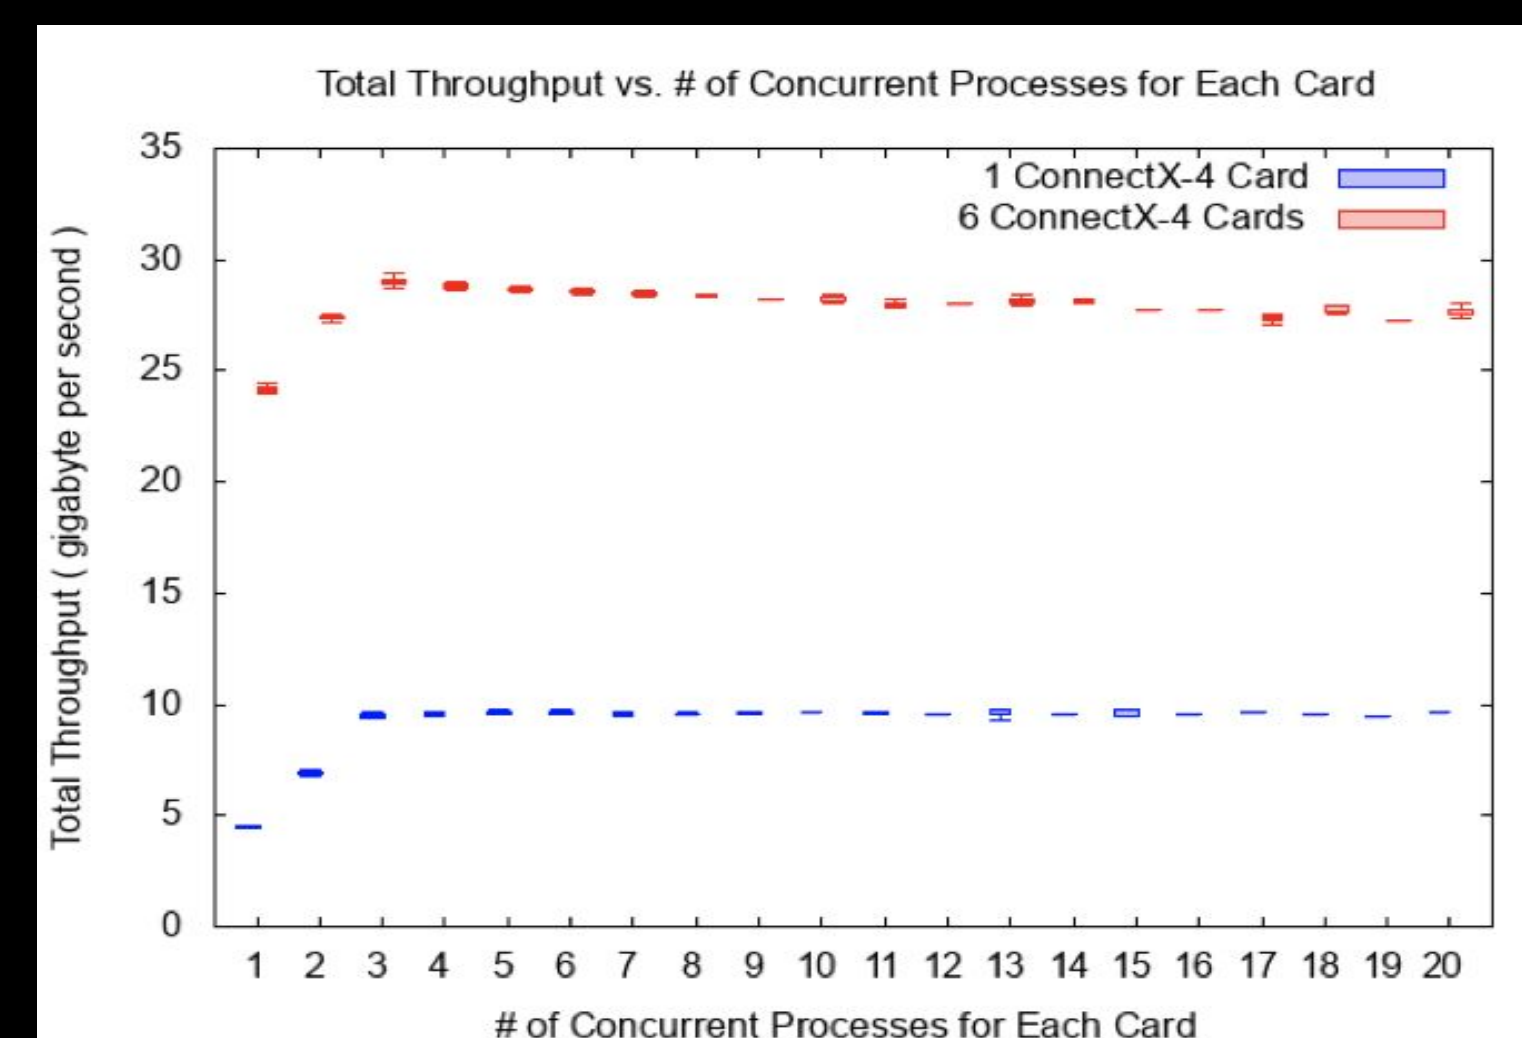
### Results



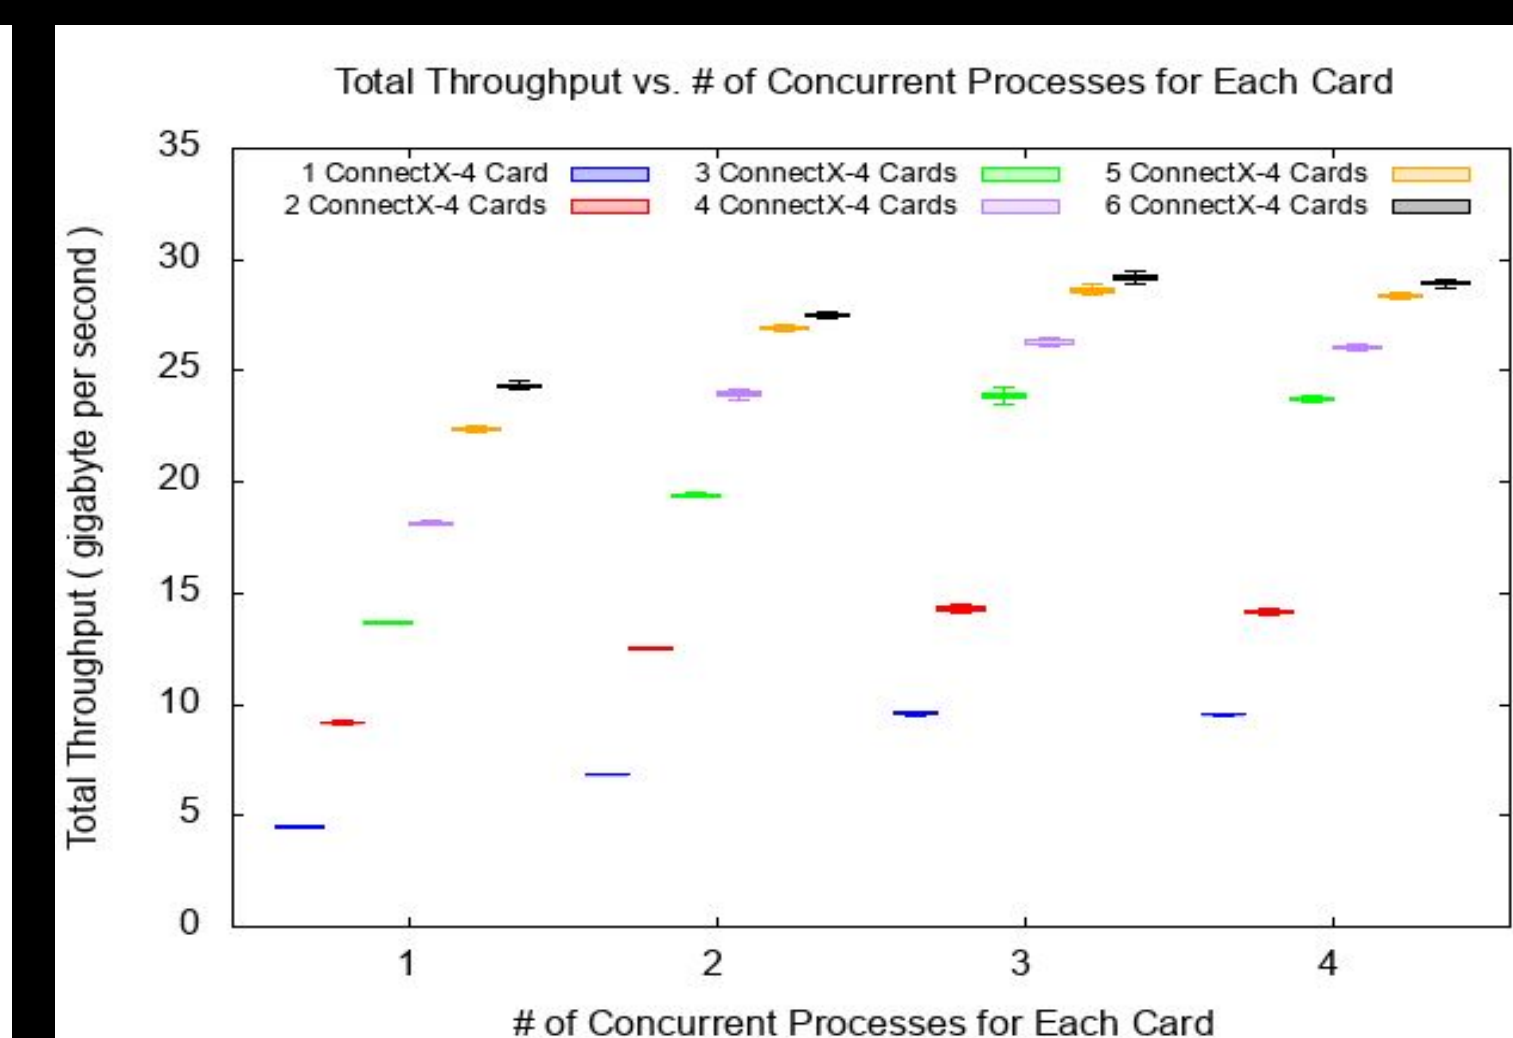
Above data shows a clear increase in throughput using offloading over ISA-L



Above data shows that initially offload offers increased throughput but as concurrency increases ISA-L exceeds the performance of the offload, however ISA-L begins to degrade soon after.



Above data shows a definitive increase in throughput when using multiple cards.



Above data shows adding additional cards always gives a performance increase; however the return on investment is diminished as the number of cards in use increases.

### Conclusion

1. ConnectX-4 cards are limited to a Galois Field of  $2^4$ , but ConnectX-5 cards can go up to  $2^8$  which enables compatibility with ISA-L encoded data
2. Mellanox IB Cards are able to encode at higher and more consistent throughput than the software based ISA-L
3. While not a linear scaling, running multiple cards on a single machine does offer a significant increase in encoding throughput for up to 3 cards per CPU

### Future Plans

- Further explore the cards' encode\_send functionality which enables them to encode data and forward it on using IB Verbs in a single blocking call
- Investigate ways to enable the cards to decode and/or rebuild data distributed across multiple servers
- Use ConnectX-5 cards to investigate effect of Galois Field used to encode data on throughput of encoding
- Look to optimize number of cards per machine for maximum throughput

### Acknowledgements

**Mentors:** David Bonnie, Jeff Inman, Will Vining  
**Instructor:** Joan Lucas  
**Assistant:** Hunter Easterday