

In conclusion

- **libCEED** can accelerate matrix-free operations for high-order discretizations
 - * Many backends targeting different hardware
- Recently added (compile-time) **single precision** support
- In progress: **mixed-precision computations** for libCEED operators
 - * Minimize memory movement for memory bound kernels
 - * Allow decoupling of memory/computation precision where advantageous
 - ❖ Interface design is a WIP, ideas and feedback welcome