Towards Access Pattern Aware Checkpointing For Kokkos Applications

Nigel Tan[1] (student)
Bogdan Nicolae[2] (advisor), Nicolas Morales[3] (advisor), Keita Teranishi[3] (advisor), Sanjukta Bhowmick[4] (advisor), Michela Taufer[1] (advisor)
[1]University of Tennessee Knoxville, [2]Argonne National Laboratory, [3]Sandia National Laboratories, [4]University of North Texas

Contributions:

- Highlight the limitations of the common checkpoint philosophy on increasingly heterogeneous systems
- Demonstrate a 4-8x reduction in checkpoint size by combining application and system awareness for checkpointing using sparse update patterns and incremental checkpoints as an example
- Lay the groundwork for supporting access pattern aware checkpointing in Kokkos Resilience and VeloC

Naive Application Level Checkpointing

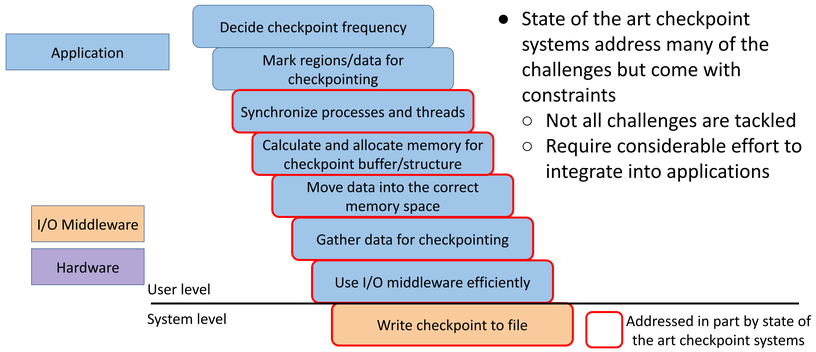

- State of the art checkpoint systems address many of the challenges but come with constraints
  - Not all challenges are tackled
  - Require considerable effort to integrate into applications

Checkpointing at the application level is complex and not portable

Limits of Current Checkpoint Philosophy

- Growing I/O and network complexity in HPC systems
  - Rising need for portability on heterogeneous platforms
- Integrating existing checkpoint systems is challenging
  - Increase in build complexity
  - Difficult to guarantee long term support
- Optimization is typically application specific
- Developing efficient checkpoint runtimes is difficult

Common resilience philosophy: checkpoint all, checkpoint frequently, do not pay attention to the update patterns

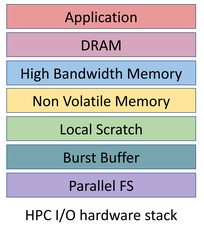

HPC I/O hardware stack

Checkpointing With Kokkos Resilience and VeloC

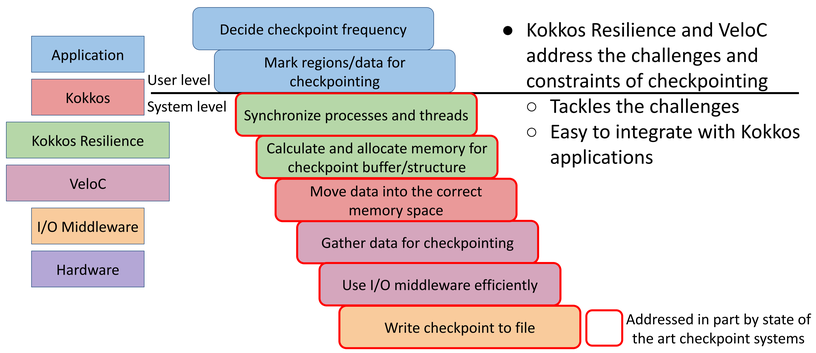

- Kokkos Resilience and VeloC address the challenges and constraints of checkpointing
  - Tackles the challenges
  - Easy to integrate with Kokkos applications

Combine application and system awareness for efficient and portable checkpoints

Applications with Sparse Data Update Patterns: the Fido Use Case

What if an application has a sparse data update pattern (i.e., few data changes at each iteration)?

- Use case: **Fido**, a graph alignment application
  - Compute a graphlet degree vector (GDV) for each vertex
  - Collect the GDVs into a matrix
  - Compare two graphs using the GDVs
- Data update pattern is highly input dependent
- **Need to checkpoint the matrix of GDVs**

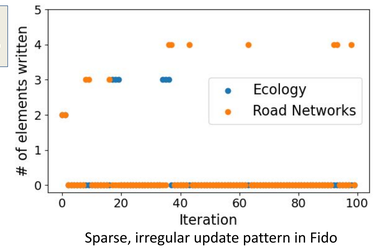

Sparse, irregular update pattern in Fido

Sparse update patterns favor incremental checkpoints

Current Solution: Full Checkpoint Strategy

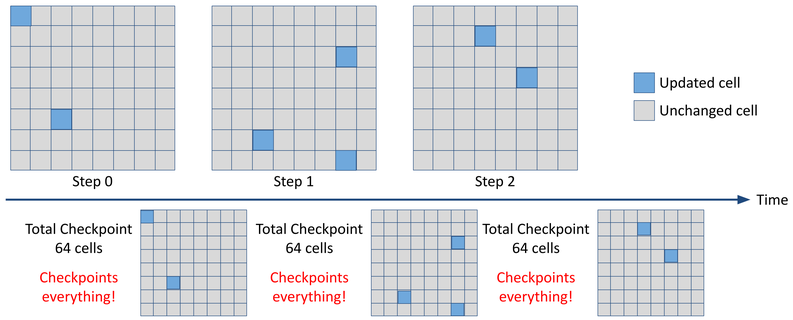

Full checkpoints are simple but can be inefficient and ignore data update patterns

A More Efficient Solution: Incremental Checkpoint Strategy

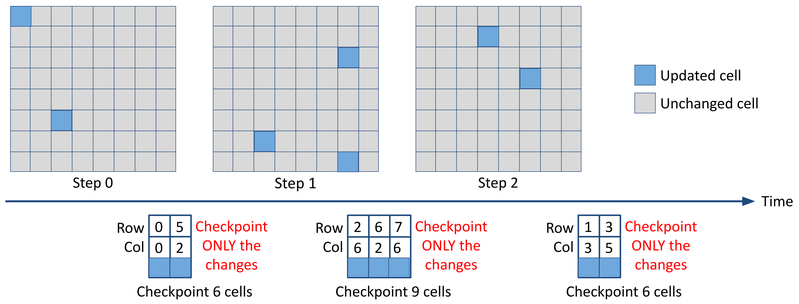

Enhance Kokkos Resilience and VeloC with access pattern aware checkpointing

Initial Implementation Of Access Pattern Aware Checkpointing

- Track the changes to the matrix of GDVs using a Map
- Serialize data for better write performance
- Run the tests on Tellico (Power9)
- Compare two different sets of graphs

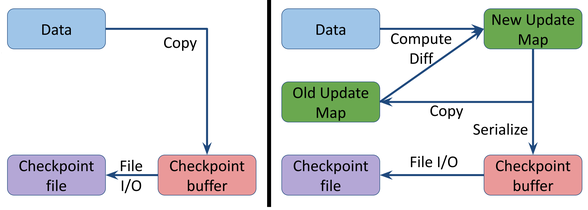

Results

Compared sets	Compared graphs	Scientific Domain	# Vertices	# Edges	# Iterations (total per graph)	Frequency (# iterations)
Set 1	<ecology1, ecology2>	Ecology data	1,000,000 999,999	4,996,000 4,995,991	6,154M 6,154M	100M
Set 2	<asia_osm, germany_osm>	Road networks	11,950,757 11,548,845	25,423,206 24,738,362	4,724M 3,126M	100M

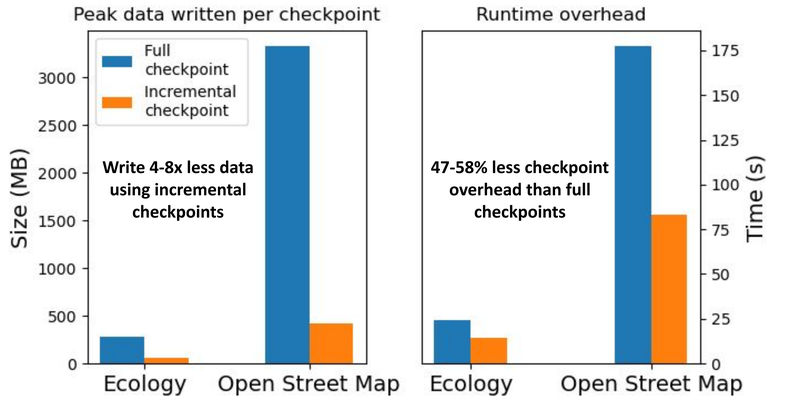

Conclusions

- Address the limits of current checkpoint philosophy with Kokkos Resilience and VeloC
- Enable a 4-8x reduction in data written for checkpointing sparsely updated data
- Build the foundation for access pattern aware checkpointing in Kokkos applications

Future Work

- Implement incremental checkpoint capabilities in Kokkos Resilience and VeloC
- Extend access pattern aware checkpointing to other patterns beyond sparse updates

References

1. Nicolae, B., Moody, A., Gonsiorowski, E., Mohror, K., & Cappello, F. (2019, May). Veloc: Towards high performance adaptive asynchronous checkpointing at large scale. In *2019 IEEE International Parallel and Distributed Processing Symposium (IPDPS)* (pp. 911-920). IEEE.
2. Morales, N., Teranishi, K., Nicolae, B., Trott, C., and Cappello, F., "Towards High Performance Resilience using Performance Portable Abstractions," to appear in *Euro-Par Conference 2021*.

Acknowledgements

Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525.

GCLab information