Progress Porting LLNL Monte Carlo Transport Codes to Nvidia GPUs

M&C 2023

August 16, 2023

M. Pozulp, R. Bleile, P. Brantley, S. Dawson,
M. McKinley, M. O'Brien, A. Robinson, M. Yang

LLNL-PRES-852689

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under contract DE-AC52-07NA27344. Lawrence Livermore National Security, LLC

Lawrence Livermore National Laboratory

This talk describes our progress porting LLNL MC codes to GPUs

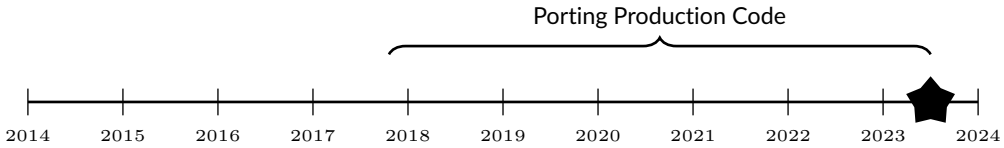

★ M&C 2023 “Progress...” by Pozulp et al.

We have significantly improved on the results in our M&C ’19 paper

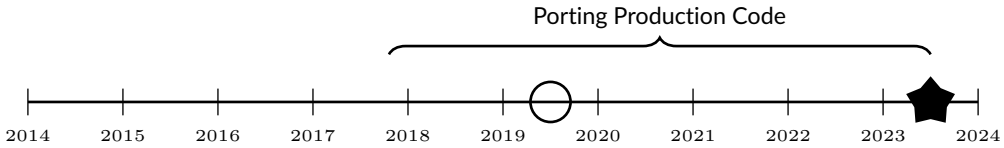

◯ M&C 2019 “Status…” by McKinley et al.

★ M&C 2023 “Progress…” by Pozulp et al.

Our M&C ’17 paper summarized an era of experimentation

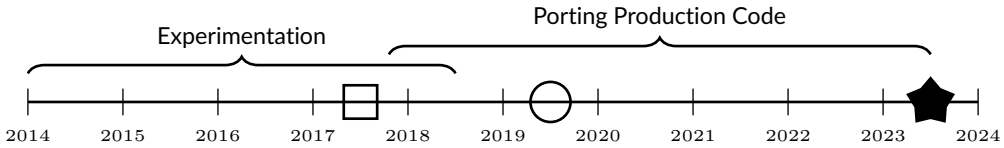

□ M&C 2017 “...Research Efforts...” by Brantley et al.

◯ M&C 2019 “Status...” by McKinley et al.

★ M&C 2023 “Progress...” by Pozulp et al.

Shared infrastructure allowed us to port two codes simultaneously

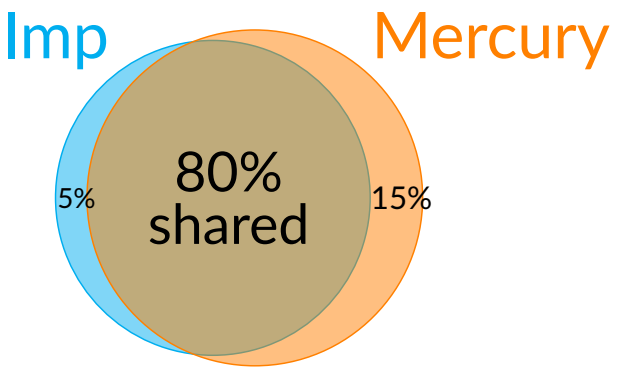

Imp ∪ Mercury = 370,000 lines C++

We compare two platforms: CTS-1 & Sierra

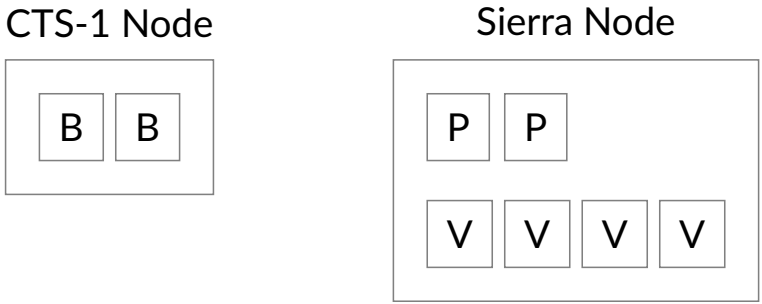

We devised a new abstraction for V100 (1/3)

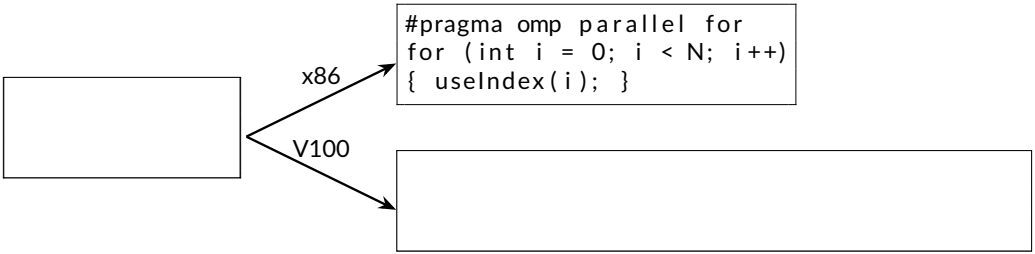

We devised a new abstraction for V100 (2/3)

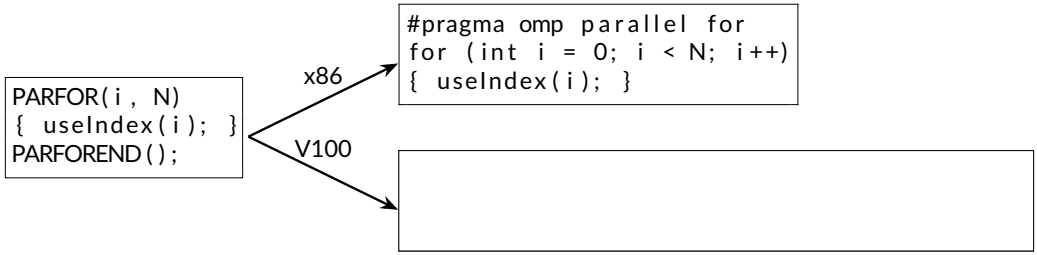

We devised a new abstraction for V100 (3/3)

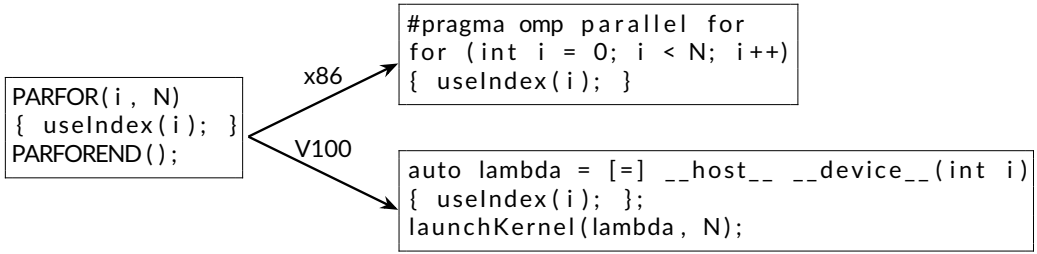

We implemented event-based tracking (1/4)

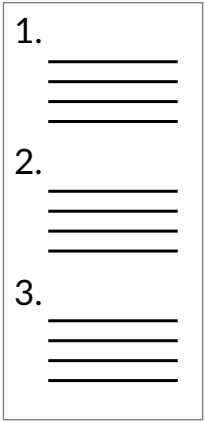

We implemented event-based tracking (2/4)

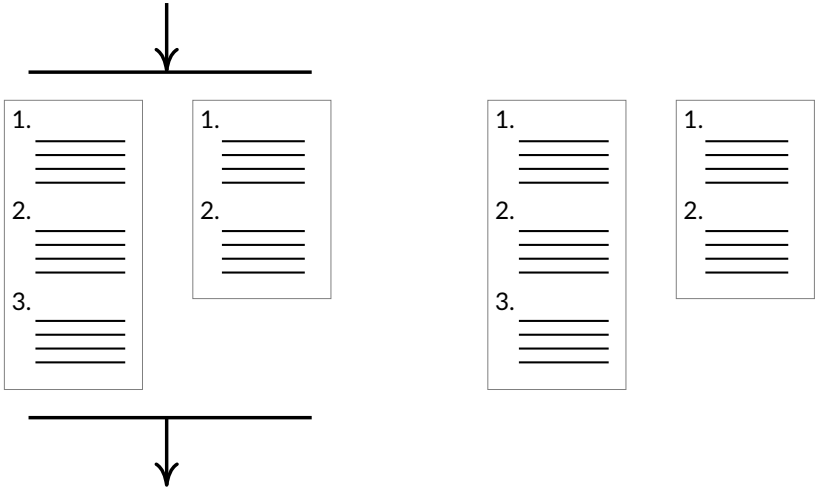

We implemented event-based tracking (3/4)

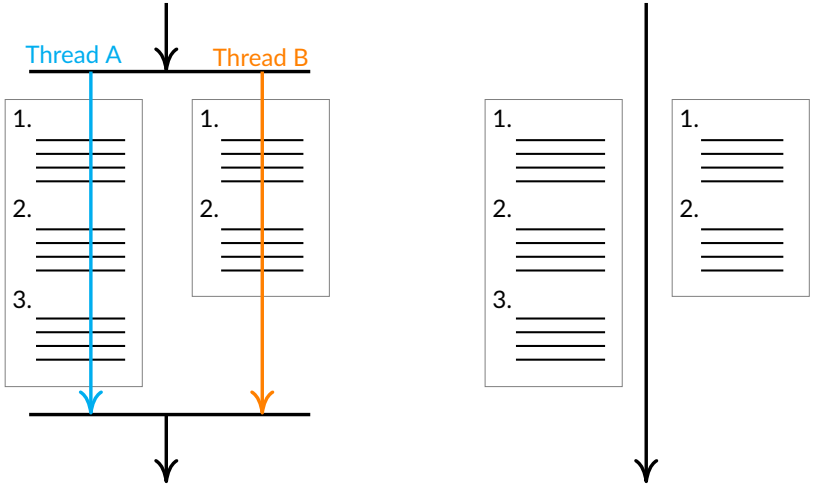

We implemented event-based tracking (4/4)

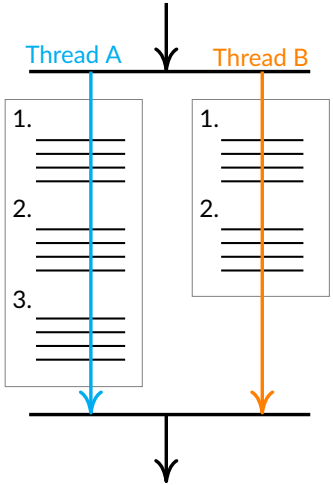

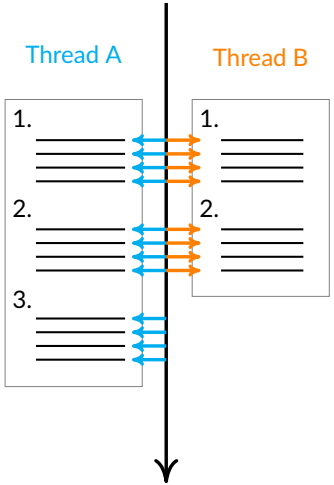

We build with link-time optimization (1/4)

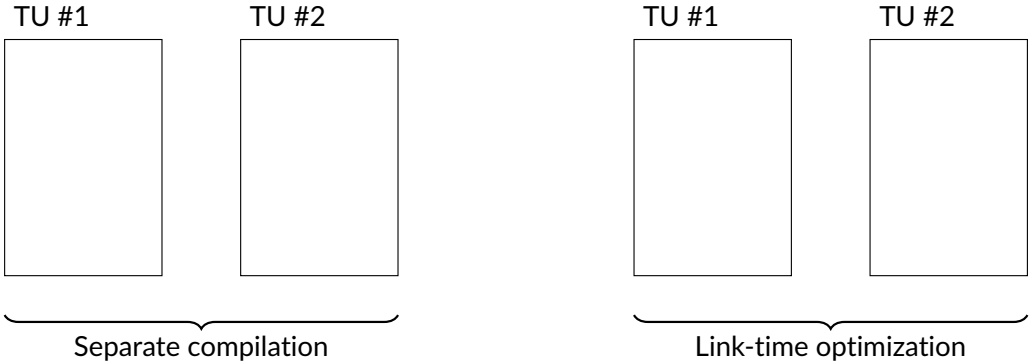

We build with link-time optimization (2/4)

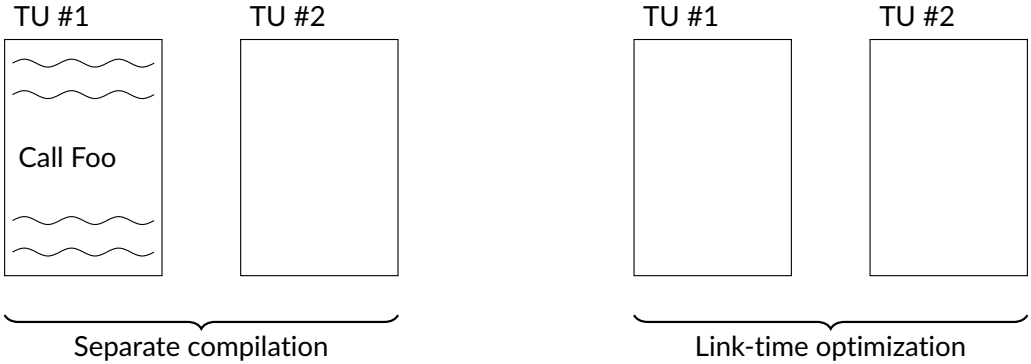

We build with link-time optimization (3/4)

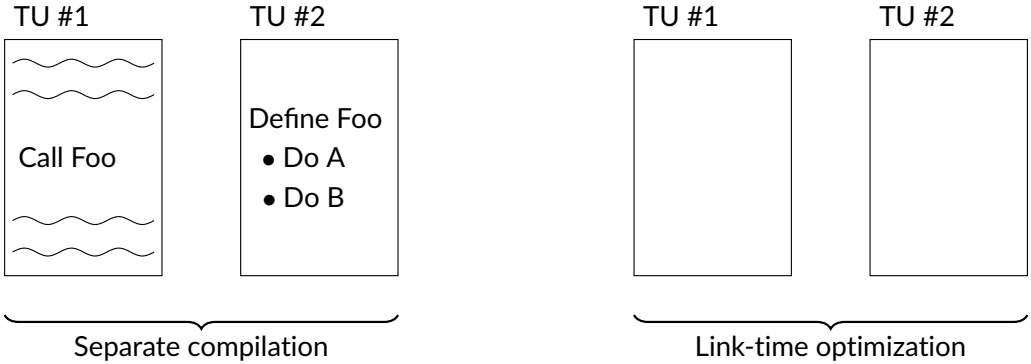

We build with link-time optimization (4/4)

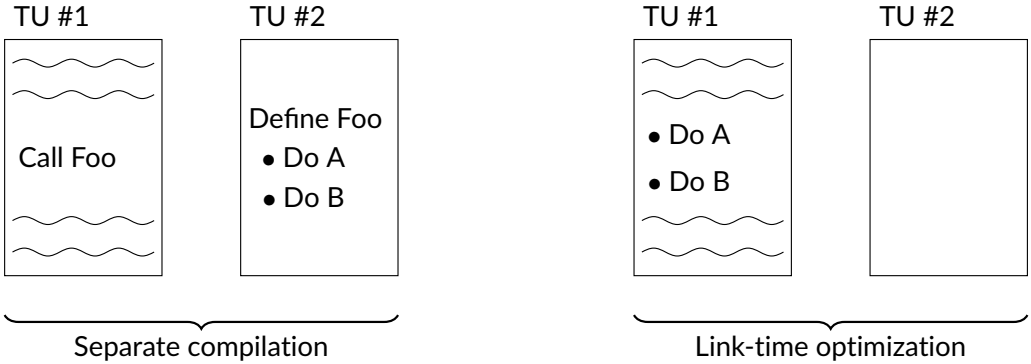

We ran two problems: Godiva in Water and Crooked Pipe (1/2)

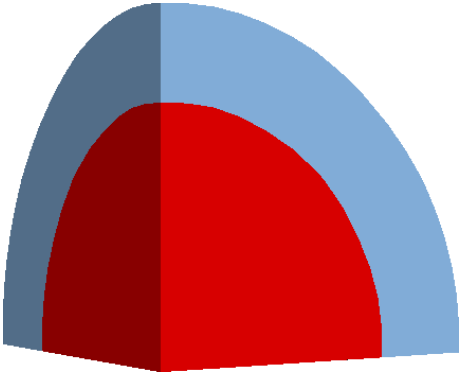

Mercury - Godiva in Water

We ran two problems: Godiva in Water and Crooked Pipe (2/2)

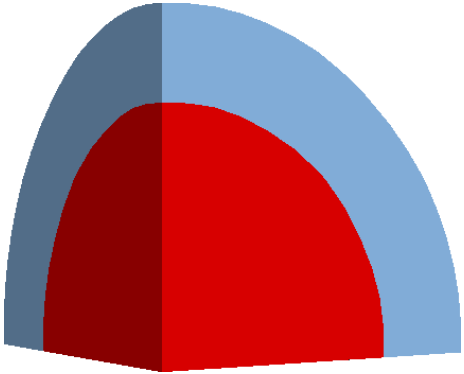

Mercury - Godiva in Water

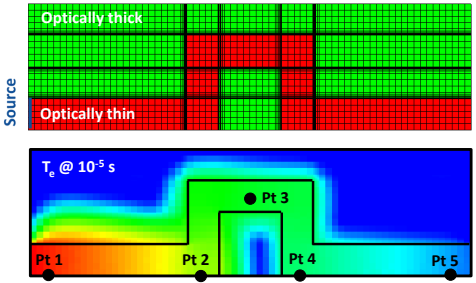

Imp - Crooked Pipe

We need 1000x more particles to saturate a Sierra node

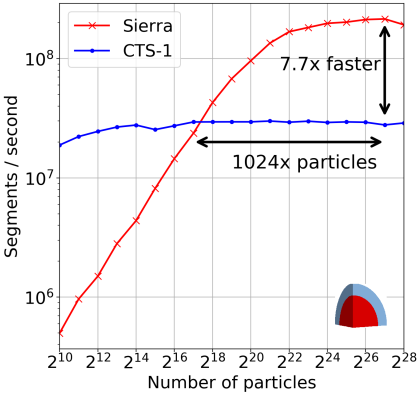

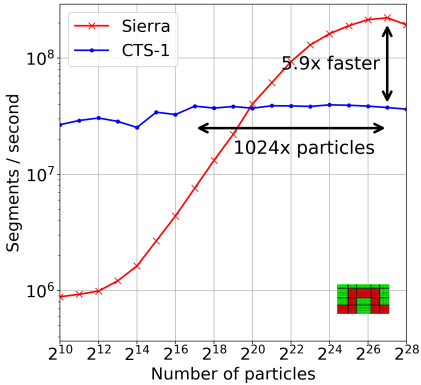

Event-based tracking and LTO provide substantial speedups (1/9)

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup

Event-based tracking and LTO provide substantial speedups (2/9)

CTS-1
Sierra
Sierra
Sierra

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup

CTS-1
Sierra
Sierra
Sierra

Event-based tracking and LTO provide substantial speedups (3/9)

CTS-1	History						
Sierra	History						
Sierra	Event						
Sierra	Event						

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup

CTS-1	History						
Sierra	History						
Sierra	Event						
Sierra	Event						

Event-based tracking and LTO provide substantial speedups (4/9)

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup
CTS-1	History	No					
Sierra	History	No					
Sierra	Event	No					
Sierra	Event	Yes					

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup
CTS-1	History	No					
Sierra	History	No					
Sierra	Event	No					
Sierra	Event	Yes					

Event-based tracking and LTO provide substantial speedups (5/9)

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup
CTS-1	History	No	6170				
Sierra	History	No	3744				
Sierra	Event	No	1085				
Sierra	Event	Yes	811				

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup
CTS-1	History	No	2711				
Sierra	History	No	913				
Sierra	Event	No	555				
Sierra	Event	Yes	467				

Event-based tracking and LTO provide substantial speedups (6/9)

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup
CTS-1	History	No	6170	6006			
Sierra	History	No	3744	3710			
Sierra	Event	No	1085	1052			
Sierra	Event	Yes	811	775			

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup
CTS-1	History	No	2711	2334			
Sierra	History	No	913	836			
Sierra	Event	No	555	477			
Sierra	Event	Yes	467	395			

Event-based tracking and LTO provide substantial speedups (7/9)

CTS-1	History	No	6170	6006	167.113		
Sierra	History	No	3744	3710	167.113		
Sierra	Event	No	1085	1052	167.113		
Sierra	Event	Yes	811	775	167.113		

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup
CTS-1	History	No	2711	2334	87.707		
Sierra	History	No	913	836	87.709		
Sierra	Event	No	555	477	87.706		
Sierra	Event	Yes	467	395	87.708		

Event-based tracking and LTO provide substantial speedups (8/9)

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup
CTS-1	History	No	6170	6006	167.113	27.8	
Sierra	History	No	3744	3710	167.113	45.0	
Sierra	Event	No	1085	1052	167.113	158.8	
Sierra	Event	Yes	811	775	167.113	215.5	

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup
CTS-1	History	No	2711	2334	87.707	37.6	
Sierra	History	No	913	836	87.709	104.9	
Sierra	Event	No	555	477	87.706	183.7	
Sierra	Event	Yes	467	395	87.708	222.3	

Event-based tracking and LTO provide substantial speedups (9/9)

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup
CTS-1	History	No	6170	6006	167.113	27.8	1.000
Sierra	History	No	3744	3710	167.113	45.0	1.619
Sierra	Event	No	1085	1052	167.113	158.8	5.707
Sierra	Event	Yes	811	775	167.113	215.5	7.745

System	Algo	LTO?	Total Time	Tracking Time	Segments	Seg/s	Speedup
CTS-1	History	No	2711	2334	87.707	37.6	1.000
Sierra	History	No	913	836	87.709	104.9	2.791
Sierra	Event	No	555	477	87.706	183.7	4.888
Sierra	Event	Yes	467	395	87.708	222.3	5.915

Achieving 7.6x and 5.8x speedups required sustained porting effort

	Total Runtime Speedup				
	FY18	FY19	FY21	FY22	FY23
Godiva in Water	0.47x	0.81x	4.43x	5.15x	7.61x
Crooked Pipe	1.45x	1.99x	4.52x	4.90x	5.81x

Ongoing work includes continued Sierra GPU porting and El Capitan GPU porting.

References I

[1] R. Procassini, D. Cullen, G. Greenman, and C. Hagmann. "Verification and Validation of Mercury: A Modern, Monte Carlo Particle Transport Code." In *Proceedings of MC2005*. Chattanooga, TN, USA (2005).

[2] J. A. Fleck and J. D. Cummings. "An implicit Monte Carlo scheme for calculating time and frequency dependent nonlinear radiation transport." *Journal of Computational Physics*, **volume 8**, pp. 313–342 (1971).

[3] P. Brantley, N. Gentile, M. Lambert, M. McKinley, M. O'Brien, and J. Walsh. "A New Implicit Monte Carlo Thermal Photon Transport Capability Developed Using Shared Monte Carlo Infrastructure." In *Proceedings of M&C 2019*, pp. 564–577. Portland, OR, USA (2019).

[4] S. S. Vazhkudai, B. R. de Supinski, A. S. Bland, A. Geist, J. Sexton, J. Kahle, C. J. Zimmer, S. Atchley, S. Oral, D. E. Maxwell, V. G. V. Larrea, and A. Bertsch. "The Design, Deployment, and Evaluation of the CORAL Pre-Exascale Systems." SC '18 (2018).

[5] M. McKinley, R. Bleile, P. Brantley, S. Dawson, M. O'Brien, M. Pozulp, and D. Richards. "Status of LLNL Monte Carlo Transport Codes on Sierra GPUs." In *Proceedings of M&C 2019*, pp. 2160–2165. Portland, OR, USA (2019).

[6] F. B. Brown and W. R. Martin. "Monte Carlo methods for radiation transport analysis on vector computers." *Progress in Nuclear Energy*, **volume 14**(3), pp. 269–299 (1984).

[7] R. M. Bergmann and J. L. Vujić. "Algorithmic choices in WARP – A framework for continuous energy Monte Carlo neutron transport in general 3D geometries on GPUs." *Annals of Nuclear Energy*, **volume 77**, pp. 176–193 (2015).

[8] R. Bergmann, K. Rowland, N. Radnović, R. Slaybaugh, and J. Vujić. "Performance and accuracy of criticality calculations performed using WARP – A framework for continuous energy Monte Carlo neutron transport in general 3D geometries on GPUs." *Annals of Nuclear Energy*, **volume 103**, p. 334–349 (2017).

References II

[9] R. Bleile, P. Brantley, M. O'Brien, and H. Childs. "Algorithmic Improvements for Portable Event-Based Monte Carlo Transport Using the Nvidia Thrust Library." *Transactions American Nuclear Society*, **volume 115**, pp. 535–538 (2016).

[10] S. P. Hamilton, S. R. Slattery, and T. M. Evans. "Multigroup Monte Carlo on GPUs: Comparison of history- and event-based algorithms." *Annals of Nuclear Energy*, **volume 113**, pp. 506–518 (2018).

[11] S. Hamilton and T. Evans. "Continuous-energy Monte Carlo neutron transport on GPUs in the Shift code." *Annals of Nuclear Energy*, **volume 128**, pp. 236–247 (2019).

[12] N. Choi, K. M. Kim, and H. Joo. "Optimization of neutron tracking algorithms for GPU-based continuous energy Monte Carlo calculation." *Annals of Nuclear Energy*, **volume 162**, p. 108508 (2021).

[13] D. A. Beckingsale, M. J. McFadden, J. P. S. Dahm, R. Pankajakshan, and R. D. Hornung. "Umpire: Application-focused management and coordination of complex hierarchical memory." *IBM Journal of Research and Development*, **volume 64**(3), pp. 1–10 (2020).

[14] S. P. Hamilton, T. M. Evans, and S. R. Slattery. "Continuous-Energy Monte Carlo Neutron Transport on GPUs in Shift." *Transactions American Nuclear Society*, **volume 118**, pp. 401–403 (2018).

[15] D. F. Richards and B. S. Ryujin. "Enhancements supporting IC usage of PEM libraries on next-gen platforms." (2021).

[16] M. Pozulp, S. Dawson, R. Bleile, P. Brantley, M. S. McKinley, M. O'Brien, and D. Richards. "Transitioning the Scientific Software Toolchain to Clang/LLVM." Paris, France (2020).

References III

[17] D. E. Cullen, C. J. Clouse, R. Procassini, and R. C. Little. "Static and Dynamic Criticality: Are They Different?" Technical report, UCRL-TR-201506 (2003).

[18] F. Graziani and J. LeBlanc. "The Crooked Pipe Test Problem." Technical report, UCRL-MI-143393 (2000).

[19] "LLNL and HPE to partner with AMD on El Capitan, projected as world's fastest supercomputer." (2020). URL `https://www.llnl.gov/news/llnl-and-hpe-partner-amd-el-capitan-projected-worlds-fastest-supercomputer`.