HYBRID METHODS FOR RADIATION TRANSPORT USING INTEGRAL DEFERRED CORRECTION

By

Michael M. Crockatt

A DISSERTATION

Submitted to
Michigan State University
in partial fulfillment of the requirements
for the degree of

Computational Mathematics, Science and Engineering – Doctor of Philosophy

ABSTRACT

HYBRID METHODS FOR RADIATION TRANSPORT USING INTEGRAL DEFERRED CORRECTION

Bv

Michael M. Crockatt

This thesis provides roughly three contributions to the study of defect correction methods and hybrid methods for radiation transport. First, a modification of traditional integral deferred correction (IDC) time-integration schemes is proposed that significantly reduces the storage requirements of the methods. These methods, which we call low-storage IDC or LS-IDC methods, require storing only one copy of each stage vector throughout the iteration process, whereas traditional IDC methods require two copies of each vector. The convergence and stability properties of the methods are examined in a variety of settings, and both analytical and experimental results are provided. A nonlinear ODE and a linear transport equation are used to compare the accuracy and storage requirements of LS-IDC integrators with other fully implicit schemes including diagonally implicit Runge-Kutta (DIRK) and space-time discontinuous Galerkin (STDG) methods. The results demonstrate that LS-IDC methods have similar accuracy but significantly reduced memory requirements compared to other fully implicit methods.

Second, extensions of a collision-based hybrid method for time-dependent radiation transport simulations are discussed. The hybrid methods are constructed by splitting the radiation flux into collided and uncollided components to which low- and high-resolution discrete ordinates approximations are applied, respectively. The use of arbitrarily high-order numerical approximations is emphasized, with particular attention paid to high-order time-integration schemes. A range of time-integrators are considered including DIRK, STDG, IDC, and LS-IDC methods of up to fifth-order accuracy. Convergence results in one spatial dimension are provided, and it is found that while the hybrid methods exhibit convergence stagnation and order reduction in certain scenarios, the overall accuracy of the hybrid approximations is comparable to standard discrete ordinates approximations in many cases. A test problem in two-dimensional xy-geometry consisting of a mock-up of a standard hohlraum configuration is used to compare the computational efficiency of the hybrid methods with standard discrete ordinates methods. It is observed that replacing a standard discrete ordinates method using an angular quadrature of order N with a hybrid discrete ordinates method using angular quadratures of order 2N and N/2 for the uncollided and collided fluxes, respectively, usually reduces solve time by a factor of 2 or more and in many cases also yields a reduction in solution error by a factor of up to 2. However, it is noted that the

specified hybrid methods require approximately 4.25 times the storage of the corresponding standard discrete ordinates methods.

Finally, two mechanisms for increasing the effectiveness of the hybrid methods are presented. The first is a reconstruction procedure for mapping between arbitrary discrete ordinates quadratures within the context of these hybrid methods. This approach, called Nyström reconstruction, is shown to be significantly more accurate than previous reconstruction methods. When applied to the two-dimensional hohlraum problem, it is observed that replacing a standard discrete ordinates method using an angular quadrature of order N with a hybrid discrete ordinates method using a Nyström reconstruction procedure and angular quadratures of order N and N/4 for the uncollided and collided fluxes, respectively, consistently reduces solve time by a factor of between 4 and 8 while increasing memory requirements by only 6% and producing little to no increase in solution error. Lastly, variations of hybrid methods using IDC integrators are described in which the hybrid approach is written as a two-grid iterative method in angle that is combined with an IDC time-integration scheme. It is demonstrated that the resulting methods are able to iteratively reduce the error due to the application of discrete ordinates quadratures of different resolutions to the collided and uncollided components.

Copyright by MICHAEL M. CROCKATT 2018

ACKNOWLEDGEMENTS

I would like to express my appreciation for the many individuals I have had the pleasure of interacting with throughout my graduate studies. In particular, I would like to thank my advisor Andrew Christlieb for his patience and guidance: I would never have been able to complete this thesis without his consistent support. I would like to express my gratitude to Cory Hauck for hosting me at Oak Ridge National Laboratory during the past five summers: I have greatly valued his mentorship and I have learned a great deal during my summers in Oak Ridge. I would also like to thank Kris Garrett, from whom I have learned a great deal about radiation transport, numerical methods, and scientific computing over the last several years. Additionally, I would like to thank all of the members of my thesis committee for lending their time, their ears, and their expertise to improving my thesis and increasing my understanding of how radiation transport models are used in practical settings. Finally, I would like to thank all of the students, postdocs, researchers, and support staff that I have interacted with at Michigan State University and Oak Ridge National Laboratory for fostering a welcoming research environment that I am grateful to have been a part of.

This work was supported in part by Michigan State University through computational resources provided by the Institute for Cyber-Enabled Research, and by Oak Ridge National Laboratory (ORNL) and Oak Ridge Associated Universities (ORAU) through the ORAU/ORNL High Performance Computing (HPC) Grant Program. This thesis and the numerical software developed during the course of this work were written using free and open-source software on a GNU/Linux system. Software development for this work extensively leveraged the GNU toolchain distributed by the Free Software Foundation, and made use of a number of numerical libraries including the Portable, Extensible Toolkit for Scientific Computation (PETSc), the Armadillo C++ library for linear algebra & scientific computing, and the Boost C++ libraries. This thesis was typeset using the TeX Live distribution of the LATEX typesetting system, with figures generated using the gnuplot graphing utility.

TABLE OF CONTENTS

LIST OF	TABLE	S	viii
LIST OF	FIGUR	RES	ix
LIST OF	ALGOI	RITHMS	xiii
СНАРТЕ	R 1	INTRODUCTION	1
1.1	Motiv	ation	1
1.2	Radia	tion transport	4
	1.2.1	Introduction	4
	1.2.2	Structure of solutions	6
	1.2.3	Standard methods	10
	1.2.4	Hybrid methods	12
1.3	Defec	t correction	14
	1.3.1	Introduction	14
	1.3.2	Integral deferred correction	17
СНАРТЕ	R 2	LOW-STORAGE INTEGRAL DEFERRED CORRECTION METHODS FOR SCIEN-	
		TIFIC COMPUTING	18
2.1	Introd	luction	18
2.2	Metho	od descriptions	19
	2.2.1	Collocation	19
	2.2.2	Integral deferred correction	20
	2.2.3	Low-storage integral deferred correction	22
2.3	Conve	ergence and stability	24
	2.3.1	Convergence in the limit of small timestep size	24
	2.3.2	Convergence for arbitrary timestep sizes	27
	2.3.3	Linear stability	30
2.4	Nume	·	32
	2.4.1	A non-stiff, non-linear ODE	36
	2.4.2		37
СНАРТЕ	R 3	HIGH-ORDER HYBRID METHODS FOR RADIATION TRANSPORT	42
3.1	The tr	ansport equation and discrete ordinates methods	42
	3.1.1	The transport equation	42
	3.1.2		43
	3.1.3		45
3.2	High-	·	48
	3.2.1	Diagonally implicit Runge-Kutta methods	50
	3.2.2	Space-time discontinuous Galerkin methods	54
	3.2.3	•	59
	3.2.4	-	64
3.3			66
	3.3.1		67
	3.3.2		73

	3.3.3	Spatial discretization	79
3.4	Nume	erical tests	82
	3.4.1	One spatial dimension	82
	3.4.2	Two spatial dimensions	88
CHAPTI	ER 4	NYSTRÖM RECONSTRUCTION AND TWO-GRID DEFECT CORRECTION FOR	
		HYBRID METHODS	127
4.1	Introd	luction	127
	4.1.1	Nyström interpolation	127
	4.1.2	Two-grid iteration	129
4.2	Metho	od descriptions	132
	4.2.1	Nyström reconstruction	132
	4.2.2	Two-grid IDC methods	142
	4.2.3	Two-grid LS-IDC methods	147
4.3	Nume	erical tests for Nyström methods	150
	4.3.1	One spatial dimension	150
	4.3.2	Two spatial dimensions	156
4.4	Nume	erical tests for two-grid methods	178
	4.4.1	One spatial dimension	
	4.4.2	Two spatial dimensions	185
CHAPTI	ER 5	CONCLUSIONS AND FUTURE WORK	230
5.1	Low-s	storage integral deferred correction methods for scientific computing	230
5.2		order hybrid methods for radiation transport	
RIRI IO	ZR A DLI	v	240

LIST OF TABLES

Table 2.1:	Contraction factors for IDC and LS-IDC methods in the stiff limit using different numbers and types of quadrature nodes with implicit Euler substeps.	29
Table 2.2:	Number of angular flux vectors required by each class of time integration scheme for solving the transport equation (1.2) using the discrete ordinates method. Here N represents the order of accuracy of the method, which may differ from the number of	
	quadrature nodes or stages used by the method	39

LIST OF FIGURES

Figure 2.1:	Contraction factors for IDC and LS-IDC methods along the negative real axis using different numbers and types of quadrature nodes with implicit Euler substeps	31
Figure 2.2:	Boundaries of the regions of absolute stability for IDC and LS-IDC methods using implicit Euler substeps on (right biased) Gauss-Radau quadrature nodes with various numbers of correction iterations. Regions of absolute stability lie to the <i>left</i> of the origin.	33
Figure 2.3:	Boundaries of the regions of absolute stability for IDC and LS-IDC methods using implicit Euler substeps on four Gauss-Lobatto quadrature nodes with various numbers of correction iterations. Regions of absolute stability lie to the <i>left</i> of the origin.	34
Figure 2.4:	Boundaries of the regions of absolute stability for IDC and LS-IDC methods using implicit Euler substeps on five Gauss-Lobatto quadrature nodes with various numbers of correction iterations. Regions of absolute stability lie to the <i>left</i> of the origin.	35
Figure 2.5:	Convergence of IDC, LS-IDC, and collocation integrators using three Gauss-Radau nodes, four Gauss-Lobatto nodes, and five right-biased equispaced nodes for the IVP (2.50). Errors are computed with respect to the ∞ -norm at time $t=10$. The number of correction iterations used with each IDC and LS-IDC method is indicated in the key of each plot	37
Figure 2.6:	Layout and sample solution for the lattice test problem for the transport equation (1.2). Color scale of sample solution is \log_{10}	38
Figure 2.7:	Memory usage vs. $L^2(d\vec{x})$ error in the scalar flux distribution for the lattice test problem.	41
Figure 3.1:	Butcher tableaus of SDIRK2 and SDIRK3 schemes	70
Figure 3.2:	Butcher tableaus of Kværnø7 and ARK8 schemes.	71
Figure 3.3:	L^2 convergence studies of hybrid-Ia and hybrid-IIa methods based on SDIRK2 and SDIRK3 integrators for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes	84

Figure 3.4:	L^2 convergence studies of hybrid-Ia and hybrid-IIa methods based on Kværnø7 and ARK8 integrators for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes	85
Figure 3.5:	L^2 convergence studies of hybrid-a methods based on STDG integrators for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes	86
Figure 3.6:	L^2 convergence studies of hybrid-I methods based on IDC5 and LS-IDC5 integrators for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.	87
Figure 3.7:	Layout and sample solution for the hohlraum test problem. Color scale of solution is \log_{10}	90
Figure 3.8:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK2 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	92
Figure 3.9:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK2 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution.	93
Figure 3.10:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK3 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	94
Figure 3.11:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK3 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution.	95

Figure 3.12	reference solution (bottom rows) of hybrid and non-hybrid Kværnø7 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.	96
Figure 3.13:	: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid Kværnø7 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution.	97
Figure 3.14:	: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid ARK8 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	98
Figure 3.15:	: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid ARK8 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	99
Figure 3.16	: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid STDG3 methods using CL and T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	100
Figure 3.17:	: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid STDG5 methods using CL and T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	101
Figure 3.18:	: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid IDC3 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	102
Figure 3.19	: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid IDC3 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	103
Figure 3.20	: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid IDC5 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	104

Figure 3.21:	reference solution (bottom rows) of hybrid and non-hybrid IDC5 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.	. 105
Figure 3.22:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid LS-IDC3 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.	. 106
Figure 3.23:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid LS-IDC3 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution.	. 107
Figure 3.24:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid LS-IDC5 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	. 108
Figure 3.25:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid LS-IDC5 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution.	. 109
Figure 3.26:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the SDIRK2 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	. 110
Figure 3.27:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the SDIRK3 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	. 111
Figure 3.28:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the Kværnø7 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	. 112
Figure 3.29:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the ARK8 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	. 113

Figure 3.30:	hybrid (red and blue) discrete ordinates methods based on the STDG3 and STDG5 integrators applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	114
Figure 3.31:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the IDC3 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	116
Figure 3.32:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the IDC5 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	117
Figure 3.33:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the LS-IDC3 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	118
Figure 3.34:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the LS-IDC5 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	119
Figure 3.35:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray) and hybrid (red and blue) discrete ordinates methods based on STDG integrators and various angular resolutions applied to the hohlraum test problem. CL (gray and blue) and T_N (black and red) angular quadratures are considered.	121
Figure 3.36:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray) and hybrid (red and blue) discrete ordinates methods based on DIRK integrators and various angular resolutions applied to the hohlraum test problem. CL (gray and blue) and T_N (black and red) angular quadratures are considered.	122
Figure 3.37:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray) and hybrid (red and blue) discrete ordinates methods based on DIRK integrators and various angular resolutions applied to the hohlraum test problem. CL (gray and blue) and T_N (black and red) angular quadratures are considered.	123

Figure 3	3.38: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray) and hybrid (red and blue) discrete ordinates methods based on IDC integrators and various angular resolutions applied to the hohlraum test problem. CL (gray and blue) and T_N (black and red) angular quadratures are considered 124
Figure 3	8.39: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray) and hybrid (red and blue) discrete ordinates methods based on LS-IDC integrators and various angular resolutions applied to the hohlraum test problem. CL (gray and blue) and T_N (black and red) angular quadratures are considered.
Figure 4	1.1: L^2 convergence studies of hybrid STDG methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes
Figure 4	1.2: L^2 convergence studies of SDIRK2 hybrid methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes 152
Figure 4	1.3: L^2 convergence studies of SDIRK3 hybrid methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes 153
Figure 4	1.4: L^2 convergence studies of Kværnø7 hybrid methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes 154
Figure 4	1.5: L^2 convergence studies of ARK8 hybrid methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes

Figure 4.6:	reference solution (bottom rows) of hybrid and non-hybrid SDIRK2 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.	158
Figure 4.7:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK2 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	159
Figure 4.8:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK3 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	160
Figure 4.9:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK3 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	161
Figure 4.10:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid Kværnø7 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	162
Figure 4.11:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid Kværnø7 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	163
Figure 4.12:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid ARK8 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	164
Figure 4.13:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid ARK8 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	165
Figure 4.14:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid STDG3 methods using CL and T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	166

Figure 4.15:	Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid STDG5 methods using CL and T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	167
Figure 4.16:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the SDIRK2 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	168
Figure 4.17:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the SDIRK3 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	169
Figure 4.18:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the Kværnø7 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	170
Figure 4.19:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the ARK8 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	171
Figure 4.20:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on STDG3 and STDG5 integrators applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales	172
Figure 4.21:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray), hybrid-a (red and blue), and hybrid-c (pink and teal) discrete ordinates methods based on STDG integrators and using various angular quadrature resolutions applied to the hohlraum test problem. CL (gray, blue, and teal) and T_N (black, red, and pink) angular quadratures are considered	175
Figure 4.22:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray), hybrid-Ia (red and blue), and hybrid-Ic (pink and teal) discrete ordinates methods based on DIRK integrators and using various angular quadrature resolutions applied to the hohlraum test problem. CL (gray, blue, and teal) and T_N (black, red, and pink) angular quadratures are considered	176

Figure 4.23:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray), hybrid-IIa (red and blue), and hybrid-IIc (pink and teal) discrete ordinates methods based on DIRK integrators and using various angular quadrature resolutions applied to the hohlraum test problem. CL (gray, blue, and teal) and T_N (black, red, and pink) angular quadratures are considered	177
Figure 4.24:	L^2 convergence studies of IDC5 hybrid-IIb and hybrid-IIc $S_{32,2}$ methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes	180
Figure 4.25:	L^2 convergence studies of LS-IDC5 hybrid-IIb and hybrid-IIc $S_{32,2}$ methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes	181
Figure 4.26:	L^2 convergence studies of IDC5 hybrid-IIb and hybrid-IIc $S_{32,4}$ methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes	182
Figure 4.27:	L^2 convergence studies of LS-IDC5 hybrid-IIb and hybrid-IIc $S_{32,4}$ methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes	183
Figure 4.28:	Deviation of scalar flux density from reference solution for hybrid-IIb IDC3 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	188
Figure 4.29:	Deviation of scalar flux density from reference solution for hybrid-IIc IDC3 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	189

Figure 4.30:	Deviation of scalar flux density from reference solution for hybrid-11b IDC3 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	190
Figure 4.31:	Deviation of scalar flux density from reference solution for hybrid-IIc IDC3 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	191
Figure 4.32:	Deviation of scalar flux density from reference solution for hybrid-IIb LS-IDC3 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	192
Figure 4.33:	Deviation of scalar flux density from reference solution for hybrid-IIc LS-IDC3 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	193
Figure 4.34:	Deviation of scalar flux density from reference solution for hybrid-IIb LS-IDC3 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution	194
Figure 4.35:	Deviation of scalar flux density from reference solution for hybrid-IIc LS-IDC3 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	195
Figure 4.36:	Deviation of scalar flux density from reference solution for hybrid-IIb IDC5 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	196
Figure 4.37:	Deviation of scalar flux density from reference solution for hybrid-IIc IDC5 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	197
Figure 4.38:	Deviation of scalar flux density from reference solution for hybrid-IIb IDC5 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	198

Figure 4.39:	Deviation of scalar flux density from reference solution for hybrid-IIc IDC5 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	.99
Figure 4.40:	Deviation of scalar flux density from reference solution for hybrid-IIb LS-IDC5 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	200
Figure 4.41:	Deviation of scalar flux density from reference solution for hybrid-IIc LS-IDC5 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	201
Figure 4.42:	Deviation of scalar flux density from reference solution for hybrid-IIb LS-IDC5 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	202
Figure 4.43:	Deviation of scalar flux density from reference solution for hybrid-IIc LS-IDC5 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution	203
Figure 4.44:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue	204
Figure 4.45:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue	205
Figure 4.46:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue	206

Figure 4.47:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue
Figure 4.48:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue
Figure 4.49:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue
Figure 4.50:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue
Figure 4.51:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue
Figure 4.52:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the LS-IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue
Figure 4.53:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the LS-IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue

	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the LS-IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue	214
	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the LS-IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue	215
	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the LS-IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue	216
Ü	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the LS-IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue	217
	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the LS-IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue	218
	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the LS-IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue	219
-	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (blue, teal, and green) discrete ordinates methods based on the IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations	222

Figure 4.61:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (red, pink, and orange) discrete ordinates methods based on the IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations	223
Figure 4.62:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (blue, teal, and green) discrete ordinates methods based on the IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations	224
Figure 4.63:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (red, pink, and orange) discrete ordinates methods based on the IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations	225
Figure 4.64:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (blue, teal, and green) discrete ordinates methods based on the LS-IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations	226
Figure 4.65:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (red, pink, and orange) discrete ordinates methods based on the LS-IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations	227
Figure 4.66:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (blue, teal, and green) discrete ordinates methods based on the LS-IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations	228
Figure 4.67:	Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (red, pink, and orange) discrete ordinates methods based on the LS-IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations	229
Figure 5.1:	Layout and sample solution for the inhomogeneous sphere problem. Color scale of solution is \log_{10}	237
Figure 5.2:	L^2 -convergence studies of hybrid-Ic SDIRK3 methods and hybrid-IIc LS-IDC3 methods applied to the inhomogeneous sphere problem with a final time of 8 and a CFL of 128. All solutions are computed using T_N quadratures	238

LIST OF ALGORITHMS

Algorithm 3.1:	Hybrid Euler timestep	49
Algorithm 3.2:	Nonhybrid DIRK timestep	51
Algorithm 3.3:	Hybrid-Ia DIRK timestep.	52
Algorithm 3.4:	Hybrid-IIa DIRK timestep	53
Algorithm 3.5:	Hybrid-a STDG timestep.	58
Algorithm 3.6:	Nonhybrid IDC timestep.	60
Algorithm 3.7:	Hybrid-Ia IDC timestep	62
Algorithm 3.8:	Hybrid-Ib IDC timestep	63
Algorithm 3.9:	Nonhybrid LS-IDC timestep.	66
Algorithm 3.10:	Hybrid-Ia LS-IDC timestep	68
Algorithm 3.11:	Hybrid-Ib LS-IDC timestep	69
Algorithm 4.1:	Hybrid Euler timestep with Nyström reconstruction	34
Algorithm 4.2:	Hybrid-I DIRK timestep with Nyström reconstruction	36
Algorithm 4.3:	Hybrid-II DIRK timestep with Nyström reconstruction	.37
Algorithm 4.4:	Hybrid STDG timestep with Nyström reconstruction	.39
Algorithm 4.5:	Hybrid-Ia and hybrid-Ic IDC timestep	41
Algorithm 4.6:	Hybrid-II IDC and LS-IDC prediction	45
Algorithm 4.7:	Hybrid-II IDC correction	48
Algorithm 4.8:	Hybrid-II LS-IDC correction	49

CHAPTER 1

INTRODUCTION

1.1 Motivation

The behavior of many systems in science and engineering can be described by a diverse array of models of various levels of complexity. Many models rely on certain assumptions about the state of a system in order to reduce the complexity of the model. When these assumptions are satisfied, or at the very least approximately satisfied, it is possible to describe the dynamics of a system that are relevant to a particular situaion using a less complex model. However, applying a model to a system that does not satisfy the assumptions required by that model often results in behavior that deviates significantly from that of the system under consideration.

The assumptions that are used to construct various models are often related to certain scales within the system. More specifically, these assumptions often describe relative differences between the magnitudes of certain physical features or the lengths of time over which certain processes occur. In many cases, different models for a given system can be arranged in sequence based on the characteristic scales for which assumptions are made, with "microscale" models for small characteristic scales on one end and "macroscale" models for large characteristic scales on the other [89]. The determination of which phenomena are accurately captured by a given model is usually governed by the relative differences between the assumed scales of the model and the characteristic scales on which each phenomenon occurs.

Using the most detailed model available might be desirable from an idealistic perspective, but this is generally not realistically feasible. While one may be able to write out a set of mathematical expressions to describe the behavior of a system using a given model, solving the resulting system of equations can be exceptionally difficult. Exact, analytic solutions are known only for relatively simple models or reduced scenarios that are significantly less complex than those found in most practical applications. Solving more complex models under realistic conditions requires using alternative strategies such as numerical approximations. Numerical approaches have the advantage that approximate solutions for very complex models can be obtained to, in theory, an arbitrary level of precision. There are nevertheless practical limitations to what can be computed. While significant advances in computing machinery have been

made since the development of the electronic computer in the twentieth century, many calculations will remain intractable for the foreseeable future due to the sheer number of operations that they require.

In light of these facts, the model that one selects for a particular application should balance the accuracy of the model under the conditions specified by the targeted application with the resources available to construct solutions (numerical or otherwise) to the equations that are used to describe the model. Numerical simulations of some systems can be performed relatively easily: there exist simple models that are sufficiently descriptive for the intended application and numerical methods are known that produce adequate approximations to the equations describing these models with a small computational cost. However, many complex systems display radically different behavior in different situations, and capturing the diverse array of phenomena that these systems exhibit using a single model and a single uniform numerical approximation is often prohibitively expensive.

This has led to the development of a variety of approaches for combining different models or different numerical methods in the hope that these combinations will be able to more efficiently describe the behavior of systems that exhibit complex behavior. These approaches are often characterized using qualifiers such as *hybrid*, *multiscale*, or *adaptive*. In many cases these methods are constructed in roughly three steps: (i) divide the problem into different components that isolate each of the important regimes, (ii) apply different models or numerical approximations to each component that are tailored towards their respective regimes, and (iii) devise a mechanism for interfacing between the different models or numerical representations that are used for each component. There are challenges associated with each of these steps that depend on the particular application and methods under consideration.

Combining methods that are tailored strongly towards disparate regimes is perhaps counterproductive when intermediate regimes are also present. While the accuracy of the methods that are used may be acceptable when conditions lie firmly within their respective regimes, accuracy can degrade rapidly as conditions move away from those regimes. Because of this, a significant loss of accuracy can occur when combined approaches are applied to problems that contain conditions between those that are the focus of the combined approximation.

There are two primary motivations for the present work. The first is to investigate extensions of a collision-based hybrid method for time-dependent radiation transport simulations that was recently described in [97]. This hybrid approach is based on the *first-collision source* technique, which divides the solution into streaming and collisional components to which different models and numerical approxi-

mations can be applied (c.f., section 1.2). The general idea is to apply a method such as ray tracing to the streaming component and something closer to a fluid approximation to the collisional component. This work extends that of [97] by considering hybrid methods using arbitrarily high-order discretizations in space and time and arbitrarily accurate angular quadratures of different types. Additionally, these methods are implemented in a high-performance C++ research code that incorporates robust solver techniques such as GMRES [168].

The second motivation for this work is to investigate the use of defect correction methods to increase the effectiveness of combined or hybrid approaches. Defect correction methods are based on the idea that an approximate solution for a problem can be iteratively refined towards a solution with a higher level of accuracy by using a method that is less accurate, but also less expensive. Our goal is to consider applying the principle of defect correction to the models that are used to describe the behavior of a system in addition to the numerical approximations that are applied to these models. In this case the desired high-accuracy solution would be described using a more detailed model of the underlying system, and the lower-accuracy solution that is used would be composed of a combination of models that are applicable to one or more limiting regimes present in the system (e.g., the streaming/collisional decomposition mentioned previously). The hybrid methods considered in this work are constructed based on a combination of high- and low-resolution numerical approximations that are applied to complementary models of radiation transport. While the impact of full physical model reduction is not explicitly explored in this work, it should be noted that the separation between reduced physical models and very coarse numerical approximations is often not fully well-defined. As such, we believe that this work provides a strong first step in the exploration of these ideas in the context of radiation transport applications. Future extensions of this work will seek to examine the impact of further reductions in model complexity in the context of defect correction methods.

This thesis provides roughly three contributions to the study of defect correction methods and hybrid methods for radiation transport, each of which corresponds to one of chapters 2 to 4. First, low-storage integral deferred correction (LS-IDC) methods are presented in chapter 2 that require approximately half the storage required by traditional integral deferred correction (IDC) methods. The importance of this contribution is related to the significant memory requirements imposed by the high-dimensional nature of certain models, including kinetic models of radiation transport. In chapter 3, extensions of a collision-based hybrid method for time dependent radiation transport simulations are discussed, with a

focus on the use of arbitrarily high-order numerical approximations including, in particular, high-order time-integration schemes. In chapter 4, two mechanisms for increasing the effectiveness of the hybrid methods are presented. The first is a reconstruction procedure for mapping between arbitrary discrete ordinates quadratures within the context of these hybrid methods. The second is a variation of the previous hybrid methods using IDC integrators in which the hybrid approach is written as a two-grid iterative method in angle that is combined with an IDC time-integration scheme. It is demonstrated that the resulting methods are able to iteratively reduce the error due to the application of discrete ordinates quadratures of different resolutions to the components of the hybrid method.

1.2 Radiation transport

1.2.1 Introduction

Radiation transport models are used to describe the behavior of charge-neutral particles as they propagate through a medium. These models are commonly applied to a diverse range of particle species including neutrons [62, 135], neutrinos [141], photons [160], and phonons [180]. In many applications, these components are responsible for a significant portion of the transfer of energy that occurs between different regions or constituent parts of the system.

Radiation transport models are typically composed of two main parts: particle streaming and particle collisions. These models assume that a particle traveling in some direction through a medium will continue to travel in that direction (particle streaming) until it collides with the material it is traveling through (particle collision). When a collision event occurs, energy is exchanged between the particle and the background material in a manner that is determined by the physical properties of the system under consideration. Interactions between the radiation particles themselves may also be considered, but in many applications these types of interactions are neglected, either because they have a relatively small impact on the overall behavior of the system or because including these effects significantly increases the complexity of the model. In many cases the importance of these effects depends on the scales of interest.

Kinetic descriptions of radiation transport are commonly expressed in terms of a first-order integrodifferential equation of the form

$$\frac{1}{\nu}\partial_{t}f\left(\vec{x},\vec{\Omega},E,t\right) + \vec{\Omega}\cdot\nabla_{\vec{x}}f\left(\vec{x},\vec{\Omega},E,t\right) = \mathcal{C}\left[f\right],\tag{1.1}$$

where the distribution function f depends on as many as seven variables: time t > 0, energy E > 0, position $\vec{x} \in \mathbb{R}^3$, and a unit vector $\vec{\Omega} \in \mathbb{S}^2$ specifying direction of travel. Here v > 0 is the particle speed, which may depend on various factors such as the type and energy E of the particle. The terms on the left of (1.1) express particle streaming and the operator C on the right of (1.1) is the *collision operator* that is used to express the behavior of collision events; i.e., radiation-material interactions, interactions between radiation particles, etc.

There are a number of factors that make solving equations of the form (1.1) difficult. In the absence of the collision term, the solution of (1.1) can be written explicitly for each direction $\vec{\Omega}$ and energy E independently following a simple ray-tracing approach. However, this approach no longer works when the collision term is included because the collision operator typically couples the particle directions and energies. Additionally, many problems of practical interest require dynamic interaction between the radiation components and the background material. For example, in radiation hydrodynamics applications the collision operator $\mathcal C$ couples the transport equation (1.1) to a hydrodynamic model for the evolution of the background material [143, 160]. The coupling between the radiation and background material is often highly nonlinear in nature, resulting in a significant degree of numerical stiffness when the radiation and background material evolve on different timescales. This combination of high-dimensionality and numerical stiffness renders radiation transport problems very computationally expensive to solve in practice.

In this work we consider a simple linear, mono-energetic, time-dependent transport equation of the form

$$\partial_t \psi + \vec{\Omega} \cdot \nabla_{\vec{X}} \psi + \sigma_t \psi = \frac{\sigma_s}{4\pi} \langle \psi \rangle + q, \qquad \langle \cdot \rangle = \int_{\mathbb{S}^2} (\cdot) \, d\vec{\Omega}, \qquad (1.2)$$

where $\psi = \psi(\vec{x}, \vec{\Omega}, t)$ is flux of radiation at time $t \ge 0$ and position $\vec{x} \in X \subset \mathbb{R}^3$ in the direction $\vec{\Omega} \in \mathbb{S}^2$. The parameters σ_s , σ_a , and $\sigma_t = \sigma_s + \sigma_a$ are the scattering, absorption, and total cross sections of the material, respectively, and $q = q(\vec{x}, \vec{\Omega}, t)$ is a source of radiation that is assumed to be known for all times $t \ge 0$. The equation (1.2) is assumed to be normalized such that particles travel with unit speed, and the background material is assumed to be fixed in time: specifically, the cross sections of the material are assumed to be non-negative functions of position only. The model (1.2) accounts for at most three forms of interaction between the radiation and background material: scattering, absorption, and emission. The

collision operator \mathcal{C} can be written in the following form that separates the contributions of these effects:

$$C[\psi] = \sigma_{\rm S} \left(\frac{1}{4\pi} \langle \psi \rangle - \psi \right) - \sigma_{\rm a} \psi + q. \tag{1.3}$$

Note that the scattering of radiation off of the background material is assumed to be isotropic.

The equation (1.2) has a relatively simple form compared to that of more advanced models, but it retains many of the same characteristics that make solving systems of this type difficult. Even though energy-dependence is excluded, this system still contains a total of six dimensions. Hence storing numerical representations of solutions can quickly become prohibitively expensive as the resolution of the numerical approximations that are applied is increased. In this work only fully implicit time discretizations will be considered. This is motivated by the fact that in many applications where coupling between the radiation and background material is present the radiation components evolve on timescales that are much faster than that of the background material: Restricting global timestep sizes based on the evolution of the radiation components is usually prohibitively expensive in such situations. However the use of implicit methods poses a challenge due to the high dimensionality of the system and the need to solve the equations that define the updated solution at each timestep. The simplicity of (1.2) does eliminate some of the difficulties present in more advanced models. Most notably, (1.2) is linear and lacks the degree of stiffness commonly encountered when nonlinear material coupling is present. It remains to be determined whether the methods introduced in this work are sufficiently robust for cases where material coupling is included.

The remainder of this section is organized as follows: Section 1.2.2 examines some mathematical properties of solutions of the transport equation (1.2) that are intimately tied to the hybrid approach used in later chapters. Section 1.2.3 provides a very brief overview of common numerical methods applied to the angular variable $\vec{\Omega}$ of (1.2). Section 1.2.4 discusses some hybrid methods that share similarities with the approach that is used in this work.

1.2.2 Structure of solutions

The purpose of this section is to outline some important properties of solutions of (1.2). Two characteristics in particular are discussed: the *diffusion limit* for strong scattering and vanishing mean-free-path, and a series form of the solution known as the *multiple collision expansion*.

The diffusion limit

As mentioned previously, the kinetic equation (1.1) is composed of two parts: particle streaming on the left and the collision term on the right. The streaming terms on the left have the form of a simple hyperbolic conservation law, while the collision term on the right serves to relax the solution of the system towards a local equilibrium state. Systems of this form can be more generally described under the umbrella of *hyperbolic conservation laws with relaxation terms* [41]. The strength of the relaxation term determines the extent to which the solution of the system is able to relax towards an equilibrium state over a given length of time.

Knowledge of the equilibrium states is essential to understanding the overall behavior of systems containing relaxation terms. One application of this knowledge is the derivation of reduced models for scenarios where the scales of interest are much larger than the *relaxation time* – the time over which the system is able to relax towards an equilibrium state. In kinetic theory, the relaxation time is related to the particle mean-free-path or mean-free-time, and a hydrodynamic model is obtained in the limit as the mean-free-path approaches zero. For example, hydrodynamic models such as the Euler or Navier-Stokes equations can be derived from Boltzmann's equation for dilute gases in this manner [38]. The standard reduced model for the transport equation (1.2) is a diffusion equation [90, 120], the derivation of which we now briefly recall.

The diffusion limit of (1.2) is used to describe the long-time behavior of systems that are dominated by particle scattering. These situations occur when the total cross section is large relative to the times of interest, and the contributions of particle absorption and emission are small. This is described formally using the following scaling:

$$\sigma_{t} \mapsto \frac{\sigma_{t}}{\varepsilon}, \qquad \sigma_{a} \mapsto \varepsilon \sigma_{a}, \qquad q \mapsto \varepsilon q, \qquad t \mapsto \frac{t}{\varepsilon}.$$
 (1.4)

Under the asymptotic scaling (1.4), the transport equation (1.2) becomes

$$\varepsilon \partial_t \psi + \vec{\Omega} \cdot \nabla_{\vec{x}} \psi + \frac{\sigma_t}{\varepsilon} \psi = \frac{1}{4\pi} \left(\frac{\sigma_t}{\varepsilon} - \varepsilon \sigma_a \right) \phi + \varepsilon q, \tag{1.5}$$

where $\phi = \langle \psi \rangle$ and all quantities in (1.5) are assumed to be O(1) with respect to ε . We have the following result for the solution of the scaled equation (1.5):

Proposition 1. If ψ solves (1.5) then ϕ satisfies the diffusion equation

$$\partial_t \phi - \nabla_{\vec{x}} \cdot \left(\frac{1}{3\sigma_t} \nabla_{\vec{x}} \phi \right) + \sigma_a \phi = \langle q \rangle \tag{1.6}$$

up to an $O(\varepsilon)$ error.

Proof. We do not consider contributions from initial or boundary layers: For more complete discussions including these components see [90, 120]. To simplify notation slightly, we let $\widetilde{\psi} = \psi - \frac{1}{4\pi} \phi$ and $\vec{J} = \langle \vec{\Omega} \psi \rangle$. The limiting behavior can be derived by substituting a Hilbert expansion in powers of ε

$$\psi = \psi^{(0)} + \varepsilon \psi^{(1)} + \varepsilon^2 \psi^{(2)} + \cdots \tag{1.7}$$

into (1.5) and equating terms for each power of ε . This gives,

at $O(\varepsilon^{-1})$:

$$\psi^{(0)} = \frac{1}{4\pi} \phi^{(0)},\tag{1.8}$$

at O(1):

$$\vec{\Omega} \cdot \nabla_{\vec{r}} \psi^{(0)} + \sigma_{t} \widetilde{\psi}^{(1)} = 0, \tag{1.9}$$

and at $O(\varepsilon)$:

$$\partial_t \psi^{(0)} + \vec{\Omega} \cdot \nabla_{\vec{x}} \psi^{(1)} + \sigma_t \widetilde{\psi}^{(2)} + \frac{\sigma_a}{4\pi} \phi^{(0)} = q. \tag{1.10}$$

Multiplying (1.9) by $\vec{\Omega}$ and integrating over \mathbb{S}^2 yields, after some algebraic manipulation,

$$\frac{1}{3\sigma_t} \nabla_{\vec{x}} \phi^{(0)} + \vec{J}^{(1)} = 0. \tag{1.11}$$

Integrating (1.10) over \mathbb{S}^2 and using (1.11) to replace the resulting $\vec{J}^{(1)}$ term yields (1.6) after some simplification.

The importance of the diffusion limit is derived from the fact that many radiation transport applications contain collision scales that are much smaller than the timescales of interest. In such cases, the solution of the transport equation (1.2) will exhibit behavior that is similar to the solution of the diffusion equation (1.6). When numerical approximations are applied to (1.2), it is often important to consider how the numerical methods that are used behave under the asymptotic scaling (1.4) in the discrete setting. It would be desirable for a consistent discretization of (1.5) to produce a consistent discretization of (1.6) in the limit as $\varepsilon \to 0$. Many methods yield an inconsistent discretization in the limit, and methods that do yield a consistent discretization may succumb to order reduction or onerous timestep restrictions in limiting regimes [28]. Methods that do not yield a consistent discretization of the limiting

equation (1.6) as $\varepsilon \to 0$ are often severely inaccurate for problems containing highly-collisional regimes [3, 27, 164].

One of the advantages of implicit time discretizations is that they are generally more well-behaved in limiting regimes than explicit methods [121]. For example, applying an implicit Euler discretization to (1.5) and examining the behavior of the scheme for small ε yields an implicit Euler approximation of the limiting equation (1.6) up to an $O(\varepsilon)$ error (the proof of this is essentially the same as that of proposition 1). This is one of the reasons (in addition to the unconditional stability of the method) that the implicit Euler method remains the preferred time discretization for many radiation transport applications.

The multiple collision expansion

The multiple collision expansion is a series form of the solution of the transport equation that has useful physical and mathematical interpretations. In the multiple collision expansion, the angular flux distribution is written as a series

$$\psi(\vec{x}, \vec{\Omega}, t) = \sum_{n=0}^{\infty} \psi_n(\vec{x}, \vec{\Omega}, t)$$
 (1.12)

where the distribution functions ψ_n satisfy

$$\partial_t \psi_0 + \vec{\Omega} \cdot \nabla_{\vec{x}} \psi_0 + \sigma_t \psi_0 = q, \tag{1.13a}$$

$$\partial_t \psi_n + \vec{\Omega} \cdot \nabla_{\vec{x}} \psi_n + \sigma_t \psi_n = \frac{\sigma_s}{4\pi} \langle \psi_{n-1} \rangle, \qquad (n \ge 1). \tag{1.13b}$$

Much of the attractiveness of this form of the solution is derived from the fact that each of the ψ_n can be interpreted physically as the distribution of particles that have experienced exactly n collision events.

Early appearances of this form of the solution, such as in [161, Section 22] and [14], were motivated by physical interpretations in terms of collision counts. Generalizations [175] and analysis [176, 177] were later provided by Syros, who showed rigorously for the case of slab geometry that the series (1.12) is absolutely convergent and equal to the solution of (1.2) whenever the solution exists. Mathematically, the series (1.12) corresponds to the Liouville-Neumann series solution of the integral form of the transport equation [69].

One of the common applications of the multiple collision approach is the derivation of analytic or semi-analytic benchmark solutions that can be used for code verification. The multiple collision approach was first used by Kholin to obtain analytic solutions to time-dependent transport problems with

anisotropic scattering [114]. Ganapol later derived a closed-form expression of the solution and applied numerical techniques to the resulting expressions to explicitly obtain benchmark solutions [83]. A number of other benchmark solutions were later derived using similar techniques [80, 81, 84, 154], including some for charged-particle systems [78, 82] and Brownian motion [79].

1.2.3 Standard methods

A wide variety of numerical methods have been devised for computing approximate solutions for the transport equation (1.2). The purpose of this section is to provide a brief overview overview of some of the more common numerical methods that are used to treat the angular or directional dependence of the solution. This discussion is intentionally brief and by no means complete, though the references given in this section should provide a reasonable starting point for those new to the field.

One of the simplest classes of methods are diffusion approximations – see, e.g., [160, Chapter III] or [134, 149, 152] and references therein for some examples. Diffusion approximations effectively reduce the system to one that does not depend on the angular direction, and hence do not permit arbitrary refinement like other methods. The simplicity of most diffusion approximations often makes them much less expensive than other methods. These approximations work well for highly collisional or optically thick regimes where the solution is nearly isotropic anyway, but are ill-suited to cases where the behavior of the solution is not dominated by particle scattering.

The two most common approaches when higher fidelity approximations are required are Monte Carlo and discrete ordinates methods. Monte Carlo methods use a stochastic approach based on a sampling of particle histories – see, e.g., [135, Chapter 7] or [52, 87, 139, 186] and references therein for more thorough discussions. Most Monte Carlo methods in use are based on variations of the implicit Monte Carlo (IMC) approach introduced by Fleck and Cummings [75]. One of the strongest advantages of Monte Carlo methods is the ability to handle continuous treatments of the problem. This permits a greater degree of fidelity for treating complicated geometries and energy-dependence of cross section values than is possible with deterministic approaches. However, because Monte Carlo methods are based on a stochastic approach, solutions contain some degree of noise and convergence of the methods is generally slow (i.e., order $N^{-1/2}$ for N particle histories).

The discrete ordinates method is a collocation approximation for the angular variable of the trans-

port equation [135, Chapters 3 & 4]. This yields a coupled system of differential equations, each of which depends only on time and space. In a physical sense, discrete ordinates methods restrict particle trajectories to a finite number of directions corresponding to the nodes of the quadrature that is used. The use of a finite number of directions naturally leads to what is called the *ray-effects phenomenon* [126, 127] wherein particle trajectories can separate, producing visible lines or rays, for example, emanating from sharp material interfaces. The popularity of discrete ordinates methods is driven by a number of factors, including the robustness of the methods, ease of enforcing boundary conditions, and the development of extremely effective iterative techniques for solving the systems that result from implicit time discretizations [4, 123]. One of the primary disadvantages of discrete ordinates methods, which the hybrid methods in this work seek to improve upon, is that the cost of computing accurate solutions for problems containing disparate collisionality regimes can be exceptionally high.

Other approaches such as moment methods are generally less common, though recent efforts have improved upon many of the difficulties inherent to these approaches. Moment methods approximate the angular variation in the solution of the transport equation using a spectral approach [138, 148, 160]. This is done by: (i) writing the solution ψ of (1.2) in series form relative to some basis of $L^2(\mathbb{S}^2)$ (the most common choice being spherical harmonics), (ii) using this representation in the original transport equation to produce an infinite, coupled system of equations, and (iii) applying a closure relation to reduce to the resulting system to a finite-dimensional approximate form. The properties of moment approximations are determined primarily by the characteristics of the closure that is applied. The simplest closure is a simple Galerkin projection in $L^2(\mathbb{S}^2)$, which yields the standard P_N method. Other closures and the addition of filtering strategies have been investigated in an attempt to improve the accuracy and robustness of moment-based approaches - see [95, 96, 119, 137, 163, 190] and references therein for some recent examples. Additionally, a good overview and comparison of different closures can be found in [85]. The most prominent disadvantage of moment methods is that ensuring positivity in the solution is generally exceptionally difficult, particularly for problems that lack a significant degree of regularity. While recent work has significantly improved the robustness of moment-based approaches, other aspects impacting the practical application of such methods, such as the development of efficient solver and preconditioning techniques, still lags behind other methods.

Finite element methods have garnered some increased attention in recent years – see [20, 37, 99, 109, 117, 185] for some examples. One of the greatest strengths of finite element methods are their flexibility

in handling approaches such as adaptive refinement [117] and sparse grids [185] in a reasonably straightforward manner. However, the development of effective solver techniques for the systems of coefficients
resulting from implicit time discretization remains is an issue. Since the number of angular flux vectors
that can be stored is often significantly limited by available memory, advanced solver techniques such
as Krylov subspace methods can be prohibitively expensive when a high level of angular fidelity is required. The flexibility afforded by finite element methods may nevertheless permit better scalability on
very large distributed-memory systems, an area where methods such as the discrete ordinates method
have faced significant challenges [5, 21].

1.2.4 Hybrid methods

All of the numerical methods mentioned in section 1.2.3 yield approximations that require some tradeoffs in accuracy or efficiency. For example, discrete ordinates methods are generally quite robust for
a diverse array of problems, but they produce ray effects in streaming-dominated regimes [126, 127]
and become ill-conditioned for problems that are very optically thick [4]. Monte Carlo methods do not
produce ray effects like discrete ordinates methods, though they do become very computationally expensive for optically thick problems [86]. In contrast, most moment-based approaches are better-suited
to highly-collisional problems, but often produce negative values and poorly conditioned systems in
streaming dominated regimes [138].

For problems that contain collision probabilities that are relatively consistent everywhere, one can select a numerical approach that is tailored towards the regime of interest. However, for problems that contain disparate collisionality regimes, no one method consistently out-performs the others. As a result, a wide variety of hybrid methods have been developed that seek to combine different numerical approaches in ways that allow a diverse array of problems to be solved more accurately and efficiently than would otherwise be possible with any one method alone. Because the effectiveness of most numerical methods depends on the degree of collisionality present in a given problem, variations of the multiple collision decomposition (1.12) and (1.13) are often used as the foundation on which hybrid methods are constructed. Loosely speaking, this is done by applying a higher-fidelity approximation to terms near the beginning of the series and a lower-fidelity, more fluid-like approximation to terms in the tail of the series.

A crucial component of this approach is determining the point at which one should switch from a higher-fidelity approximation to a lower-fidelity one. The most common approach is the first-collision source technique [8, 9], in which only the first term of the series – the uncollided flux – is treated using a higher-fidelity approximation and the remaining terms are grouped into a single component – the collided flux – that is treated using a lower-fidelity approximation. Other approaches in which the solution is split after two [88] or an arbitrarily-specified number of collision events [65, 158, 159] have also been examined. One could also go one step further and consider decomposing the angular flux into more than two components to which progressively coarser approximations are applied [181].

The other aspect that must be considered is the determination of which methods should be applied to each component. In many cases methods such as Monte Carlo [8] or high-resolution discrete ordinates approximations (or variations thereof) [11, 146, 147, 188] are selected for the higher-resolution approximation. Analytic representations may also be used for the uncollided flux when available [9]. In many cases the reduced approximation applied to the collisional component is numerical in nature: For example, a low-resolution discrete ordinates method [8, 9, 88, 188] or a moment method [146, 147] may be used. Other times a reduced physics model such as a point-kinetics [158, 159] or diffusion approximation [23, 24, 136, 142] is used in place of the kinetic model for the collisional component. Within the Monte Carlo community, hybrid methods leveraging a stochastic diffusion approximation such as Implicit Monte Carlo Diffusion [86] or Discrete Diffusion Monte Carlo (DDMC) [65, 66] after a specified number of collision events have been used with satisfactory results [1, 187].

It should be noted that methods motivated by the same considerations as the first-collision source technique have been used in isolated instances under different names. Within the phonon transport community, a similar technique is used that is often referred to as the *ballistic-diffusive* model [11, 39, 40, 102, 103, 146, 147, 188], where the *ballistic* and *diffusive* parts correspond to the uncollided and collided components, respectively. A similar technique called the Isotropic Diffusion Source Approximation (IDSA) is used within the supernova neturino transport community [23, 24, 136, 142, 153]. The IDSA splits the flux distribution into *trapped* and *streaming* components to which time-dependent diffusion and steady-state ray-tracing approximations are applied, respectively. It is interesting to note that the hybrid methods introduced in [97] which form the basis of this work share a great deal of similarity with the IDSA. The key ingredient in both methods is the introduction of a mechanism by which the collided distribution can flow into the uncollided distribution. Traditional first-collision source methods lack this

capability: the uncollided distribution flows into the collided distribution but not vice versa. In the IDSA, this mechanism is derived on the continuum level using physical arguments to introduce source terms into the split equations. In contrast, the feedback mechanism of the hybrid methods of [97] and this work (i.e., the relabeling or reconstruction procedure) is constructed on the discrete level and is coupled with the time integration scheme. Further, while this work focuses on hybrid methods using discrete ordinates approximations for both the collided and uncollided components, this choice need not be strictly enforced, since the approach is compatible with other types of discretizations.

1.3 Defect correction

1.3.1 Introduction

The terms *defect correction* and *deferred correction* are used to refer to a broad class of methods in which the accuracy of an approximate solution to a problem is improved, often in an iterative fashion. One of the simplest examples of such a method is that of *iterative refinement* for the solution of a system of linear equations [100, Chapter 12]. The iterative refinement algorithm can be constructed as follows. Suppose that a linear system

$$Ax = b \tag{1.14}$$

is to be solved, where A and b are given and x is the unknown solution. Suppose further that a matrix B is known such that $B \approx A^{-1}$ in some reasonable sense. In most applications of iterative refinement, B corresponds to the matrix obtained from a linear solution algorithm (e.g., Gaussian Elimination) executed in finite-precision arithmetic. Since B approximates A^{-1} , it can be used to define an initial approximate solution to (1.14):

$$\widetilde{x}^{(0)} = Bb. \tag{1.15}$$

The difference between this approximate solution and the exact solution x will depend on how well B approximates A^{-1} .

Now suppose that some approximate solution $\tilde{x}^{(k)}$ has been computed after k applications of the iterative process to be defined, with the initial value $\tilde{x}^{(0)}$ given by (1.15). Then the *error* $e^{(k)}$ and the

 $\operatorname{residual} r^{(k)}$ in $\widetilde{x}^{(k)}$ are defined by

$$e^{(k)} = x - \tilde{x}^{(k)},\tag{1.16a}$$

$$r^{(k)} = b - A\tilde{x}^{(k)}. ag{1.16b}$$

If the error $e^{(k)}$ were known, then the exact solution x could be obtained through

$$x = \widetilde{x}^{(k)} + e^{(k)}. ag{1.17}$$

However, the problem of determining $e^{(k)}$ is just as difficult as computing x to begin with: The error $e^{(k)}$ satisfies the equation

$$Ae^{(k)} = Ax - A\tilde{x}^{(k)} = b - A\tilde{x}^{(k)} = r^{(k)},$$
 (1.18)

which has the same form as the original problem (1.14), but with a different right-hand side. In an attempt to improve the approximation solution, an approximation $\tilde{e}^{(k)}$ of $e^{(k)}$ can be computed using B as before:

$$\widetilde{e}^{(k)} = Br^{(k)}. (1.19)$$

This approximate error can then be used to define an "improved" solution:

$$\widetilde{x}^{(k+1)} = \widetilde{x}^{(k)} + \widetilde{e}^{(k)}. \tag{1.20}$$

So long as B is a sufficiently good approximation of A^{-1} , this process will increase the accuracy of the approximate solution at each iteration and the approximate solutions $\tilde{x}^{(k)}$ will converge to the exact solution x.

The iterative refinement algorithm contains most of the crucial components of many defect correction methods: a residual is computed for a given approximate solution, which is then used to compute an error approximation that is used to update the solution in an iterative fashion. In applications of iterative refinement, the difference between A^{-1} and B is generally due to the stability of the linear solver algorithm and the accumulated effects of round-off error [106, 169], but sometimes also includes mixed-precision computations [34]. In a broader context considering other types of defect correction methods, the difference between the "exact" and approximate solutions can vary widely depending on the algorithm – see the reviews [25, 26] for an assortment of defect correction methods.

Defect correction methods for differential and integral equations have an overall structure that is similar to that of the iterative refinement algorithm, but there are significant differences due to the continuum nature of these types of operator equations. Numerical approaches for solving such operator

equations typically follow a two-step procedure in which the problem is first discretized into a system of linear or nonlinear equations that are then solved using standard techniques. The method that is applied during the discretization step (e.g., a finite difference or finite element method) is typically accompanied by some form of approximation error. Methods with smaller approximation errors are generally much more complicated and thus much more computationally expensive to apply. In this context, one goal of defect correction methods is to reduce the computational cost associated with more accurate numerical techniques by combining them with less accurate – and less expensive – approximations.

For integral equations, two-grid methods combine high- and low-resolution quadratures to construct an iterative procedure for computing the high-resolution solution [17]. These methods are often nearly identical in form to the iterative refinement algorithm, but the iterative procedure specifically targets the numerical error of the low-resolution approximation. In relation to the iterative refinement algorithm, the matrices *A* and *B* for two-grid methods are the matrix corresponding to a high-resolution quadrature and the inverse of the matrix corresponding to a low-resolution quadrature, respectively. The intent is not that the low-resolution approximation will yield the inverse of the high-resolution approximation in exact arithmetic, but instead the two-grid method can be considered more as a form of preconditioning for the high-resolution system.

Defect correction methods for differential equations typically target increasing the order of the truncation error of the numerical scheme. These approaches are most commonly used in combination with finite difference methods for boundary value problems [35, 156] and initial value problems [18, 19, 70, 76, 77], though recent work has also extended this idea to discontinuous Galerkin methods applied to the spatial variables of models of fluid dynamics [72, 73]. These methods hold the numerical grid fixed and apply a combination of high- and low-order approximations (e.g., finite difference schemes) on that grid. As such, these methods can be viewed as preconditioned iterative procedures for computing the high-order numerical solution. However, a curious phenomenon occurs: When properly formulated, these methods increase the order of accuracy of the solution at each iteration up to the order of the high-order approximation. The result is that the defect correction iteration yields the same order of accuracy as the high-order approximation in a finite number of iterations. These methods are therefore commonly not iterated until convergence, but only up to the point when the maximum order of accuracy has been attained.

1.3.2 Integral deferred correction

Integral deferred correction (IDC) methods are a class of numerical methods for solving initial value problems in which high-order accuracy is attained through iterative application of a low-order scheme within a defect correction framework. IDC methods are characterized by the use of a discrete Picard integral representation of the error based on high-order numerical quadrature. The first IDC methods were the spectral deferred correction (SDC) methods of [70], in which Euler methods (both implicit and explicit) were used to compute provisional solutions and Gauss-Legendre quadrature used to approximate the Picard integral representation of the error. The study of IDC methods has since been expanded to include additional quadratures [130], various methods for computing provisional solutions and correction iterations [45, 47, 51], time parallelism [46, 48, 50], Krylov subspace iteration [104, 105], and multilevel formulations in the style of multigrid algorithms [173, 174], among others.

When it is convergent¹, the iterative procedure on which IDC methods are based converges to the collocation solution associated with the high-order quadrature used to approximate the Picard integral representation of the error [94]. Because of this, the accuracy and stability properties of IDC methods are similar to those of collocation methods, but the coupling between internal stages is significantly reduced. This reduction in coupling can be leveraged to develop IDC methods that incorporate additional features such as high-order operator splitting [44, 49] and implicit-explicit [144, 145] forms, while maintaining the accuracy of the underlying collocation approximation. However this comes at the expense of increased storage requirements. Typical IDC methods require storing the equivalent of two copies of each stage vector: twice that of the associated collocation scheme. This can render IDC methods impractical for problems involving high-dimensional partial differential equations.

¹The iterative procedure on which IDC methods are based can be divergent in certain situations, particularly in the strongly stiff limit when large numbers of non-uniformly spaced quadrature nodes are used [162].

CHAPTER 2

LOW-STORAGE INTEGRAL DEFERRED CORRECTION METHODS FOR SCIENTIFIC COMPUTING

2.1 Introduction

In this chapter a modification of the traditional IDC approach that significantly reduces the storage requirements of the methods is presented. These methods, which we call low-storage IDC or LS-IDC methods, require storing only one copy of each stage vector throughout the iteration process. This reduction is achieved by splitting the spectral integration matrix into upper-triangular and strictly lower-triangular parts, and applying the lower-triangular portion to the most recent set of stages in the residual calculation instead of the previous set. This splitting is similar in spirit to a Gauss-Seidel iteration in the sense that updated stages are used in the residual calculation as soon as they are computed. This allows previous stage vectors to be overwritten by their corresponding updates immediately after the substep in which the update is computed.

Much of the existing literature on low-storage integrators focuses on explicit methods, where the primary concern is the number of stage vectors that must be stored at any given time (see, e.g., [111, 112, 113] and references therein). This work focuses on fully-implicit methods, where additional considerations must be made for the storage requirements of solving the system(s) that define the stage values of the methods. Collocation integrators often contain fewer stages for a given order of accuracy than other methods, but because these stages are fully coupled and must be computed simultaneously, the storage requirements of the associated implicit solve scale proportionally to the number of stages. Other methods such as diagonally implicit Runge-Kutta methods require more stages for a given order of accuracy, but these stages are only weakly coupled and can be computed sequentially. The result is that the storage requirements for the implicit solves in these methods depends primarily on the size of a single stage vector and remains relatively constant as the number of stages of the method increases. The advantage of LS-IDC methods is derived from the fact that, relative to other fully-implicit schemes, LS-IDC methods minimize both the number of stages that must be stored simultaneously and the additional storage required for solving the systems that define the stage values.

This chapter is organized as follows. In section 2.2 a review of collocation and IDC integrators is

provided to establish notation, and LS-IDC methods are introduced. In section 2.3 it is shown that LS-IDC methods converge with the same order of accuracy as traditional IDC methods for linear, constant-coefficient problems, and the linear stability of some LS-IDC methods is examined numerically. In section 2.4 LS-IDC methods are applied to a nonlinear ODE and a linear transport equation to compare their accuracy and memory efficiency with other fully-implicit methods.

2.2 Method descriptions

We consider the numerical solution of an initial value problem (IVP) of the form

$$\begin{cases} y' = f(t, y) \\ y(t_0) = y_0 \end{cases}$$
 (2.1)

where $y \in \mathbb{R}^m$. We assume that the function $f : \mathbb{R}^+ \times \mathbb{R}^m \to \mathbb{R}^m$ and initial condition $y_0 \in \mathbb{R}^m$ are known.

2.2.1 Collocation

Integrating (2.1) with respect to t yields the Picard integral equation

$$y(t) = y_0 + \int_{t_0}^{t} f(\tau, y(\tau)) d\tau, \qquad (2.2)$$

the solution of which is equivalent to that of (2.1). Let Δt denote the desired timestep size: the equation (2.2) can be written as

$$y(t) = y_0 + \Delta t \int_0^{(s-t_0)/\Delta t} f(t_0 + s\Delta t, y(t_0 + s\Delta t)) ds$$
 (2.3)

through a change of variable. Let $\{c_n\}_{n=1}^N$ be a set of distinct points in the closed interval [0,1], and let $\{\theta_n(s)\}_{n=1}^N$ defined by

$$\theta_n(s) = \int_0^s \ell_n(\zeta) \, d\zeta \tag{2.4}$$

denote the quadrature weights associated with the points $\{c_n\}$ for integrating over the interval [0, s], where the functions $\ell_n(\zeta)$ are the Lagrange basis polynomials for the points $\{c_n\}$ satisfying

$$\ell_k(c_n) = \begin{cases} 1, & n = k, \\ 0, & n \neq k. \end{cases}$$

$$(2.5)$$

Replacing the integral in (2.3) by the above quadrature yields the collocation approximation

$$Y(t) = y_0 + \Delta t \sum_{n=1}^{N} \theta_n \left(\frac{t - t_0}{\Delta t} \right) f(t_n, Y_n), \tag{2.6}$$

with $Y_n = Y(t_n)$ an approximation of $y(t_n)$ and $t_n = t_0 + c_n \Delta t$ for n = 1, ..., N. The values Y_n are defined implicitly through the nonlinear system

$$Y_n = y_0 + \Delta t \sum_{k=1}^{N} \theta_{n,k} f(t_k, Y_k), \qquad (n = 1, ..., N),$$
(2.7)

where $\theta_{n,k} = \theta_k(c_n)$. The result of the timestep update is given by $Y(t_0 + \Delta t)$, which is obtained from (2.6) once the Y_n are known.

In subsequent sections it will be useful to write the system (2.7) using the vector notation

$$\mathbf{Y} = \mathbf{y}_0 + \Delta t(\Theta \otimes I)\mathbf{f},\tag{2.8}$$

where

$$\mathbf{Y} = \begin{bmatrix} Y_1, \dots, Y_N \end{bmatrix}^T, \qquad \mathbf{f} = \begin{bmatrix} f(Y_1), \dots, f(Y_N) \end{bmatrix}^T, \tag{2.9}$$

and $\Theta = (\theta_{n,k})_{n,k=1}^N$ is the matrix of quadrature weights.

2.2.2 Integral deferred correction

We consider only the case of integral deferred correction (IDC) methods constructed using implicit Euler as the low-order method. The following notations will be used in this and subsequent sections: Latin letters will be used for exact quantities and Greek letters will be used for approximate quantities. The correction level or index of the correction iteration corresponding to a solution will be denoted using a superscript enclosed in brackets: square brackets will be used for individual point values and round brackets will be used to denote polynomial interpolants of point values.

Given a timestep size Δt and a set of points $\{c_n\}_{n=1}^N \subset [0,1]$ the IDC method begins by computing a set of provisional solutions $\{\xi_n^{[0]}\}_{n=1}^N$ using the implicit Euler method:

$$\xi_n^{[0]} = \xi_{n-1}^{[0]} + h_n \Delta t f\left(t_0 + c_n \Delta t, \, \xi_n^{[0]}\right), \qquad (n = 1, ..., N),$$
(2.10)

where

$$h_n = c_n - c_{n-1}$$
 for $n = 1, ..., N$ (2.11)

denote the substep sizes. The superscript [p] is used to denote the correction level, with p = 0 the first set of provisional solutions computed by the implicit Euler method.

The IDC correction procedure can be constructed as follows. Let $\xi^{(p-1)}(t)$ denote the polynomial interpolant (in time) of a set of provisional solutions $\{\xi_n^{[p-1]}\}_{n=1}^N$ at correction level p-1. The solution error is given by

$$e^{(p-1)}(t) = v(t) - \xi^{(p-1)}(t) \tag{2.12}$$

and the residual is given by

$$r^{(p-1)}(t) = \frac{d}{dt}\xi^{(p-1)}(t) - f\left(t, \, \xi^{(p-1)}(t)\right). \tag{2.13}$$

The derivative of the error satisfies

$$\frac{d}{dt}e^{(p-1)}(t) = \frac{d}{dt}y(t) - \frac{d}{dt}\xi^{(p-1)}(t),$$
(2.14a)

$$= f(t, y) - f(t, \xi^{(p-1)}) - r^{(p-1)}(t), \tag{2.14b}$$

$$= f\left(t, \, \xi^{(p-1)} + e^{(p-1)}\right) - f\left(t, \, \xi^{(p-1)}\right) - \frac{d}{dt} \int_{t_0}^t r^{(p-1)}(s) \, ds. \tag{2.14c}$$

Applying implicit Euler approximations to the time derivatives in (2.14c) yields

$$\epsilon_{n}^{[p-1]} = \epsilon_{n-1}^{[p-1]} + h_{n} \Delta t f\left(t_{n}, \, \xi_{n}^{[p-1]} + \epsilon_{n}^{[p-1]}\right) - h_{n} \Delta t f\left(t_{n}, \, \xi_{n}^{[p-1]}\right) - \int_{t_{n-1}}^{t_{n}} r^{(p-1)}(s) \, ds, \tag{2.15}$$

from which approximations $\{\varepsilon_n^{[p-1]}\}_{n=1}^N$ of the error in the set of provisional solutions can be obtained using the ansatz that the error is initially zero; i.e., $e^{(p-1)}(t_0)=0$ for all $p\geq 1$. These approximations of the error can be used to update the provisional solutions and define a new set of provisional solutions:

$$\xi_n^{[p]} = \xi_n^{[p-1]} + \epsilon_n^{[p-1]}, \qquad (n = 1, ..., N).$$
 (2.16)

Observing that the integral of an interpolant is equivalent to quadrature and applying the relations (2.13) and (2.16) to (2.15) yields

$$\xi_{n}^{[p]} = \xi_{n-1}^{[p]} + h_{n} \Delta t f\left(t_{n}, \xi_{n}^{[p]}\right) - h_{n} \Delta t f\left(t_{n}, \xi_{n}^{[p-1]}\right) + \Delta t \sum_{k=1}^{N} \gamma_{n,k} f\left(t_{k}, \xi_{k}^{[p-1]}\right), \tag{2.17}$$

where

$$\gamma_{n,k} = \int_{c_{n-1}}^{c_n} \ell_k(s) \, ds \tag{2.18}$$

are the normalized quadrature weights associated with the nodes $\{c_k\}_{k=1}^N$ for integrating over the intervals $[c_{n-1}, c_n]$ for n = 1, ..., N.

In practice, the values $\xi_n^{[p]}$ are computed sequentially using (2.17). Nevertheless, (2.17) defines a system of N nonlinear equations that can be written using the vector notation

$$\boldsymbol{\xi}^{[p]} = \mathbf{y}_0 + \Delta t (H \otimes I) \mathbf{f}^{[p]} + \Delta t ((\Theta - H) \otimes I) \mathbf{f}^{[p-1]}, \tag{2.19}$$

where

$$\boldsymbol{\xi}^{[p]} = \left(\xi_1^{[p]}, \dots, \xi_N^{[p]}\right)^T, \qquad \mathbf{f}^{[p]} = \left(f(\xi_1^{[p]}), \dots, f(\xi_N^{[p]})\right)^T, \tag{2.20}$$

and

$$H = \begin{pmatrix} h_1 & & & \\ h_1 & h_2 & & & \\ \vdots & \vdots & \ddots & & \\ h_1 & h_2 & \cdots & h_N \end{pmatrix}. \tag{2.21}$$

This representation of the correction procedure will be useful for the analysis in subsequent sections.

Once sufficiently many correction iterations have been performed, say P correction iterations, an approximation of the solution of (2.1) at the time $t_0 + \Delta t$ is constructed from the provisional solutions $\{\xi_n^{[P]}\}_{n=1}^N$ from the last correction iteration. For the numerical tests in section 2.4 we consider only quadrature sets with $c_N = 1$, in which case it is sufficient to use $\xi_N^{[P]}$ as the resulting approximation for the timestep.

2.2.3 Low-storage integral deferred correction

An examination of the correction equation (2.17) reveals that each of the provisional solutions on correction level p is dependent on all of the provisional solutions on correction level p-1. Because of this, the traditional IDC iteration requires storing the equivalent of two sets of provisional solutions in memory at any given time; specifically $\{\xi_n^{[p-1]}\}_{n=1}^N$ and $\{\xi_n^{[p]}\}_{n=1}^N$. This storage requirement can be prohibitively

 $^{^1}$ In practice, one commonly stores the function evaluations $\{f(t_n,\xi_n^{[p-1]})\}_{n=1}^N$ instead of the stages $\{\xi_n^{[p-1]}\}_{n=1}^N$ in order to avoid repeating the function evaluations at each substep. Note that this changes only the computational requirements of the implementation: the storage requirements of these approaches are identical.

expensive when high-order IDC methods are applied to large, complex systems of partial differential equations.

The low-storage IDC (LS-IDC) methods considered here are a simple modification of traditional IDC methods in which the updated provisional solutions are used in the computation of the residual as they become available. Traditional IDC methods evaluate the residual using the same set of provisional solutions

$$\{\xi_k^{[p-1]}\}_{k=1}^N \tag{2.22}$$

for each substep n = 1,...,N. In contrast, LS-IDC methods compute the residual for the nth substep using

$$\{\xi_k^{[p]}\}_{k=1}^{n-1} \quad \text{and} \quad \{\xi_k^{[p-1]}\}_{k=n}^N,$$
 (2.23)

which yields a different stencil for each substep. With this idea, (2.17) can be modified as:

$$\xi_{n}^{[p]} = \xi_{n-1}^{[p]} + h_{n} \Delta t f\left(t_{n}, \xi_{n}^{[p]}\right) - h_{n} \Delta t f\left(t_{n}, \xi_{n}^{[p-1]}\right) \\
+ \Delta t \sum_{k=1}^{n-1} \gamma_{n,k} f\left(t_{k}, \xi_{k}^{[p]}\right) + \Delta t \sum_{k=n}^{N} \gamma_{n,k} f\left(t_{k}, \xi_{k}^{[p-1]}\right).$$
(2.24)

The advantage of this approach is that once $\xi_n^{[p]}$ has been computed $\xi_n^{[p-1]}$ is no longer needed: hence $\xi_n^{[p-1]}$ can be overwritten by $\xi_n^{[p]}$ immediately after the nth substep. The result is that LS-IDC methods only require storing the equivalent of one correction level at any given time: half the memory requirement of traditional IDC methods.

For the analysis in subsequent sections, we note that (2.24) defines a system of K nonlinear equations of the form

$$\boldsymbol{\xi}^{[p]} = \mathbf{y}_0 + \Delta t ((H + \Theta_{\mathcal{L}}) \otimes I) \mathbf{f}^{[p]} + \Delta t ((\Theta_{\mathcal{U}} - H) \otimes I) \mathbf{f}^{[p-1]}$$
(2.25)

where Θ_L is the strictly lower-triangular portion of Θ and Θ_U is the upper-triangular portion of Θ . Constructing an approximation of the solution of (2.1) at the time $t_0 + \Delta t$ from the set of provisional solutions $\{\xi_n^{[P]}\}_{n=1}^N$ from the last correction iteration proceeds in the same manner as traditional IDC methods. For the numerical tests in section 2.4 we consider only quadrature sets with $c_N = 1$, in which case it is sufficient to use $\xi_N^{[P]}$ as the resulting approximation for the timestep.

2.3 Convergence and stability

We now examine some properties of the LS-IDC methods, focusing on the case of linear, constant coefficient systems.

2.3.1 Convergence in the limit of small timestep size

For sufficiently small Δt , it is known that IDC methods converge to the collocation solution associated with the points $\{c_n\}_{n=1}^N$ as the number of correction iterations $p \to \infty$ [94, 104, 105]. We now provide a proof that this holds also for LS-IDC methods in the case of linear, constant-coefficient ODEs. The following proposition will be used:

Proposition 2. Let A and B be arbitrary square matrices, and let $\tilde{A} = I - \Delta t A$ with $\Delta t > 0$. For Δt sufficiently small \tilde{A}^{-1} exists, and

$$\lim_{K \to \infty} \tilde{A}^{-1} \sum_{k=0}^{K} \left(\Delta t B \tilde{A}^{-1} \right)^k \tag{2.26}$$

exists and is equal to

$$\left(\tilde{A} - \Delta t B\right)^{-1}.\tag{2.27}$$

Proof. For $\Delta t < 1/\rho(A)$ (where $\rho(A)$ denotes the spectral radius of the matrix A) we have

$$\rho\left(\Delta t A\right) = \Delta t \rho\left(A\right) < 1. \tag{2.28}$$

It follows that $I - \Delta t A$ is invertible and it's inverse is given by the Neumann series expansion

$$(I - \Delta t A)^{-1} = \sum_{k=0}^{\infty} (\Delta t A)^k$$
 (2.29)

(c.f., [140, Section 7.10]). Given an operator norm $\|\cdot\|$, for

$$\Delta t < \frac{1}{\|A\| + \|B\|} \tag{2.30}$$

we have, using the absolute convergence of the Neumann series,

$$\rho(\Delta t B \tilde{A}^{-1}) \le \Delta t \|B\| \|\tilde{A}^{-1}\|,\tag{2.31a}$$

$$= \Delta t \|B\| \lim_{K \to \infty} \left\| \sum_{k=0}^{K} (\Delta t A)^k \right\|, \tag{2.31b}$$

$$\leq \Delta t \|B\| \lim_{K \to \infty} \sum_{k=0}^{K} \Delta t^k \|A\|^k, \tag{2.31c}$$

$$= \frac{\Delta t \|B\|}{1 - \Delta t \|A\|},\tag{2.31d}$$

Hence the limit in (2.26) exists and

$$\tilde{A}^{-1} \lim_{K \to \infty} \sum_{k=0}^{K} \left(\Delta t B \tilde{A}^{-1} \right)^{k} = \tilde{A}^{-1} \left(I - \Delta t B \tilde{A}^{-1} \right)^{-1} = \left(\tilde{A} - \Delta t B \right)^{-1}. \tag{2.32}$$

In the case of linear, constant-coefficient problems, we have f(t, y) = Fy for some matrix F. Then the system (2.8) for the stages of the collocation method associated with the points $\{c_n\}_{n=1}^N$ can be written as

$$\mathbf{Y} = \mathbf{y}_0 + \Delta t(\Theta \otimes F)\mathbf{Y}$$
 so that $\mathbf{Y} = [\mathbf{I} - \Delta t(\Theta \otimes F)]^{-1}\mathbf{y}_0.$ (2.33)

Proposition 3. Let f(t, y) = Fy in (2.1). Then for Δt sufficiently small the iterates $\boldsymbol{\xi}^{[p]}$ of the IDC method (2.19) satisfy

(i)
$$\lim_{n\to\infty} \xi^{[p]} = \mathbf{Y}$$
, and

(ii)
$$\|\mathbf{Y} - \boldsymbol{\xi}^{[p]}\| = O\left(\Delta t^{p+1}\right)$$
.

Proof. For f(t, y) = Fy the system (2.19) takes the form

$$\boldsymbol{\xi}^{[p]} = \mathbf{y}_0 + \Delta t (H \otimes F) \boldsymbol{\xi}^{[p]} + \Delta t ((\Theta - H) \otimes F) \boldsymbol{\xi}^{[p-1]}. \tag{2.34}$$

Letting $\mathbf{A} = H \otimes F$ and $\mathbf{B} = (\Theta - H) \otimes F$, and noting that

$$\boldsymbol{\xi}^{[0]} = \left[\mathbf{I} - \Delta t(H \otimes F)\right]^{-1} \mathbf{y}_0, \tag{2.35}$$

equation (2.34) implies that

$$\boldsymbol{\xi}^{[p]} = \tilde{\mathbf{A}}^{-1} \sum_{k=0}^{p} \left(\Delta t \mathbf{B} \tilde{\mathbf{A}}^{-1} \right)^{k}, \tag{2.36}$$

where $\tilde{\bf A} = {\bf I} - \Delta t {\bf A}$. Applying proposition 2 to (2.36) under the assumption that Δt is sufficiently small, and applying some algebraic manipulation to the result yields

$$\lim_{p \to \infty} \mathbf{\xi}^{[p]} = (\tilde{\mathbf{A}} - \Delta t \mathbf{B})^{-1} \mathbf{y}_0, \tag{2.37a}$$

$$= \left[\mathbf{I} - \Delta t (H \otimes F) - \Delta t ((\Theta - H) \otimes F) \right]^{-1} \mathbf{y}_0, \tag{2.37b}$$

$$= \left[\mathbf{I} - \Delta t(\Theta \otimes F)\right]^{-1} \mathbf{y}_0 = \mathbf{Y}. \tag{2.37c}$$

The first result implies the second upon observing that

$$\mathbf{Y} - \boldsymbol{\xi}^{[p]} = \tilde{\mathbf{A}}^{-1} \sum_{k=p+1}^{\infty} \left(\Delta t \mathbf{B} \tilde{\mathbf{A}}^{-1} \right)^k \mathbf{y}_0$$
 (2.38)

and applying relevant norms with the absolute convergence of the Neumann series.

Proposition 4. Let f(t, y) = Fy in (2.1). Then for Δt sufficiently small the iterates $\boldsymbol{\xi}^{[p]}$ of the LS-IDC method (2.25) satisfy

(i)
$$\lim_{p\to\infty} \xi^{[p]} = \mathbf{Y}$$
, and

$$(ii) \ \|\mathbf{Y} - \boldsymbol{\xi}^{[p]}\| = O\left(\Delta t^{p+1}\right).$$

Proof. The proof is similar to that of proposition 3. For f(t, y) = Fy the system (2.25) takes the form

$$\boldsymbol{\xi}^{[p]} = \mathbf{y}_0 + \Delta t ((H + \Theta_{\mathbf{L}}) \otimes F) \boldsymbol{\xi}^{[p]} + \Delta t ((\Theta_{\mathbf{U}} - H) \otimes F) \boldsymbol{\xi}^{[p-1]}. \tag{2.39}$$

Letting $\mathbf{A} = (H + \Theta_{L}) \otimes F$ and $\mathbf{B} = (\Theta_{U} - H) \otimes F$, equation (2.39) implies that

$$\boldsymbol{\xi}^{[p]} = \tilde{\mathbf{A}}^{-1} \sum_{k=0}^{p-1} \left(\Delta t \mathbf{B} \tilde{\mathbf{A}}^{-1} \right)^k \mathbf{y}_0 + \left(\Delta t \tilde{\mathbf{A}}^{-1} \mathbf{B} \right)^p \boldsymbol{\xi}^{[0]}, \tag{2.40}$$

where $\tilde{\bf A} = {\bf I} - \Delta t {\bf A}$. Assuming that Δt is sufficiently small, using proposition 2 and the fact that

$$\lim_{p \to \infty} \left(\Delta t \tilde{\mathbf{A}}^{-1} \mathbf{B} \right)^p \boldsymbol{\xi}^{[0]} = 0 \tag{2.41}$$

in (2.40) implies that

$$\lim_{p \to \infty} \boldsymbol{\xi}^{[p]} = (\tilde{\mathbf{A}} - \Delta t \mathbf{B})^{-1} \mathbf{y}_0, \tag{2.42a}$$

$$= \left[\mathbf{I} - \Delta t ((H + \Theta_{\mathcal{L}}) \otimes F) - \Delta t ((\Theta_{\mathcal{U}} - H) \otimes F) \right]^{-1} \mathbf{y}_{0}, \tag{2.42b}$$

$$= \left[\mathbf{I} - \Delta t(\Theta \otimes F) \right]^{-1} \mathbf{y}_0 = \mathbf{Y}. \tag{2.42c}$$

Now,

$$\mathbf{Y} - \boldsymbol{\xi}^{[p]} = \left(\Delta t \tilde{\mathbf{A}}^{-1} \mathbf{B}\right)^p \left(\tilde{\mathbf{A}}^{-1} \mathbf{y}_0 - \boldsymbol{\xi}^{[0]}\right) + \tilde{\mathbf{A}}^{-1} \sum_{k=p+1}^{\infty} \left(\Delta t \mathbf{B} \tilde{\mathbf{A}}^{-1}\right)^k \mathbf{y}_0, \tag{2.43a}$$

$$= \left(\Delta t \tilde{\mathbf{A}}^{-1} \mathbf{B}\right)^{p} \left[\left(\tilde{\mathbf{A}}^{-1} \mathbf{y}_{0} - \mathbf{Y} \right) + \left(\mathbf{Y} - \boldsymbol{\xi}^{[0]} \right) \right] + \tilde{\mathbf{A}}^{-1} \sum_{k=p+1}^{\infty} \left(\Delta t \mathbf{B} \tilde{\mathbf{A}}^{-1} \right)^{k} \mathbf{y}_{0}. \tag{2.43b}$$

Since

$$\tilde{\mathbf{A}}^{-1}\mathbf{y}_0 - \mathbf{Y} = \tilde{\mathbf{A}}^{-1} \sum_{k=1}^{\infty} \left(\Delta t \mathbf{B} \tilde{\mathbf{A}}^{-1} \right)^{-1} \mathbf{y}_0, \tag{2.44}$$

it follows that

$$\|\tilde{\mathbf{A}}^{-1}\mathbf{y}_0 - \mathbf{Y}\| = O(\Delta t). \tag{2.45}$$

Additionally, since $\boldsymbol{\xi}^{[0]}$ is computed using the implicit Euler method, we have

$$\|\mathbf{Y} - \boldsymbol{\xi}^{[0]}\| = O(\Delta t^2). \tag{2.46}$$

Applying relevant norms to (2.43b), and using (2.45), (2.46), and the absolute convergence of the Neumann series yields the second result.

2.3.2 Convergence for arbitrary timestep sizes

Consider the standard linear test problem

$$\begin{cases} y' = \lambda y, & \lambda \in \mathbb{C}, \\ y(t_0) = 1. \end{cases}$$
 (2.47)

When an IDC method is applied to (2.47) each correction iteration can be written in the form of a matrix-vector multiplication in which the *iteration matrix* of the method G(z) is applied to the vector of previous stage values, where $z = \lambda \Delta t$ and Δt is the timestep size. The *contraction factor* of the method is defined as the spectral radius of the iteration matrix $\rho(G(z))$. For traditional IDC methods the iteration matrix takes the form

$$G(z) = z(I - zH)^{-1}(\Theta - H),$$
 (2.48)

and the iteration matrix for LS-IDC methods takes the form

$$G(z) = z (I - z (H + \Theta_{L}))^{-1} (\Theta_{U} - H).$$
 (2.49)

The analysis discussed in the preceding section establishes convergence for IDC and LS-IDC methods when the timestep size used is sufficiently small. The arguments rely on the fact that the operator applied at each correction iteration is a contraction when the timestep size is sufficiently small. However, what constitutes "small enough" is determined by the spectrum of the iteration matrix: this spectrum is governed by the operator F of the system to be solved, and, via the matrices H and Θ , the quadrature nodes $\{c_n\}_{n=1}^N$. Upon considering the problem (2.47), it becomes clear that the results of section 2.3.1 can be portrayed as describing the behavior of the methods in the limit $|z| \to 0$ instead of $\Delta t \to 0$. While these two characterizations are mathematically equivalent, depicting this behavior in terms of $|z| \to 0$ more explicitly acknowledges the dependence on the properties of the actual system being solved.

Many problems of practical interest possess eigenvalues of sufficiently large magnitude that the behavior of methods when |z| is small is of little useful importance due to the exceptionally small step sizes required to capture this limiting behavior. Instead, it is more desirable that a method is well-behaved for |z| large. For defect correction methods, this translates into the requirement that that the contraction factor of the iteration matrix remain less than one, and ideally much less than one, even for |z| large.

The contraction factor associated with a given IDC method is strongly dependent on both the number and type of quadrature nodes used. For IDC methods constructed with equispaced nodes, the contraction factor has the desirable property that $\rho(G(z)) \to 0$ as $z \to -\infty$ with $z \in \mathbb{R}$ [76, 184]. However, schemes using equispaced nodes are not as well behaved near the imaginary axis as other schemes (they are not A-stable) and can be negatively impacted by the Runge phenomenon when large numbers of quadrature nodes are used. The situation is quite different for Gauss-Radau, Gauss-Lobatto, and Gauss-Legendre nodes, the use of which is desirable due to their spectral accuracy and superior stability properties. IDC methods constructed using these types of quadrature nodes have $\lim_{z\to -\infty} \rho(G(z)) > 0$, and, when large numbers of quadrature nodes are used, $\rho(G(z)) > 1$ away from the origin [162]. Because of this, order reduction can be observed when these methods are applied to very stiff problems.

The contraction factors of IDC and LS-IDC methods along the negative real axis are shown in figure 2.1 and the contraction factors in the strongly stiff limit are shown in table 2.1. For most quadrature nodes, the low-storage modification negatively impacts the contraction factors of the methods. This effect is most pronounced for *right-biased* quadratures: that is, quadratures that include the right endpoint of the timestep interval and not the left. The contraction factors of right-biased quadratures increase rapidly in the stiff limit as the number of nodes is increased. This effect is most pronounced

Nodes	Gauss	s-Radau	Chebys	shev-Radau	Equal Right			
	IDC	LS-IDC	IDC	LS-IDC	IDC	LS-IDC		
2	0.250	0.546	0.333	0.488	0.000	1.618		
3	0.434	0.998	0.463	0.607	0.000	4.235		
4	0.618	1.498	0.642	1.027	0.000	3.111		
5	0.737	2.078	0.760	1.612	0.001	6.9E+1		
6	0.816	2.761	0.839	2.279	0.003	5.4E+2		
7	0.873	3.568	0.895	3.113	0.008	5.4E+4		
8	0.915	4.523	0.936	4.060	0.013	3.3E+6		

(a) Contraction factors for quadrature sets including the right endpoint of the timestep interval and not the left.

Nodes	Gauss-Legendre		Gauss	-Lobatto	Che	byshev	Equal Both		
	IDC	LS-IDC	IDC	LS-IDC	IDC	LS-IDC	IDC	LS-IDC	
2	0.317	0.493	0.500	0.500	0.396	0.479	0.500	0.500	
3	0.421	0.642	0.500	0.556	0.460	0.554	0.500	0.556	
4	0.561	0.954	0.592	0.581	0.582	0.566	0.553	0.546	
5	0.665	1.305	0.684	0.591	0.679	0.595	0.616	0.856	
6	0.742	1.710	0.758	0.596	0.753	0.866	0.682	1.812	
7	0.800	2.182	0.815	0.857	0.810	1.211	0.748	3.079	
8	0.845	2.735	0.860	1.280	0.854	1.629	0.813	4.169	

(b) Contraction factors for quadrature sets containing both endpoints or neither endpoint of the timestep interval.

Table 2.1: Contraction factors for IDC and LS-IDC methods in the stiff limit using different numbers and types of quadrature nodes with implicit Euler substeps.

for equally-spaced right-biased quadratures, and, in stark contrast to the behavior of traditional IDC methods, all LS-IDC methods constructed with equally-spaced right-biased quadratures yield contraction factors greater than one in the stiff limit.

LS-IDC methods constructed with Gauss-Lobatto quadratures display notably different behavior than methods using other quadratures: When four, five, or six quadrature nodes are used, LS-IDC methods with Gauss-Lobatto quadratures produce contraction factors in the stiff limit that are smaller than those of traditional IDC methods using the same quadratures. Additionally, the contraction factors for IDC and LS-IDC methods constructed using two or three Gauss-Lobatto nodes are comparable. These results suggest that LS-IDC methods can be a competitive replacement for IDC methods when moderately high-order accuracy is desired, but may exacerbate the difficulties faced by IDC methods for very stiff problems.

2.3.3 Linear stability

We now examine the linear stability properties of LS-IDC methods. We begin by briefly reviewing some basic terminology related to linear stability theory of numerical methods for the solution of initial-value problems for ODEs before returning to the discussion of LS-IDC methods: see, e.g., [15, 92] for more detailed discussions of linear stability theory.

A timestep update with step size Δt of any one-step method applied to the test problem (2.47) can be written as $y_{n+1} = R(\lambda \Delta t)y_n$. The function R(z) is called the *stability function* of the method and the quantity |R(z)| the *amplification factor*. The *region of absolute stability* of the method is the set $\{z \in \mathbb{C} : |R(z)| \le 1\}$. A method is called *A-stable* if its region of absolute stability contains the entire left half-plane, and *L-stable* if it is *A-*stable and $|R(z)| \to 0$ as $z \to \infty$.

The stability of standard IDC methods using implicit Euler as the base method was examined numerically in [61]: it was found that it is possible to obtain IDC methods that are L-stable to within the limits of double-precision arithmetic. The construction of such methods is based on three observations: (i) any A-stable IDC method with $c_1 \neq 0$ is necessarily L-stable [130, Theorem 3.1]; (ii) Radau-IIA collocation schemes are L-stable (and hence also A-stable) at all orders of accuracy [93]; and (iii) the provisional solutions of IDC methods converge to the stage values of the collocation method corresponding to the points $\{c_n\}_{n=1}^N$ used by the IDC method as the number of correction iterations increases. L-stable IDC methods were constructed in [61] by choosing the points $\{c_n\}_{n=1}^N$ to correspond to a set of (right-biased) Gauss-Radau quadrature nodes and applying sufficiently many correction iterations until the desired stability properties were obtained. Note that these additional correction iterations do not increase the overall accuracy of the method², and serve only to refine the method's stability properties.

The boundaries of the regions of absolute stability of IDC and LS-IDC methods constructed using two and three Gauss-Radau quadrature nodes are shown in figure 2.2. The boundaries of the regions of absolute stability of IDC and LS-IDC methods constructed using four and five Gauss-Lobatto quadrature nodes are shown in figures 2.3 and 2.4, respectively. For Gauss-Radau quadratures, it is clear that the

 $^{^2}$ Additional correction iterations do not increase the overall accuracy of the IDC method once the maximum order determined by the collocation scheme is attained. For sufficiently small Δt and well-behaved problems, this maximum order of accuracy can be attained using, e.g., 2N-2 correction iterations for N Radau quadrature nodes. Additional correction iterations may help reduce the effects of order reduction observed when IDC methods are applied to very stiff problems or very large timesteps are used [104, 130, 144].

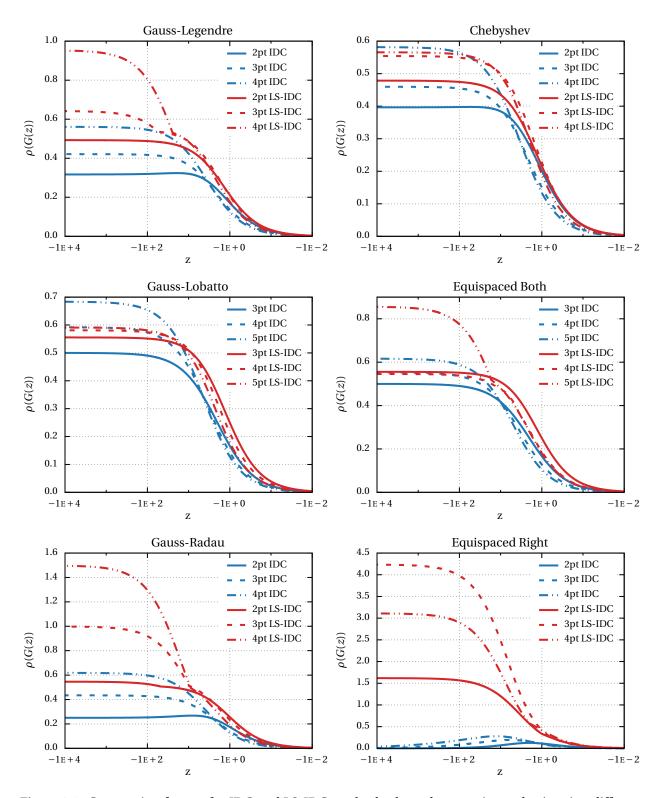


Figure 2.1: Contraction factors for IDC and LS-IDC methods along the negative real axis using different numbers and types of quadrature nodes with implicit Euler substeps.

low-storage modification negatively impacts the stability properties of the methods, and that recovering L-stability with LS-IDC methods becomes much more difficult as the order of the method increases. When two Gauss-Radau quadrature nodes are used the stability results for the LS-IDC method are similar to that of the traditional IDC method. However, the stability of the LS-IDC method using three Gauss-Radau quadrature nodes is considerably worse than the corresponding IDC method. In this case the LS-IDC method is unable to recover L-stability in the same number of correction iterations as the IDC method. For Gauss-Lobatto quadratures, the stability regions of IDC and LS-IDC methods behave similarly near the origin, but the small oscillations that appear in the boundary of the stability regions of LS-IDC methods protrude further away from the imaginary axis than the corresponding IDC methods. From a broader viewpoint, the stability of IDC and LS-IDC methods using Gauss-Lobatto quadratures is rather different. The stability regions of IDC methods using Gauss-Lobatto quadratures are unbounded when the number of correction iterations applied is small, but appear to be bounded when larger numbers of correction iterations are used. The bounded stability regions are quite large in all cases, and seem to increase in size once additional correction iterations are performed. In contrast, the stability regions of LS-IDC methods using Gauss-Lobatto quadratures are unbounded for all numbers of correction iterations considered. These results suggest that IDC and LS-IDC methods using Gauss-Lobatto quadratures are best suited to different types of problems. For problems where stability for values near the imaginary axis is more important, traditional IDC methods are likely a better choice from the perspective of stability. For problems where stability for values away from the imaginary axis is more important, LS-IDC methods seem to more consistently yield better stability when Gauss-Lobatto quadratures are used.

2.4 Numerical tests

We now consider the application of LS-IDC methods to two test problems: a nonlinear ODE for which an exact solution is known is used to study the accuracy of the methods and a linear transport equation solved by the discrete ordinates method is used to illustrate the memory efficiency of the methods.

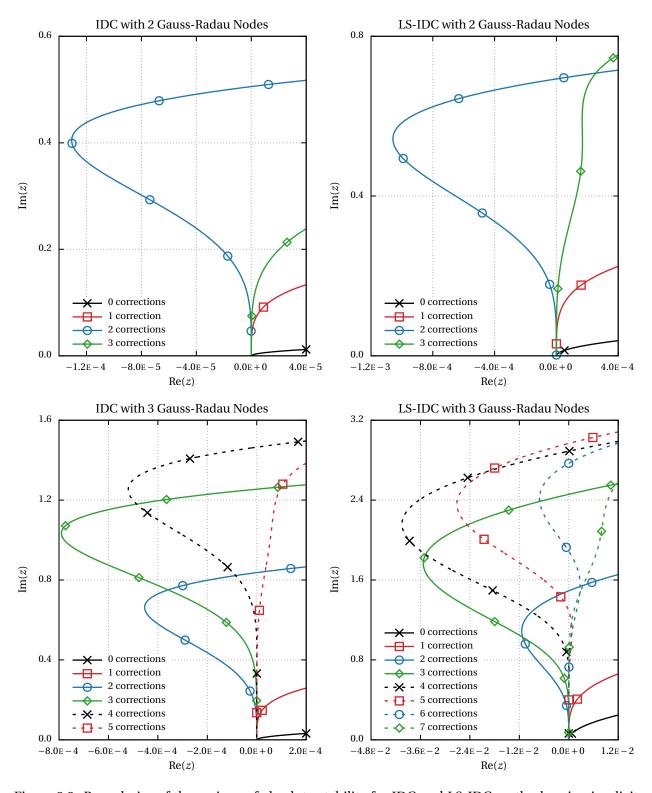


Figure 2.2: Boundaries of the regions of absolute stability for IDC and LS-IDC methods using implicit Euler substeps on (right biased) Gauss-Radau quadrature nodes with various numbers of correction iterations. Regions of absolute stability lie to the *left* of the origin.

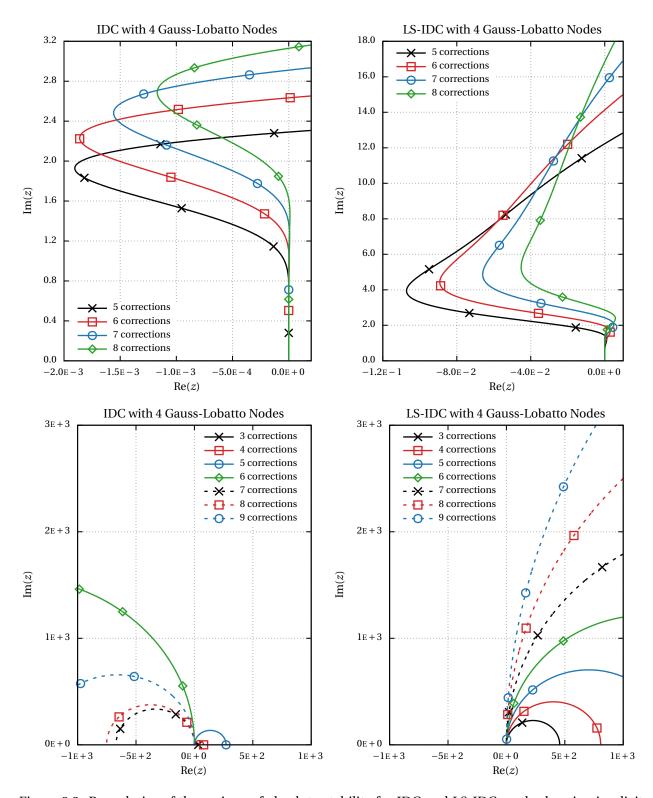


Figure 2.3: Boundaries of the regions of absolute stability for IDC and LS-IDC methods using implicit Euler substeps on four Gauss-Lobatto quadrature nodes with various numbers of correction iterations. Regions of absolute stability lie to the *left* of the origin.

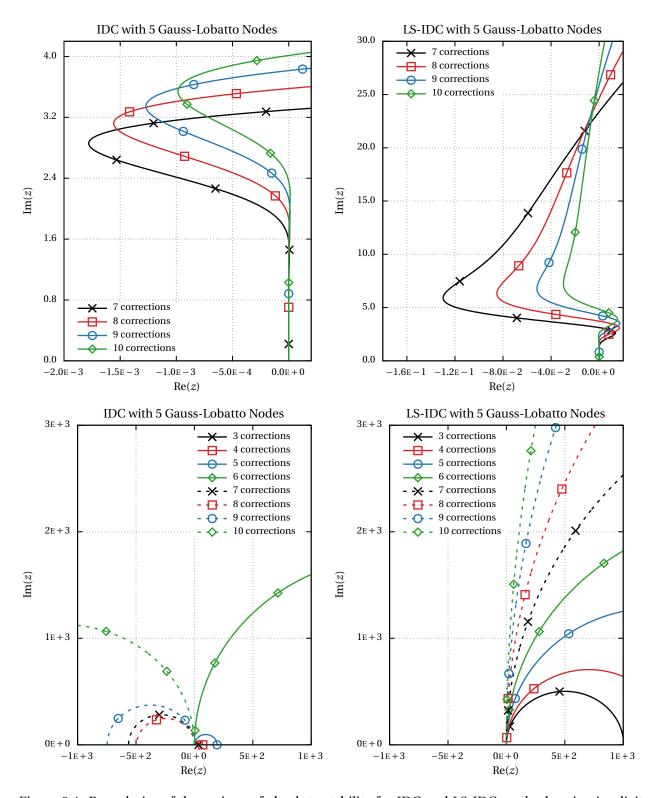


Figure 2.4: Boundaries of the regions of absolute stability for IDC and LS-IDC methods using implicit Euler substeps on five Gauss-Lobatto quadrature nodes with various numbers of correction iterations. Regions of absolute stability lie to the *left* of the origin.

2.4.1 A non-stiff, non-linear ODE

The first test problem is a simple nonlinear ODE from [18]:

$$\begin{cases} y_1' = -y_2 + y_1 \left(1 - y_1^2 - y_2^2 \right), \\ y_2' = y_1 + 3y_2 \left(1 - y_1^2 - y_2^2 \right), \\ y(0) = (1, 0)^T. \end{cases}$$
 (2.50)

The solution of (2.50) is given by $y(t) = (\cos t, \sin t)^T$.

Convergence studies of IDC and LS-IDC methods applied to the IVP (2.50) are shown in figure 2.5. Results for the collocation method corresponding to the quadrature used are included in each case for completeness. Three quadratures are considered for each method: three Gauss-Radau nodes, four Gauss-Lobatto nodes, and five right-biased equispaced nodes $\{c_n = n/5\}_{n=1}^5$. The methods are of order five, six, and five for the specified quadratures, respectively. The error in each solution is measured with respect to the ∞ -norm at time t = 10 using the noted number of timesteps.

For an IDC or LS-IDC method with a maximum order of accuracy of M, the expected order of accuracy of the method is $\min\{M, P+1\}$ after P correction iterations when implicit Euler is used as the low-order method. Thus the methods are expected to attain their maximum order of accuracy after M-1 correction iterations have been performed. For each IDC and LS-IDC method, results are shown using M-1 and M correction iterations.

There are two primary observations to be made from these results: (i) the anticipated rates of convergence are attained by all methods, and (ii) the errors in the solutions produced by the LS-IDC methods are consistently larger than those of the corresponding IDC methods. We do note that LS-IDC methods with M correction iterations appear to produce convergence rates slightly larger than those expected as the timestep size becomes very small. We believe that this effect can be attributed to a combination of: (i) the second observation above, (ii) the fact that the difference between the LS-IDC and collocation solutions can decrease at a faster rate than the overall orders of the methods (c.f. proposition 4), and (iii) the results in section 2.3.2 demonstrating that the contraction factors of LS-IDC methods decrease rapidly for small timestep sizes. Hence the LS-IDC methods yield increasingly better approximations of the collocation solution as the timestep size becomes small, producing the observed behavior.

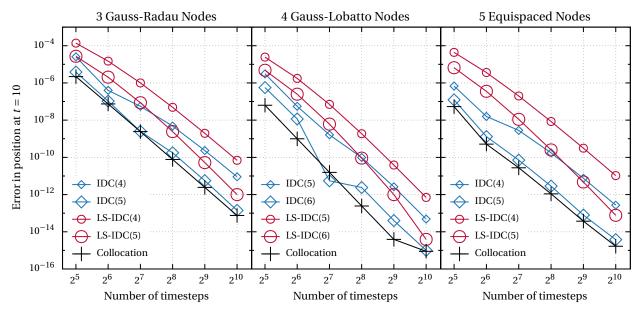


Figure 2.5: Convergence of IDC, LS-IDC, and collocation integrators using three Gauss-Radau nodes, four Gauss-Lobatto nodes, and five right-biased equispaced nodes for the IVP (2.50). Errors are computed with respect to the ∞ -norm at time t=10. The number of correction iterations used with each IDC and LS-IDC method is indicated in the key of each plot.

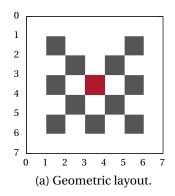
2.4.2 Linear transport equation

The second test problem is the linear transport equation (1.2). For the time being, we simply state that the system (1.2) is solved using a method of lines approach with a discontinuous Galerkin (DG) spatial discretization and a discrete ordinates angular discretization. Detailed descriptions of these methods and their application to (1.2) can be found in subsequent chapters of this thesis.

We consider a problem in two-dimensional *xy*-geometry with parameters outlined in figure 2.6 that we refer to as the *lattice problem*. A plot of the scalar flux distribution

$$\phi(\vec{x},t) = \int_{\mathbb{S}^2} \psi(\vec{x},\vec{\Omega},t) \, d\vec{\Omega} \tag{2.51}$$

is included in figure 2.6c for context. The goal of the tests conducted here is to assess the memory efficiency (i.e., error versus memory footprint) of different time integration schemes. This is done by applying each method to the lattice test problem using a fixed discrete ordinates approximation and a DG spatial discretization of the same order as the time integration scheme. Refinement studies are then performed with fixed CFL numbers of 32, 64, and 128. Two sets of studies are considered: one of third-order methods and one of fifth-order methods. An approximation of the error in each solution is computed by comparing the scalar flux distribution of each solution with that of a reference solution at the final time



Region	σ_{t}	$\sigma_{ m S}$	σ_{a}	q				
white	1	1	0	0				
gray	10	0	10	0				
red	1	1	0	1				
(b) Parameter values.								

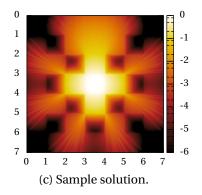


Figure 2.6: Layout and sample solution for the lattice test problem for the transport equation (1.2). Color scale of sample solution is log_{10} .

using the $L^2(d\vec{x})$ norm

$$e(t) = \left(\int_X \left(\phi(\vec{x}, t) - \phi^{\text{ref}}(\vec{x}, t) \right)^2 d\vec{x} \right)^{1/2}, \tag{2.52}$$

where ϕ^{ref} is the chosen reference solution. All solutions use a discrete ordinates approximation with a T_N quadrature [179] of order 16.

Four third-order time integration schemes are considered: an IDC method using two Gauss-Radau nodes and two correction iterations, which we refer to as IDC3; an LS-IDC method using two Gauss-Radau nodes and two correction iterations, which we refer to as LS-IDC3; the *L*-stable DIRK(3,3) method of [10], which we refer to as SDIRK3 (tableau given in figure 3.1b); and a linear DG time integration scheme that we refer to as STDG3. Five fifth-order schemes are considered: an IDC method using three Gauss-Radau nodes and four correction iterations, which we refer to as IDC5; an LS-IDC method using three Gauss-Radau nodes and four correction iterations, which we refer to as LS-IDC5; the *L*-stable, 7-stage DIRK method of [118], which we refer to as Kværnø7 (tableau given in figure 3.2a); the *L*-stable, 8-stage DIRK method of [110], which we refer to as ARK8 (tableau given in figure 3.2b); and a quadratic DG time integration scheme that we refer to as STDG5. It is useful to note that the time integration components of the STDG3 and STDG5 methods are equivalent to Radau collocation schemes using two and three quadrature nodes, respectively [125].

The tests were run on up to 16 nodes of the Laconia cluster at Michigan State University's Institute for Cyber-Enabled Research. Each Laconia node is equipped with two Intel Xeon E5-2690v4 processors and at least 128GB of RAM. Third-order methods were tested using between 56 and 896 mesh cells and fifth-order methods were tested using between 56 and 448 mesh cells, with the number of cells increasing by a factor of two in each refinement.

	STDG			IDC			LS-IDC			DIRK	
Method Order (odd):	3	5	N	3	5	N	3	5	N	3	5
Initial Condition	1	1	1	1	1	1	1	1	1	1	1
Source	1	1	1	1	1	1	1	1	1	1	1
Integrator Stages	2	3	(N+1)/2	4	6	N+1	2	3	(N+1)/2	3	6
Solver	2	3	(N+1)/2	1	1	1	1	1	1	1	1
Total	6	8	N+3	7	9	N+4	5	6	(N+1)/2+3	6	9

Table 2.2: Number of angular flux vectors required by each class of time integration scheme for solving the transport equation (1.2) using the discrete ordinates method. Here N represents the order of accuracy of the method, which may differ from the number of quadrature nodes or stages used by the method.

The memory requirements of solving (1.2) using the discrete ordinates method are determined primarily by the number of angular flux vectors that must be stored during each timestep update. The number of angular flux vectors required by each time integrator are summarized in table 2.2. There are four contributions to the number of angular flux vectors: (i) the initial condition for the timestep; (ii) a vector containing the source term q, which is assumed to be constant in time; (iii) the internal stages of the time integrator; and (iv) additional vectors required by the solver for the implicit systems. Each of the time integrators considered require one angular flux vector for the initial condition and one for the source term. The differences in memory requirements between the time integrators are due to differences in stage counts and storage requirements of the solvers. The STDG and LS-IDC integrators require storing the fewest number of stages simultaneously, and the stage counts for these methods are roughly half that of both the IDC and DIRK integrators. The number of angular flux vectors required by the solver for the implicit systems depends on the number of stages that are coupled into each solve: The solver for the IDC, LS-IDC, and DIRK integrators requires only one additional angular flux vector, irrespective of the order of the methods, and the number of additional angular flux vectors required by the solver for the STDG integrator is equal to the number of stages. The LS-IDC methods consistently produce the smallest memory footprint amongst all of the methods considered by minimizing both the number of stages that must be stored simultaneously and the number of additional angular flux vectors required by the implicit solver.

Figures 2.7a and 2.7b plot the total memory requirement of each solve for the third- and fifth-order methods, respectively, against the estimated $L^2(d\vec{x})$ error in the scalar flux distribution of each solution as measured relative to a refined reference solution. We note that the global convergence rates of the

methods in these tests are limited by the lack of regularity in the spatial variables. We observe that the reduced memory footprint of the LS-IDC methods does not appear to negatively impact the accuracy of the methods for this problem. For the third-order methods, the LS-IDC method appears to have a smaller error for a given timestep size that either the STDG or IDC method of the same order. It is not clear why this is the case, but we conjecture that the observed behavior is due to changes in the dispersive and dissipative properties of the methods when the low-storage modification is applied.

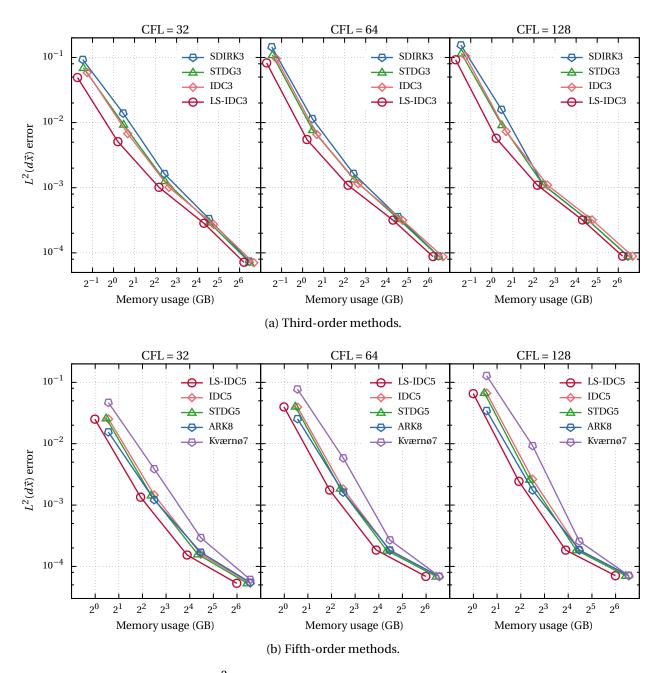


Figure 2.7: Memory usage vs. $L^2(d\vec{x})$ error in the scalar flux distribution for the lattice test problem.

CHAPTER 3

HIGH-ORDER HYBRID METHODS FOR RADIATION TRANSPORT

3.1 The transport equation and discrete ordinates methods

3.1.1 The transport equation

We consider the numerical solution of a mono-energetic, time-dependent, linear transport equation of the form (1.2). It is assumed that an inflow boundary condition ψ_B is known for all times $t \ge 0$ on the set

$$\left\{ \left(\vec{x}, \vec{\Omega}\right) \in \partial X \times \mathbb{S}^2 : \vec{n}(\vec{x}) \cdot \vec{\Omega} < 0 \right\},\tag{3.1}$$

where $\vec{n}(\vec{x})$ is the unit outward normal vector on the boundary ∂X of the spatial domain. In order to avoid complications due to problem geometry, the spatial domain X is assumed to be a convex polyhedron that is fixed for all times $t \ge 0$.

The operator notations used in this and subsequent chapters can be outlined as follows. Continuum operators will be written using non-bold, majuscule letters in a script font. For example, the equation (1.2) can be written in the form

$$\partial_t \psi = -\left(\mathcal{L} - \mathcal{SP}\right) \psi + q,\tag{3.2}$$

where

$$\mathcal{L} = \vec{\Omega} \cdot \nabla_{\vec{x}} + \sigma_{t}, \qquad \mathcal{S} = \frac{\sigma_{s}}{4\pi}, \quad \text{and} \quad \mathcal{P}(\cdot) = \int_{\mathbb{S}^{2}} (\cdot) \ d\vec{\Omega}'.$$
 (3.3)

As in (3.2), the continuum quantities to which these operators are applied will be written using non-bold, miniscule Greek or Latin letters. Approximations of operators and quantities with respect to the directional variable $\vec{\Omega}$ will be written using bold letters to denote vectors of components of the discretization that is applied; e.g., see the discrete ordinates system (3.5).

In the few cases where it is necessary to indicate that a spatial discretization has been applied, the discrete operators will be written using a Gothic font and the quantities to which these operators are applied will be written using majuscule Greek or Latin letters. Any such quantities may be written in bold or non-bold faces to indicate the discretization state of the directional variable $\vec{\Omega}$. Because Gothic typefaces are generally difficult to read and the focus of this work is the approximations with respect to

time and angle, the use of this notation will be avoided whenever possible. Instead, many discussions will focus on equations in semi-discrete form (i.e., without a spatial discretization having been applied) with the understanding that a spatial discretization would be applied before the corresponding systems are solved numerically.

3.1.2 A first-order discrete ordinates method

Given a suitable quadrature $\{\vec{\Omega}_k, \omega_k\}_{k=1}^K \subset \mathbb{S}^2 \times \mathbb{R}$, a discrete ordinates approximation for (1.2) satisfies the system of equations

$$\partial_t \psi_k + \vec{\Omega}_k \cdot \nabla_{\vec{x}} \psi_k + \sigma_t \psi_k = \frac{\sigma_s}{4\pi} \sum_{\ell=1}^K \omega_\ell \psi_\ell + q_k, \qquad (k = 1, \dots, K),$$
(3.4)

where $q_k = q(\vec{x}, \vec{\Omega}_k, t)$ and $\psi_k(\vec{x}, t)$ approximates $\psi(\vec{x}, \vec{\Omega}_k, t)$. The system (3.4) can be written in the compact form

$$\partial_t \boldsymbol{\psi} = -(\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi} + \boldsymbol{q}, \tag{3.5}$$

where

$$\boldsymbol{\psi} = [\psi_1, \dots, \psi_K]^T, \qquad \boldsymbol{q} = [q_1, \dots, q_K]^T, \tag{3.6}$$

the operator \mathcal{L} has the form

$$\mathcal{L} = \operatorname{Diag}(\mathcal{L}_1, \dots, \mathcal{L}_K) \quad \text{with} \quad \mathcal{L}_k = \vec{\Omega}_k \cdot \nabla_{\vec{x}} + \sigma_t$$
 (3.7)

for k = 1, ..., K, and the operators S and P are defined by

$$S = \left[\frac{\sigma_{S}}{4\pi}, \dots, \frac{\sigma_{S}}{4\pi}\right]^{T} \quad \text{and} \quad P = \left[\omega_{1}, \dots, \omega_{K}\right]. \tag{3.8}$$

Once a spatial discretization has been applied, (3.5) is transformed into the following initial value problem:

$$\partial_t \Psi = -(\mathfrak{L} - \mathfrak{S}\mathfrak{P})\Psi + \mathbf{Q}. \tag{3.9}$$

The quantities in (3.9) denote representations of discrete approximations to the quantities in (3.5); e.g., expansion coefficients with respect to some finite element basis for the spatial variable. The operators are derived from the original operators, and their action on the expansion coefficients is inherited from the action of the original operators.

We note that the operators \mathcal{L} and \mathfrak{L} are not linear, but affine: Inflow of radiation at the boundary of the domain results in a source term that forms the nonlinear component of the operator. The notations

$$\mathcal{L}(\cdot) = \widetilde{\mathcal{L}}(\cdot) - \psi_{B}$$
 and $\mathfrak{L}(\cdot) = \widetilde{\mathfrak{L}}(\cdot) - \Psi_{B}$ (3.10)

will be used to denote the decomposition of the operators \mathcal{L} and \mathfrak{L} into their linear components $\widetilde{\mathcal{L}}$ and $\widetilde{\mathfrak{L}}$, and boundary components ψ_B and Ψ_B , respectively, where the values of ψ_B and Ψ_B are determined from the incoming flux of radiation ψ_B on the boundary of the domain. The distinction between the linear and boundary components of these operators in not of critical importance in the present context, but will be necessary for correctly formulating the defect correction methods described later.

Applying an implicit Euler approximation to the time derivative in (3.9) with step size Δt yields a system of the form

$$(\widehat{\mathfrak{L}} - \mathfrak{S}\mathfrak{P})\Psi^n = \widehat{\mathbf{Q}}^n \tag{3.11}$$

that must be solved at each timestep, where

$$\widehat{\mathfrak{L}} = \widetilde{\mathfrak{L}} + \frac{1}{\Delta t} \mathfrak{I}, \qquad \widehat{\mathbf{Q}}^n = \frac{1}{\Delta t} \mathbf{\Psi}^{n-1} + \mathbf{Q}^n + \mathbf{\Psi}_{\mathrm{B}}^n, \tag{3.12}$$

and \mathfrak{I} is the identity operator. The linear system (3.11) is most commonly solved using one of several iterative methods built around inverting the streaming operator $\widehat{\mathfrak{L}}$ at each iteration. The traditional approach to solving (3.11) is a fixed-point method known as *source-iteration* [4]: given an initial guess $\Psi^{(0)}$ for Ψ^n , iterate

$$\mathbf{\Psi}^{(k)} = \widehat{\mathbf{\mathfrak{L}}}^{-1} \left(\mathfrak{S} \mathbf{\mathfrak{P}} \mathbf{\Psi}^{(k-1)} + \widehat{\mathbf{Q}}^n \right) \tag{3.13}$$

until the difference between successive iterates is below a given tolerance in some norm.

A more modern approach, which we employ for all numerical results in this work, is to use a Krylov subspace method. One of the primary advantages of Krylov subspace methods over traditional fixed-point methods such as source iteration is a reduction in the number of iterations required for convergence [183], particularly in strongly-scattering regimes. Using one of several matrix inversion identities [98], the solution Ψ^n of (3.11) can be written as

$$\Psi^{n} = \widehat{\mathbf{C}}^{-1} \left(\widehat{\mathbf{Q}}^{n} + \mathfrak{S} \left(\mathfrak{P} \Psi^{n} \right) \right), \tag{3.14}$$

where $\mathfrak{P}\Psi^n$ can be obtained by solving

$$\left(\mathfrak{I} - \mathfrak{P}\widehat{\mathfrak{L}}^{-1}\mathfrak{S}\right)\left(\mathfrak{P}\Psi^{n}\right) = \mathfrak{P}\widehat{\mathfrak{L}}^{-1}\widehat{\mathbf{Q}}^{n}. \tag{3.15}$$

The system (3.15) is solved using a matrix-free GMRES iteration to find $\mathfrak{P}\Psi^{n+1}$, which then yields Ψ^{n+1} through (3.14). The benefit of this approach is that the dimension of the system (3.15) to which the GMRES algorithm is applied is much smaller than the dimension of the original system (3.11). Hence the additional costs associated with the GMRES procedure, such as storage of previous iterates and construction of associated least-squares solutions, are significantly reduced.

Under some mild restrictions on the boundary conditions and the spatial discretization that is applied, direct inversion of $\widehat{\mathfrak{L}}$ can be accomplished via so-called transport sweeps. The term *transport sweep* is used to refer to a form of block forward substitution used to invert $\widehat{\mathfrak{L}}$ in which the solution in each mesh element is updated by traversing the spatial mesh in the characteristic direction determined by each ordinate direction, with the inflow boundary condition used to begin the process for each ordinate. Transport sweeps occupy the vast majority of the solve time required by each of the iterative methods described above, and a considerable amount of work has focused on efficient implementations of transport sweeps using various numerical approaches on a wide range of computer architectures – see, e.g., [5, 21, 60, 63, 101, 115, 150, 155, 157, 182, 189] and references therein for some examples.

The overall cost of these iterative solution procedures is determined by: (i) the cost of each iteration, which is dominated by the cost of the transport sweep, and (ii) the number of iterations required for the solver to converge within the desired tolerance. The interaction between these two factors is highly problem-dependent. Problems that are streaming dominated require many ordinates to mitigate the well-known "ray-effects" phenomenon [126, 127], but the ordinates are only weakly coupled and the resulting linear systems can be solved in relatively few iterations. For increasingly scattering-dominated problems the ordinates become more strongly coupled and more iterations are required. However, the increased coupling has the effect of smoothing out ray effects, so that fewer ordinates are necessary to resolve the bulk effects of the radiation. To obtain accurate solutions for problems with both streaming-and scattering-dominated regions, a uniform treatment of the angular flux requires large numbers of strongly coupled ordinates, producing systems that are very expensive to solve.

3.1.3 First-order hybrid discrete ordinates methods

We consider hybrid methods constructed following the first-collision source technique [8, 9], in which the angular flux ψ is decomposed into a collided component ψ_c and an uncollided component ψ_u with

 $\psi = \psi_{\rm u} + \psi_{\rm c}$ such that $\psi_{\rm u}$ and $\psi_{\rm c}$ satisfy

$$\partial_t \psi_{\mathbf{u}} + \vec{\Omega} \cdot \nabla_{\vec{x}} \psi_{\mathbf{u}} + \sigma_t \psi_{\mathbf{u}} = q, \tag{3.16a}$$

$$\partial_t \psi_{\mathcal{C}} + \vec{\Omega} \cdot \nabla_{\vec{x}} \psi_{\mathcal{C}} + \sigma_{\mathcal{C}} \psi_{\mathcal{C}} = \frac{\sigma_{\mathcal{S}}}{4\pi} \left[\langle \psi_{\mathcal{U}} \rangle + \langle \psi_{\mathcal{C}} \rangle \right]. \tag{3.16b}$$

In theory, there are multiple approaches that one could consider for dividing the initial and boundary conditions of the original system (1.2) between the collided and uncollided components of (3.16). Here we consider only the case where the uncollided system (3.16a) inherits the initial and boundary conditions of the original system (1.2), and the collided system (3.16b) is endowed with vacuum boundary conditions and an initial condition of zero.

The advantage of the splitting (3.16) is that it separates the limiting collisionality regimes of the original equation (1.2) in a self-consistent manner, with the free-streaming limit captured by the uncollided component and the diffusion limit captured by the collided component. Because the collision operator in this case is linear, this splitting is exact on the continuum level; that is, if ψ_u and ψ_c solve (3.16), then $\psi = \psi_u + \psi_c$ solves (1.2). However, if different numerical approximations are applied to the collided and uncollided components, then an effect that is similar to a splitting error will result.

Our intent is to utilize angular discretizations for the collided and uncollided components that are tailored to the limiting collisionality regimes captured by each component. Because solutions to problems near the free-streaming limit often contain a great deal of variation in the angular dimension, a higher-resolution angular approximation is used in the uncollided component in order to more accurately capture this behavior. The increasingly high collision rates found in problems that lie closer towards the diffusion limit have the effect of smoothing and diminishing the angular variation of the solution. Hence a lower-resolution angular approximation is used for the collided component to reduce computational cost.

In this work, discrete ordinates approximations are used for both the collided and uncollided components. Given two quadrature sets $\{\vec{\Omega}_{\mathrm{u},k},\omega_{\mathrm{u},k}\}_{k=1}^{K_{\mathrm{u}}} \subset \mathbb{S}^2 \times \mathbb{R}$ and $\{\vec{\Omega}_{\mathrm{c},k},\omega_{\mathrm{c},k}\}_{k=1}^{K_{\mathrm{c}}} \subset \mathbb{S}^2 \times \mathbb{R}$, a discrete ordinates approximation for (3.16) satisfies

$$\partial_t \boldsymbol{\psi}_{\mathbf{u}} = -\mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}} + \boldsymbol{q}, \tag{3.17a}$$

$$\partial_t \boldsymbol{\psi}_C = -(\mathcal{L}_C - \mathcal{S}_C \mathcal{P}_C) \boldsymbol{\psi}_C + \mathcal{S}_C \mathcal{P}_U \boldsymbol{\psi}_U, \tag{3.17b}$$

where

$$\psi_{\mathbf{u}} = [\psi_{\mathbf{u},1}, \dots, \psi_{\mathbf{u},K_{\mathbf{u}}}]^T, \qquad \psi_{\mathbf{c}} = [\psi_{\mathbf{c},1}, \dots, \psi_{\mathbf{c},K_{\mathbf{c}}}]^T, \quad \text{and} \quad \boldsymbol{q} = [q_1, \dots, q_{K_{\mathbf{u}}}]^T.$$
 (3.18)

Here $\psi_{\mathrm{u},k} = \psi_{\mathrm{u},k}(\vec{x},t)$ approximates $\psi_{\mathrm{u}}(\vec{x},\vec{\Omega}_{\mathrm{u},k},t)$, $\psi_{\mathrm{c},k} = \psi_{\mathrm{c},k}(\vec{x},t)$ approximates $\psi_{\mathrm{c}}(\vec{x},\vec{\Omega}_{\mathrm{c},k},t)$, and $q_k(\vec{x},t) = q(\vec{x},\vec{\Omega}_{\mathrm{u},k},t)$. The operators \mathcal{L}_{u} , \mathcal{L}_{c} , \mathcal{S}_{c} , \mathcal{P}_{u} , and \mathcal{P}_{c} are defined analogously to (3.7) and (3.8) with subscripts used to denote the quadrature associated with each operator.

The split system (3.17) is discretized using a strategy similar to the one used for the standard discrete ordinates system (3.5). Applying an implicit Euler approximation to the time derivatives in (3.17) with step size Δt yields

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n} = \frac{1}{\Delta t} \boldsymbol{\psi}_{\mathbf{u}}^{n-1} + \boldsymbol{q}^{n}, \tag{3.19a}$$

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{c}}\right) \boldsymbol{\psi}_{\mathbf{c}}^{n} = \frac{1}{\Delta t} \boldsymbol{\psi}_{\mathbf{c}}^{n-1} + \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{n}.$$
(3.19b)

Once a spatial discretization has been applied to (3.19), the resulting systems are solved sequentially using standard methods. Because (3.19a) does not contain a scattering source, the uncollided flux $\psi_{\rm u}^n$ can be computed using a single transport sweep. The collided flux $\psi_{\rm c}^n$ is obtained by solving (3.19b) using one of the iterative methods discussed in section 3.1.2. The cost of this solve is reduced by using fewer ordinates in the approximation of the collided flux. Because the number of ordinates $K_{\rm c}$ scales with the square of the quadrature order in two- and three-dimensional geometries, reducing the precision of the collided quadrature by small amounts can significantly decrease the cost of this solve.

In the absence of any external source (i.e., $q \equiv 0$ and $\psi_B \equiv 0$), the uncollided component ψ_U will decay exponentially in time at a rate determined by the total cross section σ_t . This will have the effect of reducing the accuracy of the split method (3.19) to that of the angular quadrature used in the collided component as the system is evolved in time. The approach introduced in [97] to mitigate this effect is to periodically "relabel" the collided flux as uncollided to reduce the proportion of the angular flux that is designated as collided at any given time. We consider the following generalization of this approach. Under the assumption that the initial flux distribution for each timestep is composed entirely of an uncollided component ψ_*^{n-1} and the collided flux is initially zero, (3.19) can be written as

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n} = \frac{1}{\Delta t} \boldsymbol{\psi}_{*}^{n-1} + \boldsymbol{q}^{n}, \tag{3.20a}$$

$$\left(\mathcal{L}_{c} + \frac{1}{\Delta t} - \mathcal{S}_{c} \mathcal{P}_{c}\right) \boldsymbol{\psi}_{c}^{n} = \mathcal{S}_{c} \mathcal{P}_{u} \boldsymbol{\psi}_{u}^{n}. \tag{3.20b}$$

Once the systems in (3.20) have been solved a single approximation of the resulting angular flux distribution at the time t_n , denoted by ψ^n_* , is constructed from the components ψ^n_u and ψ^n_c . In order to maintain the accuracy of the solution, ψ^n_* is constructed on the higher-resolution quadrature set used for the uncollided component. Because the approximations of the collided and uncollided components are computed using different angular quadratures, these components use different numerical representations: hence computing ψ^n_* requires a reconstruction procedure to map between these representations.

The idea introduced in [97] is to apply a reconstruction operator \mathcal{R} that maps the collided flux from the collided quadrature set to the uncollided quadrature set. The resulting flux distribution for the timestep is then defined explicitly by

$$\boldsymbol{\psi}_*^n = \boldsymbol{\psi}_1^n + \mathcal{R} \boldsymbol{\psi}_C^n. \tag{3.21}$$

A suitable definition of \mathcal{R} generally depends on the problem geometry and the choice of angular quadratures for the collided and uncollided components. The reconstruction operators used in this work are discussed in section 3.3.2.

Each timestep of the hybrid discrete ordinates method follows a three-step cadence: First solve the uncollided system, then solve the collided system, and finally apply a reconstruction to obtain the solution for the timestep. This process is outlined in detail in algorithm 3.1.

3.2 High-order time discretizations

Hybrid methods using four classes of high-order time integrators are now described. We consider diagonally-implicit Runge-Kutta (DIRK) methods, space-time discontinuous Galerkin (STDG) methods, integral deferred correction (IDC) methods, and the low-storage integral deferred correction (LS-IDC) methods introduced in chapter 2. Each time integration scheme is combined with a standard discrete ordinates approximation (we refer to these methods as *nonhybrid* methods) and several types of hybrid approximations.

A type specifier has been assigned to each hybrid method in order to allow the different hybrid approaches to be discussed in a reasonable manner. The hybrid methods for each integrator are grouped into at most two distinct types, each denoted by an upper-case Roman numeral. The type designation is used to characterize how the splitting is applied to the system. The types considered here are as follows:

Algorithm 3.1 Hybrid Euler timestep.

input: Initial condition ψ_*^{n-1} .

1: Compute ψ_{11}^{n} by solving (3.20a):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n} = \frac{1}{\Delta t} \boldsymbol{\psi}_{*}^{n-1} + \boldsymbol{q}^{n}.$$

2: Compute $\psi_{\rm C}^n$ by solving (3.20b):

$$\left(\mathcal{L}_{C} + \frac{1}{\Lambda t} - \mathcal{S}_{C} \mathcal{P}_{C}\right) \boldsymbol{\psi}_{C}^{n} = \mathcal{S}_{C} \mathcal{P}_{U} \boldsymbol{\psi}_{U}^{n}.$$

3: Apply reconstruction: $\psi_*^n \leftarrow \psi_u^n + \mathcal{R}\psi_c^n$.

return: ψ_*^n

- I. *Hybrid-I methods* are constructed by solving the continuum split system (3.17) using a standard method-of-lines approach and applying a reconstruction procedure between timesteps.
- II. *Hybrid-II methods* follow a different construction depending on the type of the time integration scheme. For more traditional time integrators such as DIRK methods, the hybrid-II methods are constructed by applying the time integrator before applying a splitting analogous to (3.16) at the semi-discrete level. For IDC and LS-IDC methods the hybrid-II splitting is unique in the sense that the splitting is wrapped inside of the defect correction iteration in a way such that the defect correction procedure is able to iteratively correct for the hybrid splitting that is employed. We will defer discussion of the hybrid-II IDC and LS-IDC methods to chapter 4.

The hybrid methods are further characterized by at most three subtypes, each denoted by a lower-case Latin letter. The subtype designation is used to characterize the reconstruction procedure that is applied. We consider only subtypes a and b in this chapter, which are characterized as follows:

- (a) *Hybrid-a methods* apply a reconstruction procedure at the end of each timestep that is based on a reconstruction operator \mathcal{R} maps from the collided quadrature set to the uncollided quadrature set directly.
- (b) *Hybrid-b methods* are specific to IDC and LS-IDC integrators. These methods apply a reconstruction procedure after each substep that is based on a reconstruction operator \mathcal{R} maps from the collided quadrature set to the uncollided quadrature set directly.

The subtype c is used in chapter 4 to designate that the method employs a Nyström reconstruction procedure.

We will not consider all possible combinations of all types of hybrid methods with all subtypes for each class of time integrator. Dirk methods of types Ia and IIa are described in section 3.2.1. For STDG integrators, the hybrid methods of type I and II are equivalent. We therefore drop the type designation for hybrid STDG methods and characterize them only by their subtype. The construction of the hybrid-a STDG methods is described in section 3.2.2 using the approach of the hybrid-I methods. Hybrid-I IDC and LS-IDC methods are described in sections 3.2.3 and 3.2.4, respectively, for subtypes a and b. As stated previously, discussions of hybrid-II IDC and LS-IDC methods is deferred to chapter 4.

3.2.1 Diagonally implicit Runge-Kutta methods

Given an initial value problem (IVP) of the form

$$y' = f(t, y), y(t_0) = y^0,$$
 (3.22)

a general *s*-stage Runge-Kutta (RK) method evolves an approximate solution of (3.22) in time through an algebraic system of equations of the form

$$y^{n} = y^{n-1} + \Delta t \sum_{i=1}^{s} b_{i} f\left(t_{n-1} + c_{i} \Delta t, Y^{(i)}\right),$$
(3.23a)

$$Y^{(i)} = y^{n-1} + \Delta t \sum_{j=1}^{s} a_{ij} f\left(t_{n-1} + c_j \Delta t, Y^{(j)}\right), \qquad (i = 1, ..., s), \qquad (3.23b)$$

where y^n approximates $y(t_n)$, $t_n = t_{n-1} + \Delta t$, and Δt denotes the timestep size. Each RK method is defined by the parameters $A = (a_{ij})_{i,j=1}^s$, $b = (b_i)_{i=1}^s$, and $c = (c_i)_{i=1}^s$, which may be chosen to satisfy various order and stability criteria (c.f., [15, 91, 92]). Diagonally-implicit RK (DIRK) methods are characterized by the fact that the coefficient matrix A is lower triangular; i.e., $a_{ij} = 0$ whenever j > i. This allows the stages $Y^{(i)}$ to be computed sequentially in a manner analogous to forward substitution.

Algorithm 3.2 Nonhybrid DIRK timestep.

input: Initial condition ψ^{n-1} .

1: **for** i = 1, ..., s **do**

2: Compute $oldsymbol{\psi}^{(i)}$ by solving (3.24b):

$$\left(\mathcal{L} + \frac{1}{\Delta t a_{ii}} - \mathcal{SP}\right) \boldsymbol{\psi}^{(i)} = \frac{1}{\Delta t a_{ii}} \boldsymbol{\psi}^{n-1} - \sum_{j=1}^{i-1} \frac{a_{ij}}{a_{ii}} \left[(\mathcal{L} - \mathcal{SP}) \, \boldsymbol{\psi}^{(j)} - \boldsymbol{q}^{(j)} \right] + \boldsymbol{q}^{(i)}$$

3: end for

4: if the DIRK scheme is stiffly accurate then

5: $\boldsymbol{\psi}^n \leftarrow \boldsymbol{\psi}^{(s)}$.

6: **else**

7:
$$\boldsymbol{\psi}^n$$
 is defined by (3.24a): $\boldsymbol{\psi}^n \leftarrow \boldsymbol{\psi}^{n-1} - \Delta t \sum_{i=1}^{s} b_i \left[(\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi}^{(i)} - \boldsymbol{q}^{(i)} \right]$.

8: **end if**

return: ψ^n

Nonhybrid DIRK method

Applying a DIRK scheme (3.23) to the standard discrete ordinates system (3.5) yields

$$\boldsymbol{\psi}^{n} = \boldsymbol{\psi}^{n-1} - \Delta t \sum_{i=1}^{s} b_{i} \left[(\mathcal{L} - \mathcal{SP}) \, \boldsymbol{\psi}^{(i)} - \boldsymbol{q}^{(i)} \right], \tag{3.24a}$$

$$\boldsymbol{\psi}^{(i)} = \boldsymbol{\psi}^{n-1} - \Delta t \sum_{j=1}^{i} a_{ij} \left[(\mathcal{L} - \mathcal{SP}) \, \boldsymbol{\psi}^{(j)} - \boldsymbol{q}^{(j)} \right], \qquad (i = 1, \dots, s),$$
(3.24b)

where $\mathbf{q}^{(i)} = \mathbf{q}(t_{n-1} + c_i \Delta t)$ for i = 1, ..., s. The nonhybrid DIRK timestep algorithm is outlined in algorithm 3.2.

Hybrid-I DIRK method

The derivation of the *hybrid-I* DIRK methods begins by applying a general DIRK scheme (3.23) to the split discrete ordinates system (3.17) with the collided flux initially zero and the initial condition ψ_*^{n-1}

Algorithm 3.3 Hybrid-Ia DIRK timestep.

input: Initial condition ψ_*^{n-1} .

1: **for** i = 1, ..., s **do**

Compute $\psi_{11}^{(i)}$ by solving (3.25c):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t a_{ii}}\right) \boldsymbol{\psi}_{\mathbf{u}}^{(i)} = \frac{1}{\Delta t a_{ii}} \boldsymbol{\psi}_{*}^{n-1} - \sum_{j=1}^{i-1} \frac{a_{ij}}{a_{ii}} \left[\mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{(j)} - \boldsymbol{q}^{(j)}\right] + \boldsymbol{q}^{(i)}.$$

Compute $\psi_c^{(i)}$ by solving (3.25d):

$$\left(\mathcal{L}_{\mathbf{C}} + \frac{1}{\Delta t a_{ii}} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{(i)} = -\sum_{i=1}^{i-1} \frac{a_{ij}}{a_{ii}} \left[\left(\mathcal{L}_{\mathbf{C}} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{(j)} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{(j)} \right] + \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{(i)}.$$

4: end for

5: if the DIRK scheme is stiffly accurate then

6:
$$\boldsymbol{\psi}_{11}^{n} \leftarrow \boldsymbol{\psi}_{11}^{(s)}$$
.

5: If the DIRK sche
6:
$$\psi_{\mathbf{u}}^{n} \leftarrow \psi_{\mathbf{u}}^{(s)}$$
.
7: $\psi_{\mathbf{c}}^{n} \leftarrow \psi_{\mathbf{c}}^{(s)}$.
8: **else**

9:
$$\boldsymbol{\psi}_{\mathbf{u}}^{n}$$
 is defined by (3.25a): $\boldsymbol{\psi}_{\mathbf{u}}^{n} \leftarrow \boldsymbol{\psi}_{*}^{n-1} - \Delta t \sum_{i=1}^{s} b_{i} \left[\mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{(i)} - \boldsymbol{q}^{(i)} \right]$.

10:
$$\boldsymbol{\psi}_{\mathbf{C}}^{n}$$
 is defined by (3.25b): $\boldsymbol{\psi}_{\mathbf{C}}^{n} \leftarrow -\Delta t \sum_{i=1}^{S} b_{i} \left[\left(\mathcal{L}_{\mathbf{C}} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}} \right) \boldsymbol{\psi}_{\mathbf{C}}^{(i)} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{(i)} \right].$

11: end if

12:
$$\boldsymbol{\psi}_*^n \leftarrow \boldsymbol{\psi}_u^n + \mathcal{R} \boldsymbol{\psi}_c^n$$
.

return: ψ_*^n

for the uncollided component. This yields

$$\boldsymbol{\psi}_{\mathbf{u}}^{n} = \boldsymbol{\psi}_{*}^{n-1} - \Delta t \sum_{i=1}^{s} b_{i} \left[\mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{(i)} - \boldsymbol{q}^{(i)} \right], \tag{3.25a}$$

$$\boldsymbol{\psi}_{c}^{n} = -\Delta t \sum_{i=1}^{s} b_{i} \left[\left(\mathcal{L}_{c} - \mathcal{S}_{c} \mathcal{P}_{c} \right) \boldsymbol{\psi}_{c}^{(i)} - \mathcal{S}_{c} \mathcal{P}_{u} \boldsymbol{\psi}_{u}^{(i)} \right], \tag{3.25b}$$

$$\psi_{\mathbf{u}}^{(i)} = \psi_{*}^{n-1} - \Delta t \sum_{j=1}^{i} a_{ij} \left[\mathcal{L}_{\mathbf{u}} \psi_{\mathbf{u}}^{(j)} - q^{(j)} \right], \qquad (i = 1, ..., s), \qquad (3.25c)$$

$$\boldsymbol{\psi}_{c}^{(i)} = -\Delta t \sum_{j=1}^{i} a_{ij} \left[\left(\mathcal{L}_{c} - \mathcal{S}_{c} \mathcal{P}_{c} \right) \boldsymbol{\psi}_{c}^{(j)} - \mathcal{S}_{c} \mathcal{P}_{u} \boldsymbol{\psi}_{u}^{(j)} \right], \qquad (i = 1, ..., s).$$
 (3.25d)

The stages $\{\psi_{\mathbf{u}}^{(i)}\}_{i=1}^s$ and $\{\psi_{\mathbf{c}}^{(i)}\}_{i=1}^s$ are computed using an alternating cadence between the uncollided and collided stages; i.e., in the order $\pmb{\psi}_u^{(1)}, \pmb{\psi}_c^{(1)}, \pmb{\psi}_u^{(2)}, \pmb{\psi}_c^{(2)}, \ldots$

The *hybrid-Ia* DIRK method is obtained by applying an operator \mathcal{R} to map the collided component

Algorithm 3.4 Hybrid-IIa DIRK timestep.

input: Initial condition ψ_*^{n-1} .

- Compute $\psi_{11}^{(i)}$ by solving (3.26a):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t a_{ii}}\right) \boldsymbol{\psi}_{\mathbf{u}}^{(i)} = \frac{1}{\Delta t a_{ii}} \boldsymbol{\psi}_{*}^{n-1} - \sum_{j=1}^{i-1} \frac{a_{ij}}{a_{ii}} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}}\right) \boldsymbol{\psi}_{*}^{(j)} - \boldsymbol{q}^{(j)} \right] + \boldsymbol{q}^{(i)}.$$

Compute $\psi_c^{(i)}$ by solving (3.26b):

$$\left(\mathcal{L}_{C} + \frac{1}{\Delta t a_{ii}} - \mathcal{S}_{C} \mathcal{P}_{C}\right) \psi_{C}^{(i)} = \mathcal{S}_{C} \mathcal{P}_{U} \psi_{U}^{(i)}.$$

- 4: $\boldsymbol{\psi}_{*}^{(i)} \leftarrow \boldsymbol{\psi}_{u}^{(i)} + \mathcal{R} \boldsymbol{\psi}_{c}^{(i)}$. 5: end for
- 6: if the DIRK scheme is stiffly accurate then

- $\boldsymbol{\psi}_{*}^{n}$ is defined by (3.28): $\boldsymbol{\psi}_{*}^{n} \leftarrow \boldsymbol{\psi}_{*}^{n-1} \Delta t \sum_{i=1}^{s} b_{i} \left[\left(\boldsymbol{\mathcal{L}}_{u} \boldsymbol{\mathcal{S}}_{u} \boldsymbol{\mathcal{P}}_{u} \right) \boldsymbol{\psi}_{*}^{(i)} \boldsymbol{q}^{(i)} \right].$
- 10: end if

return: ψ_{*}^{n}

at the final time to the uncollided quadrature set to obtain the reconstructed solution. Once $\psi^n_{\rm u}$ and $\psi^n_{\rm c}$ are known – either through (3.25a) and (3.25b), or directly from $\psi_{\rm II}^{(s)}$ and $\psi_{\rm C}^{(s)}$ if the DIRK scheme is stiffly accurate – the reconstructed solution ψ_*^n is given by (3.21). The hybrid-Ia DIRK timestep algorithm is outlined in algorithm 3.3.

Hybrid-II DIRK method

The hybrid-II DIRK methods are derived by changing the order in which the collision-based splitting and the time discretization are applied. In the derivation of the hybrid-I DIRK methods, the splitting was applied to the continuum transport equation (1.2) to obtain the split system (3.16), which was then discretized in angle, space, and time. For the hybrid-II DIRK methods, the time discretization is applied first. A collision-based splitting that is analogous to (3.16) is then applied to the equations that define the stage values, with the understanding that each such equation has the form of a steady-state transport equation. The resulting systems are then discretized in angle and space in a manner similar to the methods described earlier.

We omit a detailed derivation of the hybrid-II DIRK methods and simply state the final result. Each stage of the hybrid-II DIRK method is composed of three parts: an uncollided component $\boldsymbol{\psi}_{\mathrm{u}}^{(i)}$, a collided component $\boldsymbol{\psi}_{\mathrm{c}}^{(i)}$, and a reconstructed value $\boldsymbol{\psi}_{*}^{(i)}$ that is computed from $\boldsymbol{\psi}_{\mathrm{u}}^{(i)}$ and $\boldsymbol{\psi}_{\mathrm{c}}^{(i)}$. The components $\boldsymbol{\psi}_{\mathrm{u}}^{(i)}$ and $\boldsymbol{\psi}_{\mathrm{c}}^{(i)}$ satisfy

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t a_{ii}}\right) \boldsymbol{\psi}_{\mathbf{u}}^{(i)} = \frac{1}{\Delta t a_{ii}} \boldsymbol{\psi}_{*}^{n-1} - \sum_{j=1}^{i-1} \frac{a_{ij}}{a_{ii}} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}}\right) \boldsymbol{\psi}_{*}^{(j)} - \boldsymbol{q}^{(j)} \right] + \boldsymbol{q}^{(i)}, \quad (3.26a)$$

$$\left(\mathcal{L}_{C} + \frac{1}{\Delta t a_{ii}} - \mathcal{S}_{C} \mathcal{P}_{C}\right) \boldsymbol{\psi}_{C}^{(i)} = \mathcal{S}_{C} \mathcal{P}_{U} \boldsymbol{\psi}_{U}^{(i)}. \tag{3.26b}$$

Since the collided and uncollided components of each stage depend on the reconstructed values of the previous stages, the reconstruction procedure must be applied between stages. This is in contrast to the hybrid-Ia DIRK method, where the reconstruction procedure was applied only to the approximations at the final time after all of the stages had been computed.

The *hybrid-IIa* DIRK method is obtained by applying an operator \mathcal{R} to map the collided component for each stage to the uncollided quadrature set to obtain the reconstructed stage values. Once the components $\boldsymbol{\psi}_{\mathrm{u}}^{(i)}$ and $\boldsymbol{\psi}_{\mathrm{c}}^{(i)}$ for a stage are known, the reconstructed stage value $\boldsymbol{\psi}_{*}^{(i)}$ is given by

$$\boldsymbol{\psi}_{*}^{(i)} = \boldsymbol{\psi}_{\mathsf{u}}^{(i)} + \mathcal{R}\boldsymbol{\psi}_{\mathsf{c}}^{(i)}. \tag{3.27}$$

The solution for the timestep is computed using the following modified form of (3.24a)

$$\boldsymbol{\psi}_{*}^{n} = \boldsymbol{\psi}_{*}^{n-1} - \Delta t \sum_{i=1}^{s} b_{i} \left[\left(\mathcal{L}_{u} - \mathcal{S}_{u} \mathcal{P}_{u} \right) \boldsymbol{\psi}_{*}^{(i)} - \boldsymbol{q}^{(i)} \right], \tag{3.28}$$

or is obtained directly from $\psi_*^{(s)}$ if the DIRK scheme is stiffly accurate. The hybrid-IIa DIRK timestep algorithm is outlined in algorithm 3.4.

3.2.2 Space-time discontinuous Galerkin methods

We now consider traditional and hybrid discrete ordinates methods using space-time discontinuous Galerkin (STDG) methods of arbitrary order. Assume the spatial domain X is a convex polyhedron and let \mathcal{T}_h be a partition of X into non-overlapping polyhedral cells \mathcal{K} of diameter less than h. Given a time t_n for which an approximate solution ψ^n is known and a timestep size Δt , we consider advancement of the solution to the time $t_{n+1} = t_n + \Delta t$ using an orthogonal space-time mesh obtained by extruding the spatial mesh \mathcal{T}_h through the time interval $I_n = [t_n, t_{n+1}]$. This yields a space-time mesh $\widetilde{\mathcal{T}}_{h,n}$ that partitions the space-time slab $X \times I_n$ into polytopal cells of the form $\widetilde{\mathcal{K}} = \mathcal{K} \times I_n$ for each $\mathcal{K} \in \mathcal{T}_h$. Let $V_{h,n}$ be

a finite-dimensional space of functions of both space and time such that the restriction of each $v \in V_{h,n}$ to each cell $\widetilde{\mathcal{K}} \in \widetilde{\mathcal{T}}_{h,n}$ is continuous.

Nonhybrid STDG method

For each ordinate $\vec{\Omega}_k$, the boundary ∂X of the spatial domain can be decomposed into two components $\partial X = \partial X_k^- \cup \partial X_k^+$ where

$$\partial X_k^- = \left\{ \vec{x} \in \partial X : \vec{n}(\vec{x}) \cdot \vec{\Omega}_k < 0 \right\} \quad \text{and} \quad \partial X_k^+ = \left\{ \vec{x} \in \partial X : \vec{n}(\vec{x}) \cdot \vec{\Omega}_k \ge 0 \right\}$$
 (3.29)

are the inflow and outflow components of the boundary for the kth ordinate, respectively. Let \mathcal{E}_h denote the set of interior edges of the partition \mathcal{T}_h with $(\bigcup_{e \in \mathcal{E}_h} e) \cap \partial X = \emptyset$, and to each edge $e \in \mathcal{E}_h$ associate a unit vector \vec{n}_e normal to the edge whose orientation may be chosen arbitrarily. On each edge $e \in \partial X_k^{\pm}$ the associated unit normal vector \vec{n}_e is specified to be the outward normal.

For each ordinate, the STDG approximation of the solution ψ_k of the discrete ordinates system (3.5) is given by $\psi_k^{h,n} \in V_{h,n}$ satisfying

$$\mathcal{B}_{k}^{h,n}\left(\psi_{k}^{h,n},v^{h,n}\right) = \bar{\mathcal{B}}^{h,n}\left(\phi^{h,n},v^{h,n}\right) + \mathcal{A}_{k}^{h,n}\left(v^{h,n}\right), \qquad (k=1,\ldots,K),$$
(3.30)

for every $v^{h,n} \in V_{h,n}$ with the bilinear and linear forms

$$\mathcal{B}_{k}^{h,n}(u,v) = -\sum_{\widetilde{\mathcal{K}} \in \widetilde{\mathcal{T}}_{h,n}} \int_{\widetilde{\mathcal{K}}} \left(\partial_{t} v + \vec{\Omega}_{k} \cdot \nabla_{\vec{x}} v + \sigma_{t} v \right) u \, d\vec{x} \, dt + \sum_{\mathcal{K} \in \mathcal{T}_{h}} \int_{\mathcal{K}} (vu) \, \Big|_{t=t_{n+1}} \, d\vec{x}$$

$$-\sum_{e \in \mathcal{E}_{h}} \int_{e \times \mathbf{I}_{n}} \left(\vec{\Omega}_{k} \cdot \vec{n}_{e} \right) [v] \, \widehat{u} \, ds \, dt + \sum_{e \in \partial X_{k}^{+}} \int_{e \times \mathbf{I}_{n}} \left(\vec{\Omega}_{k} \cdot \vec{n}_{e} \right) v^{-} u^{-} \, ds \, dt,$$

$$(3.31a)$$

$$\bar{\mathcal{B}}^{h,n}(u,v) = \frac{1}{4\pi} \sum_{\widetilde{\mathcal{K}} \in \widetilde{\mathcal{T}}_{h,n}} \int_{\widetilde{\mathcal{K}}} \sigma_{s} v u \, d\vec{x} \, dt, \tag{3.31b}$$

$$\begin{split} \mathcal{A}_{k}^{h,n}(v) &= \sum_{\widetilde{\mathcal{K}} \in \widetilde{\mathcal{T}}_{h,n}} \int_{\widetilde{\mathcal{K}}} q_{k} v \, d\vec{x} \, dt - \sum_{e \in \partial X_{k}^{-}} \int_{e \times \mathbf{I}_{n}} \left(\vec{\Omega}_{k} \cdot \vec{n}_{e} \right) \psi_{\mathrm{B},k} v^{-} \, ds \, dt \\ &+ \sum_{\mathcal{K} \in \mathcal{T}_{n}} \int_{\mathcal{K}} \psi_{k}^{n} v \left(\vec{x}, t^{n} \right) \, d\vec{x}, \end{split} \tag{3.31c}$$

where

$$\boldsymbol{\psi}^{h,n} = \left[\psi_1^{h,n}, \dots, \psi_K^{h,n}\right]^T, \qquad \phi^{h,n} = \mathcal{P}\boldsymbol{\psi}^{h,n}, \tag{3.32}$$

 $\psi_{\mathrm{B},k} = \psi_{\mathrm{B}}(\vec{x},\vec{\Omega}_k,t)$ is the value of the inflow boundary condition specified on ∂X_k^- for the ordinate $\vec{\Omega}_k$,

$$v^{\pm}(\vec{x},t) = \lim_{\varepsilon \to 0^{+}} v(\vec{x} \pm \varepsilon \vec{n}_{e}, t), \qquad [v] = v^{+} - v^{-}, \tag{3.33}$$

and \hat{v} is a numerical flux specified in terms of v^+ and v^- . In order to maintain the direction of information flow imposed by the continuous operator \mathcal{L} , we use the *upwind flux*

$$\widehat{v}(\vec{x},t) = \lim_{\varepsilon \to 0^+} v(\vec{x} - \varepsilon \vec{\Omega}_k, t). \tag{3.34}$$

In practice (3.30) is written as a linear system relative to some basis of $V_{h,n}$. If $B_{h,n}$ is a vector whose elements form a basis of $V_{h,n}$ and $\Psi_k^{h,n}$ is the vector of coefficients such that $\psi_k^{h,n} = B_{h,n}^T \Psi_k^{h,n}$, then (3.30) can be written as

$$\mathfrak{L}^{h,n}\Psi^{h,n} = \mathfrak{S}^{h,n}\mathfrak{P}^{h,n}\Psi^{h,n} + \mathbf{Q}^{h,n},\tag{3.35}$$

where

$$\left(\mathbf{\Psi}^{h,n}\right)^{T} = \left[\left(\mathbf{\Psi}_{1}^{h,n}\right)^{T}, \dots, \left(\mathbf{\Psi}_{K}^{h,n}\right)^{T}\right] \quad \text{and} \quad \left(\mathbf{Q}^{h,n}\right)^{T} = \left[\left(Q_{1}^{h,n}\right)^{T}, \dots, \left(Q_{K}^{h,n}\right)^{T}\right] \quad (3.36)$$

with $Q_k^{h,n} = \mathcal{A}_k^{h,n}(B_{h,n})$; the operator $\mathfrak{L}^{h,n}$ has the form

$$\mathfrak{L}^{h,n} = \operatorname{Diag}\left(\mathfrak{L}_{1}^{h,n}, \dots, \mathfrak{L}_{K}^{h,n}\right) \quad \text{with} \quad \mathfrak{L}_{k}^{h,n} = \mathcal{B}_{k}^{h,n}\left(B_{h,n}^{T}, B_{h,n}\right)$$
(3.37)

for k = 1,...,K; and the operators $\mathfrak{S}^{h,n}$ and $\mathfrak{P}^{h,n}$ are defined by

$$\mathfrak{S}^{h,n} = \mathbb{1}_K \otimes \tilde{\mathfrak{S}}^{h,n} \quad \text{and} \quad \mathfrak{P}^{h,n} \Psi^{h,n} = \sum_{k=1}^K \omega_k \Psi_k^{h,n}, \tag{3.38}$$

where

$$\tilde{\mathfrak{S}}^{h,n} = \bar{\mathcal{B}}^{h,n} \left(B_{h,n}^T, B_{h,n} \right) \tag{3.39}$$

and $\mathbb{1}_K$ is the *K*-dimensional column vector of ones.

As was the case with the implicit time integration schemes in the previous sections, the STDG discretization yields a linear system (3.35) that must be solved at each timestep. The system (3.35) can be solved using the same iterative techniques described in section 3.1.2 where the streaming operator $\mathfrak{L}^{h,n}$ is inverted through transport sweeps. The approximate solution at the time t_{n+1} is given by evaluating $\psi^{h,n}$ at t_{n+1} .

Hybrid STDG method

Unlike the DIRK methods discussed in section 3.2.1, STDG integrators produce identical methods irrespective of the order in which the collision-based splitting and the time discretization are applied. This

is due to the fact that the temporal degrees of freedom in the STDG method are fully coupled, preventing the reconstruction procedure from being applied at intermediate stages. Here we consider discretization of the split system (3.16) in a manner similar to the STDG discrete ordinates method described previously.

Most of the previous discussion applies here with minor modifications to notation. The decomposition of the boundary $\partial X = \partial X_k^- \cup \partial X_k^+$ becomes $\partial X = \partial X_{\mathbf{u},k}^- \cup \partial X_{\mathbf{u},k}^+$ and $\partial X = \partial X_{\mathbf{c},k}^- \cup \partial X_{\mathbf{c},k}^+$ for the uncollided and collided ordinates respectively, where the components are defined similarly to (3.29) with $\vec{\Omega}_{\mathbf{u},k}$ and $\vec{\Omega}_{\mathbf{c},k}$ replacing $\vec{\Omega}_k$. For each ordinate, the STDG approximations of the uncollided and collided components $\psi_{\mathbf{u},k}$ and $\psi_{\mathbf{c},k}$ of the split system (3.16) are given by $\psi_{\mathbf{u},k}^{h,n}, \psi_{\mathbf{c},k}^{h,n} \in V_{h,n}$ satisfying

$$\mathcal{B}_{\mathbf{u},k}^{h,n} \left(\psi_{\mathbf{u},k}^{h,n}, \nu^{h,n} \right) = \mathcal{A}_{\mathbf{u},k}^{h,n} \left(\nu^{h,n} \right), \qquad (k = 1, \dots, K_{\mathbf{u}}), \tag{3.40a}$$

$$\mathcal{B}_{c,k}^{h,n}\left(\psi_{c,k}^{h,n}, v^{h,n}\right) = \bar{\mathcal{B}}^{h,n}\left(\phi^{h,n}, v^{h,n}\right), \qquad (k = 1, \dots, K_{\mathbb{C}}), \tag{3.40b}$$

for every $v^{h,n} \in V_{h,n}$. The bilinear forms $\mathcal{B}^{h,n}_{\mathbf{u},k}$ and $\mathcal{B}^{h,n}_{\mathbf{c},k}$ are defined analogously to (3.31a); i.e.,

$$\mathcal{B}_{\mathbf{u},k}^{h,n}(u,v) = -\sum_{\widetilde{\mathcal{K}} \in \widetilde{\mathcal{T}}_{h,n}} \int_{\widetilde{\mathcal{K}}} \left(\partial_{t} v + \vec{\Omega}_{\mathbf{u},k} \cdot \nabla_{\vec{x}} v + \sigma_{t} v \right) u \, d\vec{x} \, dt + \sum_{\mathcal{K} \in \mathcal{T}_{h}} \int_{\mathcal{K}} (vu) \, \Big|_{t=t_{n+1}} \, d\vec{x}$$

$$-\sum_{e \in \mathcal{E}_{h}} \int_{e \times \mathbf{I}_{n}} \left(\vec{\Omega}_{\mathbf{u},k} \cdot \vec{n}_{e} \right) [v] \, \widehat{u} \, ds \, dt + \sum_{e \in \partial X_{\mathbf{u},k}^{+}} \int_{e \times \mathbf{I}_{n}} \left(\vec{\Omega}_{\mathbf{u},k} \cdot \vec{n}_{e} \right) v^{-} u^{-} \, ds \, dt$$

$$(3.41)$$

and similarly for $\mathcal{B}^{h,n}_{\mathrm{c},k}$. The linear form $\mathcal{A}^{h,n}_{\mathrm{u},k}$ is defined by

$$\mathcal{A}_{\mathbf{u},k}^{h,n}(v) = \sum_{\widetilde{\mathcal{K}} \in \widetilde{\mathcal{T}}_{h,n}} \int_{\widetilde{\mathcal{K}}} q_k v \, d\vec{x} \, dt - \sum_{e \in \partial X_{\mathbf{u},k}^-} \int_{e \times \mathbf{I}_n} (\vec{\Omega}_{\mathbf{u},k} \cdot \vec{n}_e) \psi_{\mathbf{B},k} v^- \, ds \, dt + \sum_{\mathcal{K} \in \mathcal{T}_n} \int_{\mathcal{K}} \psi_{*,k}^n v(\vec{x},t^n) \, d\vec{x},$$

$$(3.42)$$

where

$$\psi_{*k}^{n} = \psi_{*}^{n}(\vec{x}, \vec{\Omega}_{U,k}, t), \qquad q_{k} = (\vec{x}, \vec{\Omega}_{U,k}, t), \qquad \text{and} \qquad \psi_{B,k} = \psi_{B}(\vec{x}, \vec{\Omega}_{U,k}, t),$$
 (3.43)

for $k=1,\ldots,K_{\rm U}$, are the initial condition, source, and boundary condition for the timestep, respectively, given on the uncollided discrete ordinates set. The bilinear form $\bar{\mathcal{B}}^{h,n}$ is again defined by (3.31b), but the scalar flux $\phi^{h,n}$ is defined here by

$$\phi^{h,n} = \mathcal{P}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{h,n} + \mathcal{P}_{\mathbf{c}} \boldsymbol{\psi}_{\mathbf{c}}^{h,n}, \tag{3.44}$$

Algorithm 3.5 Hybrid-a STDG timestep.

input: Initial condition ψ_*^n .

1: Compute $\Psi_{\rm U}^{h,n}$ by solving (3.46a):

$$\mathfrak{L}_{\mathbf{u}}^{h,n}\Psi_{\mathbf{u}}^{h,n}=\mathbf{Q}^{h,n}.$$

2: Compute $\Psi_{\rm C}^{h,n}$ by solving (3.46b):

$$\left(\mathfrak{L}_{c}^{h,n}-\mathfrak{S}_{c}^{h,n}\mathfrak{P}_{c}^{h,n}\right)\Psi_{c}^{h,n}=\mathfrak{S}_{c}^{h,n}\mathfrak{P}_{c}^{h,n}\Psi_{u}^{h,n}.$$

3: $\Psi_*^{h,n} \leftarrow \Psi_{11}^{h,n} + \Re \Psi_{C}^{h,n}$.

return: $\psi^{h,n}_* \Big|_{t=t_{n+1}}$

with

$$\boldsymbol{\psi}_{\mathbf{u}}^{h,n} = \left[\psi_{\mathbf{u},1}^{h,n}, \dots, \psi_{\mathbf{u},K_{\mathbf{u}}}^{h,n} \right]^{T} \quad \text{and} \quad \boldsymbol{\psi}_{\mathbf{c}}^{h,n} = \left[\psi_{\mathbf{c},1}^{h,n}, \dots, \psi_{\mathbf{c},K_{\mathbf{c}}}^{h,n} \right]^{T}.$$
 (3.45)

If $B_{h,n}$ is a vector whose elements form a basis of $V_{h,n}$ and $\Psi_{\mathbf{u},k}^{h,n}$ and $\Psi_{\mathbf{c},k}^{h,n}$ are vectors of coefficients such that $\psi_{\mathbf{u},k}^{h,n} = B_{h,n}^T \Psi_{\mathbf{u},k}^{h,n}$ and $\psi_{\mathbf{c},k}^{h,n} = B_{h,n}^T \Psi_{\mathbf{c},k}^{h,n}$, then (3.40) can be written as

$$\mathfrak{L}_{11}^{h,n}\Psi_{11}^{h,n} = \mathbf{Q}^{h,n},\tag{3.46a}$$

$$\left(\mathfrak{L}_{c}^{h,n} - \mathfrak{S}_{c}^{h,n}\mathfrak{P}_{c}^{h,n}\right)\Psi_{c}^{h,n} = \mathfrak{S}_{c}^{h,n}\mathfrak{P}_{u}^{h,n}\Psi_{u}^{h,n},\tag{3.46b}$$

where $\Psi_{\rm u}^{h,n}$, $\Psi_{\rm c}^{h,n}$, and $Q^{h,n}$ are defined analogously to (3.36); the operators $\mathfrak{L}_{\rm u}^{h,n}$ and $\mathfrak{L}_{\rm c}^{h,n}$ are defined analogously to (3.37); and $\mathfrak{S}_{\rm c}^{h,n}$, $\mathfrak{P}_{\rm c}^{h,n}$, and $\mathfrak{P}_{\rm u}^{h,n}$ are defined analogously to (3.38) and (3.39). The systems in (3.46) are solved in a manner similar to the previous hybrid methods: The uncollided system (3.46a) is solved using a single transport sweep, and then the collided system (3.46b) is solved using one of the iterative methods discussed in section 3.1.2.

Once the components $\psi_{\mathbf{u}}^{h,n}$ and $\psi_{\mathbf{c}}^{h,n}$ have been computed, a reconstruction is performed to obtain a single approximation $\psi_{*}^{h,n}$ to the angular flux with

$$\boldsymbol{\psi}_{*}^{h,n} = \left[\psi_{*,1}^{h,n}, \dots, \psi_{*,K_{\mathbf{u}}}^{h,n} \right]^{T}, \qquad \psi_{*,k}^{h,n} = B_{h,n}^{T} \Psi_{*,k}^{h,n}, \qquad \text{for } k = 1, \dots, K_{\mathbf{u}},$$
(3.47)

and $\Psi^{h,n}_*$ defined analogously to (3.36). The *hybrid-a* STDG method is obtained by applying an operator \mathcal{R} to map the collided component to the uncollided quadrature set to obtain the coefficients of the reconstructed approximation:

$$\Psi_*^{h,n} = \Psi_{11}^{h,n} + \Re \Psi_{C}^{h,n}. \tag{3.48}$$

The approximate solution at the time t_{n+1} is given by evaluating $\psi_*^{h,n}$ at t_{n+1} . The hybrid-a STDG timestep algorithm is outlined in algorithm 3.5.

3.2.3 Integral deferred correction methods

A detailed derivation of integral deferred correction (IDC) methods can be found in section 2.2.2. The notation used in this section deviates from that of section 2.2.2 in two respects. First, the substep index and correction iteration index are both written as superscripts; i.e., $\xi^{n,[0]}$ instead of $\xi^{[0]}_n$. Second, the conventions for Latin and Greek letters stated at the beginning of section 2.2.2 are not imposed here. The remainder of the notation introduced in section 2.2.2 for IDC methods is preserved. In this section the symbol $\delta_{n,\ell}$ denotes the Kronecker delta:

$$\delta_{n,\ell} = \begin{cases} 1, & \text{if } n = \ell, \\ 0, & \text{if } n \neq \ell. \end{cases}$$
(3.49)

Nonhybrid IDC method

Following the discussion in section 2.2.2, the IDC iteration begins by computing a set of provisional solutions $\{\psi^{n,[0]}\}_{n=1}^N$ satisfying an implicit Euler approximation of the discrete ordinates system (3.9), which can be written as

$$\boldsymbol{\psi}^{n,[0]} = \boldsymbol{\psi}^{n-1,[0]} - h_n \Delta t \left[(\mathcal{L} - \mathcal{SP}) \, \boldsymbol{\psi}^{n,[0]} - \boldsymbol{q}^n \right], \qquad (n = 1, ..., N), \tag{3.50}$$

where ψ^0 is the initial condition for the timestep, $\psi^{0,[p]} = \psi^0$ for all $p \ge 0$, and $q^n = q(t_n)$. Applying (2.17) to (3.9) yields the correction equation

$$\boldsymbol{\psi}^{n,[p]} = \boldsymbol{\psi}^{n-1,[p]} - h_n \Delta t (\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi}^{n,[p]} + h_n \Delta t (\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi}^{n,[p-1]}$$

$$- \Delta t \sum_{\ell=1}^{N} \gamma_{n,\ell} \left[(\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi}^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right],$$
(3.51)

for n = 1,...,N. The IDC timestep algorithm for the discrete ordinates system (3.9) is outlined in algorithm 3.6.

Algorithm 3.6 Nonhybrid IDC timestep.

input: Initial condition ψ^0 , number of correction iterations *P*.

1: **for** p = 0, ..., P **do**

⊳ Set initial condition.

2: $\psi^{0,[p]} \leftarrow \psi^0$.

3: end for

4: **for** n = 1, ..., N **do**

▶ Prediction steps.

5: Compute $\psi^{n,[0]}$ by solving (3.50):

$$\left(\mathcal{L} + \frac{1}{h_n \Delta t} - \mathcal{SP}\right) \boldsymbol{\psi}^{n,[0]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}^{n-1,[0]} + \boldsymbol{q}^n.$$

6: end for

7: **for** p = 1, ..., P **do**

➤ Correction iterations.

8: **for** n = 1, ..., N **do**

9: Compute $\psi^{n,[p]}$ by solving (3.51):

$$\left(\mathcal{L} + \frac{1}{h_n \Delta t} - \mathcal{SP}\right) \boldsymbol{\psi}^{n,[p]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}^{n-1,[p]} + \sum_{\ell=1}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_n}\right) (\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi}^{\ell,[p-1]} + \sum_{\ell=1}^{N} \frac{\gamma_{n,\ell}}{h_n} \boldsymbol{q}^{\ell}.$$

10: end for

11: end for

return: $\psi^{N,[P]}$

Hybrid-I IDC methods

Two subtypes of hybrid-I IDC methods are considered in this section: a and b. The construction of the hybrid-Ia IDC method proceeds as follows. Using the notation of (2.10), an implicit Euler approximation of (3.17) for the initial set of provisional solutions $\{\psi_{\mathbf{u}}^{n,[0]}\}_{n=1}^{N}$ and $\{\psi_{\mathbf{c}}^{n,[0]}\}_{n=1}^{N}$ takes the form

$$\boldsymbol{\psi}_{\mathbf{u}}^{n,[0]} = \boldsymbol{\psi}_{\mathbf{u}}^{n-1,[0]} - h_n \Delta t \left[\boldsymbol{\mathcal{L}}_{\mathbf{u}} \boldsymbol{\psi}^{n,[0]} - \boldsymbol{q}^n \right], \tag{3.52a}$$

$$\boldsymbol{\psi}_{\mathrm{C}}^{n,[0]} = \boldsymbol{\psi}_{\mathrm{C}}^{n-1,[0]} - h_{n} \Delta t \left[(\boldsymbol{\mathcal{L}}_{\mathrm{C}} - \boldsymbol{\mathcal{S}}_{\mathrm{C}} \boldsymbol{\mathcal{P}}_{\mathrm{C}}) \boldsymbol{\psi}_{\mathrm{C}}^{n,[0]} - \boldsymbol{\mathcal{S}}_{\mathrm{C}} \boldsymbol{\mathcal{P}}_{\mathrm{U}} \boldsymbol{\psi}_{\mathrm{U}}^{n,[0]} \right], \tag{3.52b}$$

for $n=1,\ldots,N$. It is assumed that the solution at the initial time for the step is entirely uncollided and that the collided flux is initially zero; i.e., $\psi_{\rm u}^{0,[p]}=\psi_*^0$ and $\psi_{\rm c}^{0,[p]}=0$ for all $p\geq 0$, where ψ_*^0 is the initial condition for the timestep.

Applying (2.17) to (3.17) yields the correction equation

$$\boldsymbol{\psi}_{\mathbf{u}}^{n,[p]} = \boldsymbol{\psi}_{\mathbf{u}}^{n-1,[p]} - h_{n} \Delta t \boldsymbol{\mathcal{L}}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{n,[p]} + h_{n} \Delta t \boldsymbol{\mathcal{L}}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{n,[p-1]} - \Delta t \sum_{\ell=1}^{N} \gamma_{n,\ell} \left[\boldsymbol{\mathcal{L}}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right], \quad (3.53a)$$

$$\boldsymbol{\psi}_{c}^{n,[p]} = \boldsymbol{\psi}_{c}^{n-1,[p]} - h_{n} \Delta t \left[\left(\mathcal{L}_{c} - \mathcal{S}_{c} \mathcal{P}_{c} \right) \boldsymbol{\psi}_{c}^{n,[p]} - \mathcal{S}_{c} \mathcal{P}_{u} \boldsymbol{\psi}_{u}^{n,[p]} \right]$$

$$+ h_{n} \Delta t \left[\left(\mathcal{L}_{c} - \mathcal{S}_{c} \mathcal{P}_{c} \right) \boldsymbol{\psi}_{c}^{n,[p-1]} - \mathcal{S}_{c} \mathcal{P}_{u} \boldsymbol{\psi}_{u}^{n,[p-1]} \right]$$

$$- \Delta t \sum_{\ell=1}^{N} \gamma_{n,\ell} \left[\left(\mathcal{L}_{c} - \mathcal{S}_{c} \mathcal{P}_{c} \right) \boldsymbol{\psi}_{c}^{\ell,[p-1]} - \mathcal{S}_{c} \mathcal{P}_{u} \boldsymbol{\psi}_{u}^{\ell,[p-1]} \right].$$

$$(3.53b)$$

The provisional solutions for the prediction level and correction levels are computed using an alternating cadence between the uncollided and collided components; i.e., in the order

$$\psi_{\mathbf{u}}^{1,[0]}, \psi_{\mathbf{c}}^{1,[0]}, \quad \psi_{\mathbf{u}}^{2,[0]}, \psi_{\mathbf{c}}^{2,[0]}, \dots, \quad \psi_{\mathbf{u}}^{1,[1]}, \psi_{\mathbf{c}}^{1,[1]}, \quad \psi_{\mathbf{u}}^{2,[1]}, \psi_{\mathbf{c}}^{2,[1]}, \dots$$
 (3.54)

The final approximation for the timestep is obtained by applying a reconstruction operator to map the approximation of the collided component from the last substep of the final correction iteration to the uncollided quadrature set:

$$\psi_*^{N,[P]} = \psi_{11}^{N,[P]} + \mathcal{R}\psi_{C}^{N,[P]}. \tag{3.55}$$

Here *P* denotes the number of correction iterations performed using (3.53). The hybrid-Ia IDC timestep algorithm is outlined in algorithm 3.7.

The hybrid-Ib IDC method is a modification of the hybrid-Ia IDC method in which the reconstruction operator \mathcal{R} is applied after each substep of the IDC prediction and correction iterations. This yields reconstructed approximations $\psi_*^{n,[p]}$ defined by

$$\psi_{*}^{n,[p]} = \psi_{11}^{n,[p]} + \mathcal{R}\psi_{C}^{n,[p]}, \tag{3.56}$$

for each n = 1,...,N and p = 0,...,P. Equations (3.52) and (3.53) are modified such that the initial conditions of the uncollided and collided fluxes are $\psi_*^{n-1,[p]}$ and zero, respectively, for each substep. More specifically, (3.52) becomes

$$\boldsymbol{\psi}_{\mathbf{u}}^{n,[0]} = \boldsymbol{\psi}_{*}^{n-1,[0]} - h_{n} \Delta t \left[\boldsymbol{\mathcal{L}}_{\mathbf{u}} \boldsymbol{\psi}^{n,[0]} - \boldsymbol{q}^{n} \right], \tag{3.57a}$$

$$\boldsymbol{\psi}_{c}^{n,[0]} = -h_{n}\Delta t \left[(\mathcal{L}_{c} - \mathcal{S}_{c}\mathcal{P}_{c}) \boldsymbol{\psi}_{c}^{n,[0]} - \mathcal{S}_{c}\mathcal{P}_{u} \boldsymbol{\psi}_{u}^{n,[0]} \right], \tag{3.57b}$$

Algorithm 3.7 Hybrid-Ia IDC timestep.

input: Initial condition ψ^0_* , number of correction iterations *P*.

⊳ Set initial condition.

- 1: **for** p = 0,...,P **do**2: $\psi_{\mathbf{u}}^{0,[p]} \leftarrow \psi_{*}^{0}$.
 3: $\psi_{\mathbf{c}}^{0,[p]} \leftarrow 0$.

5: **for** n = 1, ..., N **do**

▶ Prediction steps.

Compute $\psi_{\rm u}^{n,[0]}$ by solving (3.52a):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{h_n \Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n,[0]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}_{\mathbf{u}}^{n-1,[0]} + \boldsymbol{q}^n.$$

Compute $\psi_{C}^{n,[0]}$ by solving (3.52b):

$$\left(\mathcal{L}_{\mathbf{C}} + \frac{1}{h_n \Delta t} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{n,[0]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}_{\mathbf{U}}^{n-1,[0]} + \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{n,[0]}.$$

- 8: end for
- 9: **for** p = 1, ..., P **do**

▶ Correction iterations.

- **for** n = 1, ..., N **do** 10:
- Compute $\psi_{11}^{n,[p]}$ by solving (3.53a): 11:

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{h_n \Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n,[p]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}_{\mathbf{u}}^{n-1,[p]} + \sum_{\ell=1}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_n}\right) \mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{\ell,[p-1]} + \sum_{\ell=1}^{N} \frac{\gamma_{n,\ell}}{h_n} \boldsymbol{q}^{\ell}.$$

Compute $\psi_{C}^{n,[p]}$ by solving (3.53b): 12:

$$\begin{split} \left(\mathcal{L}_{\mathbf{C}} + \frac{1}{h_{n}\Delta t} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{n,[p]} &= \sum_{\ell=1}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_{n}}\right) \left[(\mathcal{L}_{\mathbf{C}} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}) \, \boldsymbol{\psi}_{\mathbf{C}}^{\ell,[p-1]} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{\ell,[p-1]} \right] \\ &+ \frac{1}{h_{n}\Delta t} \boldsymbol{\psi}_{\mathbf{C}}^{n-1,[p]} + \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{n,[p]}. \end{split}$$

- end for 13:
- 14: **end for**
- 15: $\psi_*^{N,[P]} \leftarrow \psi_{11}^{N,[P]} + \mathcal{R}\psi_{C}^{N,[P]}$

return: $\psi_*^{N,[P]}$

Algorithm 3.8 Hybrid-Ib IDC timestep.

input: Initial condition ψ^0_* , number of correction iterations *P*.

- 1: **for** p = 0, ..., P **do**
 - $\boldsymbol{\psi}_{*}^{0,[p]} \leftarrow \boldsymbol{\psi}_{*}^{0}$.
- 3: end for
- 4: **for** n = 1, ..., N **do**

▶ Prediction steps.

⊳ Set initial condition.

5: Compute $\psi_{\mathrm{U}}^{n,[0]}$ by solving (3.57a):

$$\left(\mathcal{L}_{\mathrm{U}} + \frac{1}{h_n \Delta t}\right) \boldsymbol{\psi}_{\mathrm{U}}^{n,[0]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}_{*}^{n-1,[0]} + \boldsymbol{q}^{n}.$$

6: Compute $\boldsymbol{\psi}_{\text{C}}^{n,[0]}$ by solving (3.57b):

$$\left(\mathcal{L}_{\mathbf{C}} + \frac{1}{h_{n}\Delta t} - \mathcal{S}_{\mathbf{C}}\mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{n,[0]} = \mathcal{S}_{\mathbf{C}}\mathcal{P}_{\mathbf{U}}\boldsymbol{\psi}_{\mathbf{U}}^{n,[0]}.$$

- 7: $\boldsymbol{\psi}_{*}^{n,[0]} \leftarrow \boldsymbol{\psi}_{\mathrm{u}}^{n,[0]} + \mathcal{R}\boldsymbol{\psi}_{\mathrm{c}}^{n,[0]}$.
- 8: end for
- 9: **for** p = 1,...,P **do**

▶ Correction iterations.

- 10: **for** n = 1, ..., N **do**
- 11: Compute $\psi_{\mathbf{u}}^{n,[p]}$ by solving (3.58a):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{h_n \Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n,[p]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}_{*}^{n-1,[p]} + \sum_{\ell=1}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_n}\right) \mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{\ell,[p-1]} + \sum_{\ell=1}^{N} \frac{\gamma_{n,\ell}}{h_n} \boldsymbol{q}^{\ell}.$$

12: Compute $\psi_{\rm C}^{n,[p]}$ by solving (3.58b):

$$\begin{split} \left(\mathcal{L}_{\mathbf{C}} + \frac{1}{h_{n}\Delta t} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{n,[p]} &= \sum_{\ell=1}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_{n}}\right) \left[\left(\mathcal{L}_{\mathbf{C}} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{\ell,[p-1]} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{\ell,[p-1]} \right] \\ &+ \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{n,[p]}. \end{split}$$

- 13: $\boldsymbol{\psi}_{*}^{n,[p]} \leftarrow \boldsymbol{\psi}_{\mathrm{u}}^{n,[p]} + \mathcal{R}\boldsymbol{\psi}_{\mathrm{c}}^{n,[p]}$.
- **14: end for**
- 15: **end for**

return: $\pmb{\psi}_*^{N,[P]}$

and (3.53) becomes

$$\boldsymbol{\psi}_{\mathbf{u}}^{n,[p]} = \boldsymbol{\psi}_{*}^{n-1,[p]} - h_{n} \Delta t \mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{n,[p]} + h_{n} \Delta t \mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{n,[p-1]} - \Delta t \sum_{\ell=1}^{N} \gamma_{n,\ell} \left[\mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right], \quad (3.58a)$$

$$\boldsymbol{\psi}_{\mathbf{c}}^{n,[p]} = -h_{n} \Delta t \left[\left(\mathcal{L}_{\mathbf{c}} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{c}} \right) \boldsymbol{\psi}_{\mathbf{c}}^{n,[p]} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{n,[p]} \right] + h_{n} \Delta t \left[\left(\mathcal{L}_{\mathbf{c}} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{c}} \right) \boldsymbol{\psi}_{\mathbf{c}}^{n,[p-1]} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{n,[p-1]} \right] \quad (3.58b)$$

for n = 1, ..., N. In this case, the provisional solutions are computed using a three-step cadence between the uncollided, collided, and reconstructed components; i.e., in the order

 $-\Delta t \sum_{\ell=1}^{N} \gamma_{n,\ell} \left[\left(\mathcal{L}_{c} - \mathcal{S}_{c} \mathcal{P}_{c} \right) \boldsymbol{\psi}_{c}^{\ell,[p-1]} - \mathcal{S}_{c} \mathcal{P}_{u} \boldsymbol{\psi}_{u}^{\ell,[p-1]} \right],$

$$\psi_{\mathbf{u}}^{1,[0]}, \psi_{\mathbf{c}}^{1,[0]}, \psi_{*}^{1,[0]}, \quad \psi_{\mathbf{u}}^{2,[0]}, \psi_{\mathbf{c}}^{2,[0]}, \psi_{*}^{2,[0]}, \dots, \\ \dots, \quad \psi_{\mathbf{u}}^{1,[1]}, \psi_{\mathbf{c}}^{1,[1]}, \psi_{*}^{1,[1]}, \quad \psi_{\mathbf{u}}^{2,[1]}, \psi_{\mathbf{c}}^{2,[1]}, \psi_{*}^{2,[1]}, \dots$$

$$(3.59)$$

The hybrid-Ib IDC timestep algorithm is outlined in algorithm 3.8.

3.2.4 Low-storage integral deferred correction methods

A detailed derivation of low-storage IDC (LS-IDC) methods can be found in section 2.2.3. The methods in this section differ from those in the preceding section (section 3.2.3) only in the way that the residuals are evaluated during the correction iterations. Because of the similarity of the IDC and LS-IDC methods, we give only very brief descriptions of the LS-IDC methods.

Nonhybrid LS-IDC method

The prediction step of the nonhybrid LS-IDC method is the same as that of the nonhybrid IDC method: An initial set of provisional solutions $\{\psi^{n,[0]}\}_{n=1}^{N}$ is computed using the implicit Euler approximation (3.50) of the discrete ordinates system (3.9). The LS-IDC correction equation is obtained through a simple modification of the residual term of (3.51):

$$\boldsymbol{\psi}^{n,[p]} = \boldsymbol{\psi}^{n-1,[p]} - h_n \Delta t (\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi}^{n,[p]} + h_n \Delta t (\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi}^{n,[p-1]}$$

$$- \Delta t \sum_{\ell=1}^{n-1} \gamma_{n,\ell} \left[(\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi}^{\ell,[p]} - \boldsymbol{q}^{\ell} \right] - \Delta t \sum_{\ell=n}^{N} \gamma_{n,\ell} \left[(\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi}^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right],$$
(3.60)

for n = 1,...,N. The LS-IDC timestep algorithm for the discrete ordinates system (3.9) is outlined in algorithm 3.9.

Hybrid-I LS-IDC methods

As before, two subtypes of hybrid-I LS-IDC methods are considered in this section: a and b. The prediction steps are the same as those of the hybrid-Ia and hybrid-Ib IDC methods: The initial provisional solutions $\{\psi_{\mathbf{u}}^{n,[0]}\}_{n=1}^{N}$ and $\{\psi_{\mathbf{c}}^{n,[0]}\}_{n=1}^{N}$ are computed using (3.52) and (3.57) for the hybrid-Ia and hybrid-Ib methods, respectively. The correction equations for the hybrid-Ia LS-IDC method are obtained by modifying the residual terms of (3.53):

$$\psi_{\mathbf{u}}^{n,[p]} = \psi_{\mathbf{u}}^{n-1,[p]} - h_{n} \Delta t \mathcal{L}_{\mathbf{u}} \psi_{\mathbf{u}}^{n,[p]} + h_{n} \Delta t \mathcal{L}_{\mathbf{u}} \psi_{\mathbf{u}}^{n,[p-1]}$$

$$- \Delta t \sum_{\ell=1}^{n-1} \gamma_{n,\ell} \left[\mathcal{L}_{\mathbf{u}} \psi_{\mathbf{u}}^{\ell,[p]} - q^{\ell} \right] - \Delta t \sum_{\ell=n}^{N} \gamma_{n,\ell} \left[\mathcal{L}_{\mathbf{u}} \psi_{\mathbf{u}}^{\ell,[p-1]} - q^{\ell} \right],$$

$$\psi_{\mathbf{c}}^{n,[p]} = \psi_{\mathbf{c}}^{n-1,[p]} - h_{n} \Delta t \left[\left(\mathcal{L}_{\mathbf{c}} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{c}} \right) \psi_{\mathbf{c}}^{n,[p]} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{u}} \psi_{\mathbf{u}}^{n,[p]} \right] + h_{n} \Delta t \left[\left(\mathcal{L}_{\mathbf{c}} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{c}} \right) \psi_{\mathbf{c}}^{n,[p-1]} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{u}} \psi_{\mathbf{u}}^{n,[p-1]} \right]$$

$$- \Delta t \sum_{\ell=1}^{n-1} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{c}} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{c}} \right) \psi_{\mathbf{c}}^{\ell,[p]} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{u}} \psi_{\mathbf{u}}^{\ell,[p]} \right]$$

$$- \Delta t \sum_{\ell=n}^{N} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{c}} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{c}} \right) \psi_{\mathbf{c}}^{\ell,[p-1]} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{u}} \psi_{\mathbf{u}}^{\ell,[p-1]} \right].$$
(3.61b)

Similarly, the correction equations for the hybrid-Ib LS-IDC method are obtained by modifying the residual terms of (3.58):

$$\psi_{\mathbf{u}}^{n,[p]} = \psi_{*}^{n-1,[p]} - h_{n} \Delta t \mathcal{L}_{\mathbf{u}} \psi_{\mathbf{u}}^{n,[p]} + h_{n} \Delta t \mathcal{L}_{\mathbf{u}} \psi_{\mathbf{u}}^{n,[p-1]}$$

$$- \Delta t \sum_{\ell=1}^{n-1} \gamma_{n,\ell} \left[\mathcal{L}_{\mathbf{u}} \psi_{\mathbf{u}}^{\ell,[p]} - \mathbf{q}^{\ell} \right] - \Delta t \sum_{\ell=n}^{N} \gamma_{n,\ell} \left[\mathcal{L}_{\mathbf{u}} \psi_{\mathbf{u}}^{\ell,[p-1]} - \mathbf{q}^{\ell} \right],$$

$$\psi_{\mathbf{c}}^{n,[p]} = -h_{n} \Delta t \left[\left(\mathcal{L}_{\mathbf{c}} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{c}} \right) \psi_{\mathbf{c}}^{n,[p]} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{u}} \psi_{\mathbf{u}}^{n,[p]} \right]$$

$$+ h_{n} \Delta t \left[\left(\mathcal{L}_{\mathbf{c}} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{c}} \right) \psi_{\mathbf{c}}^{n,[p-1]} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{u}} \psi_{\mathbf{u}}^{n,[p-1]} \right]$$

$$- \Delta t \sum_{\ell=1}^{n-1} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{c}} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{c}} \right) \psi_{\mathbf{c}}^{\ell,[p]} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{u}} \psi_{\mathbf{u}}^{\ell,[p]} \right]$$

$$- \Delta t \sum_{\ell=n}^{N} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{c}} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{c}} \right) \psi_{\mathbf{c}}^{\ell,[p-1]} - \mathcal{S}_{\mathbf{c}} \mathcal{P}_{\mathbf{u}} \psi_{\mathbf{u}}^{\ell,[p-1]} \right].$$

$$(3.62b)$$

The provisional solutions for the prediction level and correction levels are computed using the same alternating cadences as the corresponding IDC methods (i.e., (3.54) and (3.59)), and the reconstructed solutions are defined similarly as well (i.e., using (3.55) and (3.56)). The hybrid-Ia and hybrid-Ib LS-IDC timestep algorithms are outlined in algorithms 3.10 and 3.11, respectively.

Algorithm 3.9 Nonhybrid LS-IDC timestep.

input: Initial condition ψ^0 , number of correction iterations *P*.

for
$$p = 0,..., P$$
 do $\psi^{0,[p]} \leftarrow \psi^0$.

▶ Set initial condition.

end for

for n = 1, ..., N **do**

▶ Prediction steps.

Compute $\psi^{n,[0]}$ by solving (3.50):

$$\left(\mathcal{L} + \frac{1}{h_n \Delta t} - \mathcal{SP}\right) \boldsymbol{\psi}^{n,[0]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}^{n-1,[0]} + \boldsymbol{q}^n.$$

end for

for p = 1, ..., P **do**

➤ Correction iterations.

for n = 1, ..., N **do**

Compute $\psi^{n,[p]}$ by solving (3.60):

$$\begin{split} \left(\mathcal{L} + \frac{1}{h_n \Delta t} - \mathcal{SP}\right) \boldsymbol{\psi}^{n,[p]} &= \frac{1}{h_n \Delta t} \boldsymbol{\psi}^{n-1,[p]} - \sum_{\ell=1}^{n-1} \frac{\gamma_{n,\ell}}{h_n} \left(\mathcal{L} - \mathcal{SP}\right) \boldsymbol{\psi}^{\ell,[p]} \\ &+ \sum_{\ell=n}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_n}\right) \left(\mathcal{L} - \mathcal{SP}\right) \boldsymbol{\psi}^{\ell,[p-1]} + \sum_{\ell=1}^{N} \frac{\gamma_{n,\ell}}{h_n} \boldsymbol{q}^{\ell} \end{split}$$

end for end for

return: $\psi^{N,[P]}$

3.3 Implementation details

This section provides a detailed description of the numerical methods used in our implementations and the numerical tests in subsequent sections. Solving the kinetic equation (1.2) in full generality requires significant computational resources. In order to reduce the computational demands of the tests conducted for this work, we consider test cases in the following reduced geometries [135, Section 1-3]:

(i) One spatial dimension (slab geometry): In the case that $\partial_{\chi}\psi = \partial_{\gamma}\psi = 0$, (1.2) can be reduced to

$$\partial_t \psi + \mu \partial_z \psi + \sigma_t \psi = \frac{\sigma_s}{2} \int_{-1}^1 \psi(\mu') d\mu' + q$$
 (3.63)

where $\mu \in [-1, 1]$ is the *z*-component of $\vec{\Omega} = (\xi, \eta, \mu) \in \mathbb{S}^2$.

(ii) Two spatial dimensions (xy-geometry): In the case that $\partial_z \psi = 0$ the angular flux is considered as a function of xy-position only, reducing (1.2) to

$$\partial_t \psi + \xi \partial_x \psi + \eta \partial_y \psi + \sigma_t \psi = \frac{\sigma_s}{4\pi} \int_{S^2} \psi \left(\vec{\Omega}'\right) d\vec{\Omega}' + q. \tag{3.64}$$

Here the angular space is not reduced as it is for slab geometry. While it is possible to reduce the angular space to one hemisphere (e.g., $\mu \ge 0$) due to symmetry in the z-dimension, this reduces the number of ordinates by the same constant factor (one half) for all discrete ordinates sets considered in this work. Hence relative performance comparisons between discrete ordinates approximations of different resolutions should not depend significantly on whether this reduction is performed or not. We do not consider this type of reduction so that discussions of the numerical treatment of the angular variable in this geometry extend naturally to the full problem in three spatial dimensions.

Throughout this section, $N_{\rm U}$ and $N_{\rm C}$ are used to represent positive integers that denote the resolution of the high- and low-resolution angular quadratures used for the uncollided and collided components, respectively. The number of ordinates $K_{\rm U}$ and $K_{\rm C}$ in each quadrature corresponding to the values of $N_{\rm U}$ and $N_{\rm C}$ will depend on the type of the quadrature that is used. For each of the quadratures considered here a relationship of the form $K_{\rm U} = f(N_{\rm U})$ can be established where f is a simple algebraic expression.

3.3.1 Specifics of high-order time discretizations

The high-order time discretizations described in section 3.2 are discussed in general forms: each class requires that certain parameters be specified in order to generate a particular scheme of that type. We now provide detailed descriptions of the schemes considered for the numerical tests conducted for this work.

DIRK schemes

We consider four DIRK schemes:

• The *L*-stable DIRK(2,2) scheme of [10], which we refer to as SDIRK2. The scheme is second order and stiffly accurate. The Butcher tableau of the scheme is given in figure 3.1, with the parameter

$$\alpha = 1 - \frac{\sqrt{2}}{2}.\tag{3.65}$$

• The L-stable DIRK(3,3) scheme of [10], which we refer to as SDIRK3. The scheme is third order and

Algorithm 3.10 Hybrid-Ia LS-IDC timestep.

input: Initial condition ψ^0_* , number of correction iterations *P*.

$$\label{eq:posterior} \begin{aligned} & \text{for } p = 0, \dots, P \text{ do} \\ & \pmb{\psi}_{u}^{0,[p]} \leftarrow \pmb{\psi}_{*}^{0}; \quad \pmb{\psi}_{c}^{0,[p]} \leftarrow 0. \\ & \text{end for} \end{aligned}$$

> Set initial condition.

for n = 1, ..., N **do**

▶ Prediction steps.

Compute $\psi_{11}^{n,[0]}$ by solving (3.52a):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{h_n \Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n,[0]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}_{\mathbf{u}}^{n-1,[0]} + \boldsymbol{q}^n.$$

Compute $\psi_{\rm C}^{n,[0]}$ by solving (3.52b):

$$\left(\mathcal{L}_{\mathsf{C}} + \frac{1}{h_n \Delta t} - \mathcal{S}_{\mathsf{C}} \mathcal{P}_{\mathsf{C}}\right) \psi_{\mathsf{C}}^{n,[0]} = \frac{1}{h_n \Delta t} \psi_{\mathsf{u}}^{n-1,[0]} + \mathcal{S}_{\mathsf{C}} \mathcal{P}_{\mathsf{u}} \psi_{\mathsf{u}}^{n,[0]}.$$

end for

for p = 1, ..., P **do**

➤ Correction iterations.

for n = 1, ..., N **do**

Compute $\psi_{\mathrm{II}}^{n,[p]}$ by solving (3.61a):

$$\begin{split} \left(\mathcal{L}_{\mathbf{u}} + \frac{1}{h_{n}\Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n,[p]} &= \frac{1}{h_{n}\Delta t} \boldsymbol{\psi}_{\mathbf{u}}^{n-1,[p]} - \sum_{\ell=1}^{n-1} \frac{\gamma_{n,\ell}}{h_{n}} \mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{\ell,[p]} + \sum_{\ell=n}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_{n}}\right) \mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{\ell,[p-1]} \\ &+ \sum_{\ell=1}^{N} \frac{\gamma_{n,\ell}}{h_{n}} \boldsymbol{q}^{\ell}. \end{split}$$

Compute $\psi_{\rm C}^{n,[p]}$ by solving (3.61b):

$$\begin{split} \left(\mathcal{L}_{\mathrm{C}} + \frac{1}{h_{n}\Delta t} - \mathcal{S}_{\mathrm{C}}\mathcal{P}_{\mathrm{C}}\right) \boldsymbol{\psi}_{\mathrm{C}}^{n,[p]} &= \frac{1}{h_{n}\Delta t} \boldsymbol{\psi}_{\mathrm{C}}^{n-1,[p]} + \mathcal{S}_{\mathrm{C}}\mathcal{P}_{\mathrm{U}}\boldsymbol{\psi}_{\mathrm{U}}^{n,[p]} \\ &- \sum_{\ell=1}^{n-1} \frac{\gamma_{n,\ell}}{h_{n}} \left[\left(\mathcal{L}_{\mathrm{C}} - \mathcal{S}_{\mathrm{C}}\mathcal{P}_{\mathrm{C}}\right) \boldsymbol{\psi}_{\mathrm{C}}^{\ell,[p]} - \mathcal{S}_{\mathrm{C}}\mathcal{P}_{\mathrm{U}}\boldsymbol{\psi}_{\mathrm{U}}^{\ell,[p]} \right] \\ &+ \sum_{\ell=n}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_{n}}\right) \left[\left(\mathcal{L}_{\mathrm{C}} - \mathcal{S}_{\mathrm{C}}\mathcal{P}_{\mathrm{C}}\right) \boldsymbol{\psi}_{\mathrm{C}}^{\ell,[p-1]} - \mathcal{S}_{\mathrm{C}}\mathcal{P}_{\mathrm{U}}\boldsymbol{\psi}_{\mathrm{U}}^{\ell,[p-1]} \right]. \end{split}$$

end for

end for

$$\boldsymbol{\psi}_{*}^{N,[P]} \leftarrow \boldsymbol{\psi}_{\mathrm{u}}^{N,[P]} + \mathcal{R}\boldsymbol{\psi}_{\mathrm{c}}^{N,[P]}$$

return: $\pmb{\psi}_*^{N,[P]}$

Algorithm 3.11 Hybrid-Ib LS-IDC timestep.

input: Initial condition ψ_*^0 , number of correction iterations *P*.

for
$$p = 0,...,P$$
 do $\psi_*^{0,[p]} \leftarrow \psi_*^0$. end for

⊳ Set initial condition.

for n = 1, ..., N **do**

Compute $\psi_{11}^{n,[0]}$ by solving (3.57a):

 $\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{h_n \Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n,[0]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}_{*}^{n-1,[0]} + \boldsymbol{q}^{n}.$

Compute $\psi_{C}^{n,[0]}$ by solving (3.57b):

$$\left(\mathcal{L}_{\mathbf{C}} + \frac{1}{h_n \Delta t} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}\right) \psi_{\mathbf{C}}^{n,[0]} = \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \psi_{\mathbf{U}}^{n,[0]}.$$

$$\boldsymbol{\psi}_{*}^{n,[0]} \leftarrow \boldsymbol{\psi}_{\mathrm{u}}^{n,[0]} + \mathcal{R}\boldsymbol{\psi}_{\mathrm{c}}^{n,[0]}.$$
 end for

for p = 1, ..., P **do**

▶ Correction iterations.

for n = 1,...,N do
Compute $\psi_{11}^{n,[p]}$ by solving (3.62a):

$$\begin{split} \left(\mathcal{L}_{\mathbf{u}} + \frac{1}{h_{n}\Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n,[p]} &= \frac{1}{h_{n}\Delta t} \boldsymbol{\psi}_{*}^{n-1,[p]} - \sum_{\ell=1}^{n-1} \frac{\gamma_{n,\ell}}{h_{n}} \mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{\ell,[p]} + \sum_{\ell=n}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_{n}}\right) \mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{\ell,[p-1]} \\ &+ \sum_{\ell=1}^{N} \frac{\gamma_{n,\ell}}{h_{n}} \boldsymbol{q}^{\ell}. \end{split}$$

Compute $\psi_c^{n,[p]}$ by solving (3.62b):

$$\begin{split} \left(\mathcal{L}_{\mathbf{C}} + \frac{1}{h_{n}\Delta t} - \mathcal{S}_{\mathbf{C}}\mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{n,[p]} &= \mathcal{S}_{\mathbf{C}}\mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{n,[p]} - \sum_{\ell=1}^{n-1} \frac{\gamma_{n,\ell}}{h_{n}} \left[\left(\mathcal{L}_{\mathbf{C}} - \mathcal{S}_{\mathbf{C}}\mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{\ell,[p]} - \mathcal{S}_{\mathbf{C}}\mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{\ell,[p]} \right] \\ &+ \sum_{\ell=n}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_{n}}\right) \left[\left(\mathcal{L}_{\mathbf{C}} - \mathcal{S}_{\mathbf{C}}\mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{\ell,[p-1]} - \mathcal{S}_{\mathbf{C}}\mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{\ell,[p-1]} \right]. \end{split}$$

$$\boldsymbol{\psi}_*^{n,[p]} \leftarrow \boldsymbol{\psi}_u^{n,[p]} + \mathcal{R} \boldsymbol{\psi}_c^{n,[p]}.$$
 end for

end for

return: $\psi_*^{N,[P]}$

(a) SDIRK2 with α given by (3.65).

(b) SDIRK3 with parameters given by (3.66).

Figure 3.1: Butcher tableaus of SDIRK2 and SDIRK3 schemes.

stiffly accurate. The Butcher tableau of the scheme is given in figure 3.1, with the parameters

$$\alpha = 0.435866521508459, \tag{3.66a}$$

$$\tau = \frac{\alpha + 1}{6},\tag{3.66b}$$

$$b_1 = -\frac{6\alpha^2 - 16\alpha + 1}{4},\tag{3.66c}$$

$$b_2 = \frac{6\alpha^2 - 20\alpha + 5}{4}. ag{3.66d}$$

- The *L*-stable, 7-stage DIRK method of [118], which we refer to as Kværnø7. The scheme is fifth order and stiffly accurate. The Butcher tableau of the scheme is given in figure 3.2.
- The *L*-stable, 8-stage DIRK method of [110], which we refer to as ARK8. The scheme is fifth order and stiffly accurate. The Butcher tableau of the scheme is given in figure 3.2.

While many of these schemes are constructed with embedded pairs to provide error estimates for adaptive timestepping, we do not make use of these capabilities and hence the embedded pairs are omitted.

IDC and LS-IDC schemes

Each IDC and LS-IDC scheme is generated by specifying the quadrature nodes $\{c_n\}$ which form the substeps of the method. The substep sizes h_n and the quadrature weights for evaluating the residual $\gamma_{n,k}$ are computed from the positions of the quadrature nodes using (2.11) and (2.18), respectively (see section 2.2.2 for a thorough description of IDC methods). We consider IDC and LS-IDC methods constructed with two quadratures: the two- and three-point Gauss-Radau quadratures. These quadratures yield IDC and LS-IDC methods of order three and five, respectively, and we refer to these methods as IDC3 and IDC5, and LS-IDC3 and LS-IDC5, respectively.

0.0	0.0						
0.52	0.26	0.26					
1.230333209967908	0.13	0.84033320996790809	0.26				
0.895765984350076	0.22371961478320505	0.47675532319799699	-0.06470895363112615	0.26			
0.436393609858648	0.16648564323248321	0.10450018841591720	0.03631482272098715	-0.13090704451073998	0.26		
1.0	0.13855640231268224	0.0	-0.04245337201752043	0.02446657898003141	0.61943039072480676	0.26	
1.0	0.13659751177640291	0.0	-0.05496908796538376	-0.04118626728321046	0.62993304899016403	0.06962479448202728	0.26
	0.13659751177640291	0.0	-0.05496908796538376	-0.04118626728321046	0.62993304899016403	0.06962479448202728	0.26

(a) Butcher tableau of the Kværnø7 scheme.

0.0	0										
$\frac{41}{100}$	$\frac{41}{200}$	$\frac{41}{200}$									
$\frac{2935347310677}{11292855782101}$	$\frac{41}{400}$	$-\frac{567603406766}{11931857230679}$	$\frac{41}{200}$								
$\tfrac{1426016391358}{7196633302097}$	683785636431 9252920307686	0	$-\frac{110385047103}{1367015193373}$	$\frac{41}{200}$							
$\frac{92}{100}$	3016520224154 10081342136671	0	$\frac{30586259806659}{12414158314087}$	$-\frac{22760509404356}{11113319521817}$	$\frac{41}{200}$						
$\frac{24}{100}$	218866479029 1489978393911	0	638256894668 5436446318841	$-\frac{1179710474555}{5321154724896}$	$-\frac{60928119172}{8023461067671}$	$\frac{41}{200}$					
<u>3</u> 5	1020004230633 5715676835656	0	25762820946817 25263940353407	$-\frac{2161375909145}{9755907335909}$	$-\frac{211217309593}{5846859502534}$	$-\frac{4269925059573}{7827059040749}$	$\frac{41}{200}$				
1	$-\frac{872700587467}{9133579230613}$	0	0	22348218063261 9555858737531	$-\frac{1143369518992}{8141816002931}$	$-\frac{39379526789629}{19018526304540}$	$\tfrac{32727382324388}{42900044865799}$	$\frac{41}{200}$			
	$-\frac{872700587467}{9133579230613}$	0	0	22348218063261 9555858737531	$-\frac{1143369518992}{8141816002931}$	$-\frac{39379526789629}{19018526304540}$	32727382324388 42900044865799	$\frac{41}{200}$			
(b) ARK8											

Figure 3.2: Butcher tableaus of Kværnø7 and ARK8 schemes.

STDG schemes

Each STDG scheme is generated by specifying the space $V_{h,n}$ in which the approximate solution lies. Because the space-time mesh is obtained by extruding a spatial mesh through the timestep interval, the form of the spatial and temporal components may be specified separately: the resulting space-time approximation space is composed of the tensor product of the spatial and temporal components.

We consider two STDG schemes, which we refer to as STDG3 and STDG5. The STDG3 scheme uses a tensor product basis in each mesh element that consists of a polynomial of degree 2 in each spatial dimension and a polynomial of degree 1 in the time dimension. The STDG5 scheme uses a tensor product basis in each mesh element that consists of a polynomial of degree 4 in each spatial dimension and a polynomial of degree 2 in the time dimension. These methods are formally third- and fifth-order accurate in space and second- and third-order accurate in time, respectively, at all space-time points. Additionally, the schemes are third- and fifth-order superconvergent in time, respectively, at the timestep points (similar superconvergence results can be shown for values along certain boundaries of the spatial mesh) [122, 124, 125]. Further, the time-integration components of the methods can be shown to be equivalent to Radau collocation schemes, which are *L*-stable at all orders of accuracy [64, 125].

The approximation spaces used for the STDG methods are chosen such that the order of accuracy of the methods aligns with the choices made for DIRK and IDC methods. For each test problem considered in this work, solution errors are measured at the final simulation time by comparing the scalar flux distribution of each approximate solution with that of a reference solution using a continuum L^p distance metric with respect to the spatial variables; e.g., (2.52). The DG spatial discretizations that are used (c.f. section 3.3.3) are not superconvergent in this norm. Since the timestep sizes are always chosen such that the final simulation time is an integer multiple of the timestep size, the final simulation time is a point of superconvergence of the DG time discretizations. Hence using a polynomial basis of lower degree in the time dimension and higher degree in the spatial dimensions yields STDG methods with the same order of accuracy in both space and time with respect to the metrics used to evaluate solution errors.

The motivation for this choice is to maintain consistency between the different classes of time integration schemes, and further, by the relationship between IDC, collocation, and DG methods. The IDC and STDG methods used are chosen such that the IDC solutions converge to the STDG solutions in the limit of infinitely many correction iterations. This observation is based on the following facts: (i) IDC

methods converge to the collocation solution corresponding to the choice of quadrature nodes that is used in the limit of infinitely many correction iterations; (ii) the IDC methods considered in this work are constructed using Gauss-Radau quadratures; and (iii) the DG time integration schemes used here are equivalent to Radau collocation schemes.

3.3.2 Angular quadratures and relabeling

Slab geometry

In slab geometry the angular variable takes values $\mu \in [-1,1]$ and the ordinates (and their associated weights) are obtained from Gauss-Legendre quadrature. As was done in [61, 97], the angular reconstruction is performed at each spatial point by evaluating the Lagrange interpolant (over the angular variable) of the collided flux at each of the uncollided ordinates. This method defines a linear map from the collided quadrature set to the uncollided quadrature set that can be applied using a single matrix multiplication. This yields a reconstruction procedure requiring $2K_{\rm U}K_{\rm C}$ operations per spatial point.

xy-geometry

In two-dimensional xy-geometry the angular variable takes values $\vec{\Omega} \in \mathbb{S}^2$. It is known that optimal Gauss-type quadratures over \mathbb{S}^2 do not exist as they do in the one-dimensional case [170]. Because of this, a large number of quadratures have been developed for use in discrete ordinates calculations – c.f. [2, 6, 7, 32, 33, 71, 116, 132, 167] for some examples. The design of most angular quadratures is motivated by the desire to optimize one or more particular characteristics of the quadrature. Many quadratures emphasize the exact integration of as many spherical harmonics as possible [6, 33, 131, 132, 133]. This restriction is often accompanied by a requirement on the symmetry of the quadrature: Most quadratures are chosen such that they are symmetric with respect to octahedral reflections and rotations aligned with the coordinate axes [33, 131, 132, 133], though quadratures using other forms of symmetry such as those corresponding to icosahedral rotations have been described [6]. Other quadratures are designed around geometric considerations that permit arbitrary levels of refinement, local refinement, and a distribution of nodes over the sphere that is reasonably close to uniform [107, 129, 178, 179]

One of the focal points of this work is the use of methods that permit an arbitrary level of accuracy, and in particular, methods of arbitrarily high order. As such, the use of angular quadratures that per-

mit an arbitrary level of refinement is desirable. Unfortunately, many quadratures designed for discrete ordinates calculations have limitations that either make them ill-suited to the construction of arbitrarily accurate solvers or to the construction of high-order reconstruction procedures: For example, some quadratures require negative weights as the number of nodes is increased [128]. These limitations are a result of the fact that certain features of the quadrature must be sacrificed in order to optimize other features. We do not seek to conduct an exhaustive (or even particularly thorough) comparison of angular quadratures in this work. However, we consider two kinds of angular quadratures for multidimensional geometries for the sake of variety: CL tensor-product quadratures and the T_N quadratures of [178, 179]. Both quadrature families permit an arbitrary level of refinement in the quadrature, but the choices that are made in the design of these quadratures emphasize different characteristics in each case.

Chebyshev-Legendre (CL) quadratures

The Chebyshev-Legendre (CL) quadratures are constructed using a tensor product approach in which an N-point Gauss-Legendre quadrature is applied to the cosine of the polar angle of the sphere (i.e., the z Cartesian variable) and a 2N-point composite midpoint rule is applied to the azimuthal angle at each latitude of the sphere corresponding to one of the Gauss-Legendre nodes. The CL quadrature of order N thus contains a total of $2N^2$ quadrature nodes. A detailed definition of the quadrature is given in definition 1. The advantages of CL quadratures are that: (i) they are relatively easy to construct; (ii) they possess strictly positive quadrature weights for all orders of accuracy; and (iii) the number of spherical harmonics that are integrated exactly by the quadrature increases proportionally to the order of the quadrature: More specifically, the CL quadrature rule of order N on \mathbb{S}^2 integrates all spherical polynomials of degree less than 2N exactly [16, Theorem 5.4].

Definition 1 (Chebyshev-Legendre (CL) quadrature of order N). Let $\{\mu_1, ..., \mu_N\}$ and $\{w_1, ..., w_N\}$ denote the Nth set of Gauss-Legendre nodes and weights, respectively, over the interval [-1,1], and let

$$\phi_k = \frac{(2k-1)\pi}{2N}, \qquad (k=1,\dots,2N).$$
 (3.67)

Then the Chebyshev-Legendre quadrature rule of order N on \mathbb{S}^2 is given by

$$\left\{ \left(\vec{\Omega}_{k,\ell}, \omega_{k,\ell}\right) : k = 1, \dots, 2N \text{ and } \ell = 1, \dots, N \right\}$$
(3.68)

where

$$\vec{\Omega}_{k,\ell} = \left(\cos\phi_k\sqrt{1-\mu_\ell^2}, \sin\phi_k\sqrt{1-\mu_\ell^2}, \mu_\ell\right) \quad and \quad \omega_{k,\ell} = \frac{\pi}{N}w_\ell. \tag{3.69}$$

In contrast to the slab-geometry case, a single, straightforward interpolation approach cannot be used to define reconstruction procedures for hybrid methods in higher-dimensional settings. One of the difficulties in this regard is selecting the interpolation space: the choice of interpolation space on \mathbb{S}^2 depends strongly on the quadrature that is used, and in most cases an appropriate choice is not obvious. This is further complicated by the fact that solutions to interpolation problems on higher-dimensional manifolds such as \mathbb{S}^2 do not always exist [170], and the interpolants that do exist are often too ill-conditioned to be of practical use. For example, CL quadratures are constructed using a tensor product of one-dimensional Gauss-Legendre and composite midpoint quadratures. This construction can be used to define interpolants for CL nodes that is based on a product of one-dimensional polynomial and trigonometric interpolations. While the conditioning of both one-dimensional interpolations is satisfactory in most cases, we have found that the two-dimensional product is often inaccurate due to the commonly encountered "pole problem" [29] on the sphere.

One alternative is a technique called *hyperinterpolation* wherein a continuous reconstruction is performed using half as many spherical harmonics as nodes in the CL quadrature. The spherical harmonics are eigenfunctions of the Laplace-Beltrami operator on \mathbb{S}^2 , and form an orthonormal basis of spherical polynomials in $L^2(\mathbb{S}^2)$. Following [16], the (real) spherical harmonics can be written as

$$Y_{n,1}(\mu,\phi) = c_n \mathcal{P}_n(\mu) \qquad n \ge 0 \tag{3.70a}$$

$$Y_{n,2m}(\mu,\phi) = c_{n,m} \mathcal{P}_n^m(\mu) \cos(m\phi)$$
 $n \ge 1; m = 1,...,n$ (3.70b)

$$Y_{n,2m+1}(\mu,\phi) = c_{n,m} \mathcal{P}_n^m(\mu) \sin(m\phi)$$
 $n \ge 1; m = 1,...,n$ (3.70c)

where \mathcal{P}_n^m is the associated Legendre function and the normalization constants c_n and $c_{n,m}$ are given by

$$c_n = \sqrt{\frac{2n+1}{4\pi}}, \qquad c_{n,m} = \sqrt{\frac{2n+1}{2\pi} \frac{(n-m)!}{(n+m)!}}.$$
 (3.71)

A reconstruction procedure for hybrid methods using CL quadratures can be defined as follows: Given a collided flux solution defined on a set of CL nodes of order N_c , compute a discrete L^2 projection onto the basis of spherical harmonics of degree less than N_c using the quadrature rule (this is the hyperinterpolation step). The values of the reconstructed collided flux solution are obtained by evaluating the expansion of spherical harmonics generated by the hyperinterpolation procedure at each node of the higher-resolution quadrature. Since the CL quadrature of order N contains $2N^2$ nodes and there are only N^2 spherical harmonics of degree less than N, this hyperinterpolation method does not in general interpolate the nodal values. However, hyperinterpolation on CL nodes is stable in the uniform norm provided that the solution is sufficiently smooth [171]. Additionally, it is straightforward to show that this reconstruction procedure preserves the angular moments of the collided flux that are used as part of the hyperinterpolation step; that is, the angular moments of degree less than $N_{\rm C}$ in the reconstructed collided flux distribution are the same as those of the original collided flux distribution on the low-resolution quadrature.

As in the one-dimensional case, this reconstruction procedure for CL quadratures based on hyperinterpolation defines a linear map from the collided quadrature set to the uncollided quadrature set. This map can be decomposed into two components: one that maps from the nodal values on the collided quadrature set to the coefficients of the spherical harmonics expansion, and one that maps from the coefficients of the spherical harmonics expansion to the nodal values on the uncollided quadrature set. Because the dimension of the vector containing the coefficients of the spherical harmonics expansion is less than the dimensions of the vectors containing the nodal values on either quadrature set, this reconstruction is most efficiently applied using a sequence of two matrix multiplications, one for each component. This results in a reconstruction procedure requiring $4N_1^2N_1^2 + 4K_2^4$ operations 1.

The tensor-product construction of CL quadratures makes them easy to construct and their properties straightforward to describe, but also limits their efficiency. Since each latitudinal component of these quadratures contains the same number of nodes around the circumference of the sphere, these quadratures produce a significant degree of clustering near the poles of the sphere. This clustering renders them rather inefficient relative to other quadratures in streaming-dominated regimes, where the degree of uniformity in the distribution of quadrature nodes is one of the primary determining factors of the accuracy of the discrete ordinates method. This characteristic motivates our second choice of quadratures: the T_N quadratures [178, 179].

If the reconstruction procedure is applied as a single matrix multiplication, then $8N_{\rm u}^2N_{\rm c}^2$ operations are required: $4N_{\rm c}^2\left(N_{\rm u}^2-N_{\rm c}^2\right)$ operations more than the two-matrix approach.

T_N quadratures

The T_N quadratures are based on tessellations applied to the faces of an octahedron inscribed within the sphere that are projected onto the sphere's surface. The result is essentially a generalized midpoint rule in which the quadrature nodes consist of the projections of the centroids of the triangles onto the sphere and the quadrature weights are computed as the surface areas of the projected triangles on the sphere. Each T_N quadrature can be defined by considering a single octant and then reflecting the quadrature nodes from that octant across each of the coordinate planes to the other octants. Here, as in [178], we consider only the first octant. The tessellation is performed in the context of a reference triangle called the *basal triangle*. This triangle is embedded within a two-dimensional plane called the *basal plane* in which coordinates are denoted using an ordered pair (u, v): The vertices of the basal triangle lie at the coordinates (-1/2,0), (1/2,0), and (0,1).

A uniform tessellation is applied to the basal triangle according to the order N of the quadrature. This tessellation consists of N rows of triangles indexed n = 0, ..., N-1, where the nth row contains 2n + 1 triangles that are indexed by k = 0, ..., 2n. The coordinates of the vertices of the kth triangle in the nth row are given by

$$\left(\frac{k-n}{2N}, \frac{N-n}{N}\right), \quad \left(\frac{k-n-1}{2N}, \frac{N-n-1}{N}\right), \quad \left(\frac{k-n+1}{2N}, \frac{N-n-1}{N}\right)$$
 (3.72a)

for k even and

$$\left(\frac{k-n}{2N}, \frac{N-n-1}{N}\right), \quad \left(\frac{k-n-1}{2N}, \frac{N-n}{N}\right), \quad \left(\frac{k-n+1}{2N}, \frac{N-n}{N}\right)$$
 (3.72b)

for k odd. The coordinates (u^*, v^*) of the centroid of each triangle in the tessellation of the basal triangle are given by

$$u^* = \frac{1}{3} (u_1 + u_2 + u_3), \qquad v^* = \frac{1}{3} (v_1 + v_2 + v_3), \tag{3.73}$$

where $\{(u_i, v_i)\}_{i=1}^3$ are the coordinates of the vertices of the triangle.

The projection from the basal triangle to the surface of the sphere follows a two-step process. The first step maps each ordered pair $(u, v) \in \mathbb{R}^2$ in the basal plane to an ordered triple $(x, y, z) \in \mathbb{R}^3$ for the corresponding point on the face of the octahedron centered at the origin with vertices at $(\pm 1, 0, 0)$, $(0, \pm 1, 0)$, and $(0, 0, \pm 1)$. This map is defined explicitly by

$$x = \frac{1}{2} - u - \frac{1}{2}v,$$
 $y = \frac{1}{2} + u - \frac{1}{2}v,$ $z = v.$ (3.74)

The second step projects each point $(x, y, z) \in \mathbb{R}^3$ on the octahedron to a point $(\xi, \eta, \mu) \in \mathbb{S}^2$. This projection is defined explicitly by

$$\xi = \frac{x}{\sqrt{x^2 + y^2 + z^2}}, \qquad \eta = \frac{y}{\sqrt{x^2 + y^2 + z^2}}, \qquad \mu = \frac{z}{\sqrt{x^2 + y^2 + z^2}}. \tag{3.75}$$

The quadrature weights are computed from the surface area on the sphere of the projection of each triangle in the tessellation, which can be written explicitly using results from elementary trigonometry. Let $\vec{\alpha}, \vec{\beta}, \vec{\gamma} \in \mathbb{S}^2$ denote the vectors containing the coordinates of the vertices of a projected triangle. The arc lengths of each edge of the projected triangle are given by

$$a = \cos^{-1}(\vec{\beta} \cdot \vec{\gamma}), \qquad b = \cos^{-1}(\vec{\alpha} \cdot \vec{\gamma}), \qquad c = \cos^{-1}(\vec{\alpha} \cdot \vec{\beta}), \qquad (3.76)$$

and the semi-perimeter of the projected triangle is given by

$$s = \frac{1}{2} (a+b+c). {(3.77)}$$

The surface area of the triangle is given by

$$A + B + C - \pi \tag{3.78}$$

where

$$A = 2\tan^{-1}\left(\frac{\tan r}{\sin(s-a)}\right), \qquad B = 2\tan^{-1}\left(\frac{\tan r}{\sin(s-b)}\right), \qquad C = 2\tan^{-1}\left(\frac{\tan r}{\sin(s-c)}\right), \tag{3.79}$$

and

$$\tan r = \sqrt{\frac{\sin(s-a)\sin(s-b)\sin(s-c)}{\sin s}}.$$
(3.80)

A reconstruction procedure for hybrid methods using T_N quadratures can be defined based on a piecewise-constant interpolation over the tessellation used to construct the low-resolution quadrature. Evaluating this interpolant at each of the high-resolution quadrature nodes yields the reconstructed values of the collided flux. One advantage of this reconstruction procedure is that it is very inexpensive to apply: a simple copy operation is all that is required once the mapping between nodes has been determined. When the order $N_{\rm U}$ of the high-resolution quadrature is an integer multiple of the order $N_{\rm C}$ of the low-resolution quadrature, the tessellation on which the high-resolution quadrature is based is a refinement of the tessellation on which the low-resolution quadrature is based. In this case, determining the mapping between nodes is reasonably straightforward, and, due to the way the quadrature

weights are computed, the resulting reconstruction procedure is guaranteed to preserve the scalar flux of the solution. In this work we consider only hybrid methods using T_N quadratures that differ by integer multiples.

At this point, we would like to emphasize that the properties of the reconstruction procedures defined for the CL and T_N quadratures are quite different. Whereas the reconstruction for CL quadratures preserves as many angular moments as are used in the hyperinterpolation step, the reconstruction for T_N quadratures preserves only the zeroth angular moment of the solution exactly. In particular, the T_N reconstruction does not preserve the solution's current (first angular moments), which may impact the accuracy of solutions near material interfaces [74]. Because of the nonlinearities found in thermal radiative transfer problems, it is often necessary to enforce positivity in the scalar flux distribution. While the T_N reconstruction will preserve positivity in reconstructed solutions, the CL reconstruction can introduce negative values due to oscillations in the spherical harmonics expansion, even if the original solution is non-negative at each low-resolution quadrature node.

3.3.3 Spatial discretization

Each high-order time discretization is paired with an arbitrarily high-order discontinuous Galerkin spatial discretization, which is constructed in essentially the same manner as the STDG methods described in section 3.2.2, but without the components in the time dimension. The discontinuous Galerkin (DG) method was first proposed by Reed and Hill for approximation of spatial variables in steady-state neutron transport problems [166]. Formal convergence results were later established by Lesaint and Raviart [125] and Johnson and Pitkäranta [108]. The method has subsequently been generalized to time-dependent first-order systems of hyperbolic conservation laws by Cockburn and Shu, et al. [53, 55, 57, 58, 59]. Extensions to second-order elliptic and parabolic equations [12, 13, 22, 30, 36, 54, 56, 68] and other higher-order equations [42] have also been investigated.

One of the primary advantages of the DG method is the lack of strict continuity enforcement across boundaries of the finite element mesh. This permits the construction of finite element bases in which the support of each basis element is contained in a single cell. Combining such localized bases with "upwind" numerical fluxes (which preserve the characteristic dependence of the continuous operator \mathcal{L}) allows efficient sweep algorithms of arbitrarily high order to be constructed.

For all of the Runge-Kutta integrators (i.e., the Euler method, the DIRK methods, and the IDC and LS-IDC methods) solving the discrete time-dependent system for each stage value is equivalent to solving a steady-state system with a specific form. Hence we consider a steady-state version of the discrete ordinates system (3.5), which an be written in the form

$$(\mathcal{L} - \mathcal{SP}) \psi = q. \tag{3.81}$$

Assume the spatial domain X is a convex polyhedron and let \mathcal{T}_h be a partition of X into nonoverlapping polyhedral cells \mathcal{K} of diameter less than h. For each ordinate $\vec{\Omega}_k$, the boundary of the spatial domain can be decomposed into two components $\partial X = \partial X_k^- \cup \partial X_k^+$ where

$$\partial X_k^- = \left\{ \vec{x} \in \partial X : \vec{n}(\vec{x}) \cdot \vec{\Omega}_k < 0 \right\} \quad \text{and} \quad \partial X_k^+ = \left\{ \vec{x} \in \partial X : \vec{n}(\vec{x}) \cdot \vec{\Omega}_k \ge 0 \right\}$$
 (3.82)

are the inflow and outflow components of the boundary for the kth ordinate, respectively. Let \mathcal{E}_h denote the set of interior edges of the partition \mathcal{T}_h with $\left(\bigcup_{e\in\mathcal{E}_h}e\right)\cap\partial X=\emptyset$, and to each edge $e\in\mathcal{E}_h$ associate a unit vector \vec{n}_e normal to the edge whose orientation may be chosen arbitrarily. On each edge $e\in\partial X_k^\pm$ the associated unit normal vector \vec{n}_e is specified to be the outward normal.

For each ordinate, the DG approximation of the solution ψ_k of the steady-state discrete ordinates system (3.81) is given by $\psi_k^h \in V_h$ satisfying

$$\mathcal{B}_{k}^{h}\left(\psi_{k}^{h}, v^{h}\right) = \bar{\mathcal{B}}^{h}\left(\phi^{h}, v^{h}\right) + \mathcal{A}_{k}^{h}\left(v^{h}\right), \qquad (k = 1, \dots, K)$$

$$(3.83)$$

for every $v^h \in V_h$ with the bilinear and linear forms

$$\begin{split} \mathcal{B}_{k}^{h}(u,v) &= -\sum_{\mathcal{K} \in \mathcal{T}_{h}} \int_{\mathcal{K}} \left(\vec{\Omega}_{k} \cdot \nabla_{\vec{x}} v + \sigma_{\mathsf{t}} v \right) u \, d\vec{x} - \sum_{e \in \mathcal{E}_{h}} \int_{e} \left(\vec{\Omega}_{k} \cdot \vec{n}_{e} \right) [v] \, \widehat{u} \, ds \\ &+ \sum_{e \in \partial X_{k}^{+}} \int_{e} \left(\vec{\Omega}_{k} \cdot \vec{n}_{e} \right) v^{-} u^{-} \, ds, \end{split} \tag{3.84a}$$

$$\bar{\mathcal{B}}^{h}(u,v) = \frac{1}{4\pi} \sum_{\mathcal{K} \in \mathcal{T}_{h}} \int_{\mathcal{K}} \sigma_{\mathbf{S}} v u \, d\vec{x}, \tag{3.84b}$$

$$\mathcal{A}_{k}^{h}(v) = \sum_{\mathcal{K} \in \mathcal{T}_{h}} \int_{\mathcal{K}} q_{k} v \, d\vec{x} - \sum_{e \in \partial X_{k}^{-}} \int_{e} \left(\vec{\Omega}_{k} \cdot \vec{n}_{e}\right) \psi_{\mathrm{B},k} v^{-} \, ds, \tag{3.84c}$$

where

$$\boldsymbol{\psi}^h = \begin{bmatrix} \psi_1^h, \dots, \psi_K^h \end{bmatrix}^T, \qquad \phi^h = \mathcal{P}\boldsymbol{\psi}^h, \tag{3.85}$$

 $\psi_{\mathrm{B},k} = \psi_{\mathrm{B}}(\vec{x},\vec{\Omega}_k)$ is the value of the inflow condition specified on ∂X_k^- for the ordinate $\vec{\Omega}_k$,

$$v^{\pm}(\vec{x}) = \lim_{\varepsilon \to 0^{+}} v(\vec{x} \pm \varepsilon \vec{n}_{e}), \qquad [v] = v^{+} - v^{-}, \tag{3.86}$$

and \hat{v} is a numerical flux specified in terms of v^+ and v^- . In order to maintain the direction of information flow imposed by the continuous operator \mathcal{L} , we use the *upwind flux*

$$\widehat{v} = \lim_{\varepsilon \to 0^+} v \left(\vec{x} - \varepsilon \vec{\Omega}_k \right). \tag{3.87}$$

In practice (3.83) is written as a linear system relative to some basis of V_h . If B_h is a vector whose elements form a basis for V_h and ψ_k^h be the vector of coefficients such that $\psi_k^h = B_h^T \Psi_k^h$, then (3.83) can be written as

$$\mathfrak{L}^h \mathbf{\Psi}^h = \mathfrak{S}^h \mathfrak{P}^h \mathbf{\Psi}^h + \mathbf{Q}^h, \tag{3.88}$$

where

$$\left(\mathbf{\Psi}^{h}\right)^{T} = \left[\left(\mathbf{\Psi}_{1}^{h}\right)^{T}, \dots, \left(\mathbf{\Psi}_{K}^{h}\right)^{T}\right] \quad \text{and} \quad \left(\mathbf{Q}^{h}\right)^{T} = \left[\left(Q_{1}^{h}\right)^{T}, \dots, \left(Q_{K}^{h}\right)^{T}\right] \quad (3.89)$$

with $Q_k^h = \mathcal{A}_k^h(B_h)$: the operator \mathfrak{L}^h has the form

$$\mathfrak{L}^{h} = \operatorname{Diag}\left(\mathfrak{L}_{1}^{h}, \dots, \mathfrak{L}_{K}^{h}\right) \quad \text{with} \quad \mathfrak{L}_{k}^{h} = \mathcal{B}_{k}^{h}\left(B_{h}^{T}, B_{h}\right)$$
(3.90)

for k = 1,...,K; and the operators \mathfrak{S}^h and \mathfrak{P}^h are defined by

$$\mathfrak{S}^{h} = \mathbb{1}_{K} \otimes \bar{\mathfrak{S}}^{h} \quad \text{and} \quad \mathfrak{P}^{h} \Psi^{h} = \sum_{k=1}^{K} \omega_{k} \Psi_{k}^{h}, \tag{3.91}$$

where

$$\tilde{\mathfrak{S}}^h = \bar{\mathcal{B}}^h \left(B_h^T, B_h \right) \tag{3.92}$$

and $\mathbb{1}_K$ is the *K*-dimensional column vector of ones.

For most of the numerical tests considered in this work, each time integrator is paired with a DG spatial discretization of the same order of accuracy with respect to space as the order of the time integrator: All exceptions are explicitly noted. We consider only uniform Cartesian meshes where each spatial cell is assigned a tensor product basis of polynomials of degree at most one less than the order of the time integration scheme.

3.4 Numerical tests

Two sets of test problems are considered to assess the accuracy and efficiency of the hybrid methods. The first set consists of a sequence of convergence tests for smooth problems in one-dimensional slab geometry under the diffusive scaling (1.4). The aim of the second set of tests is to compare the computational efficiency (i.e., error versus solve time) of hybrid and standard discrete ordinates methods in two-dimensional xy-geometry. Throughout this section S_N will be used to denote a standard discrete ordinates approximation using an angular quadrature of order N with K total ordinates, and $S_{N_{\rm U},N_{\rm C}}$ will be used to denote a hybrid discrete ordinates approximation using angular quadratures of order $N_{\rm U}$ and $N_{\rm C}$ with $K_{\rm U}$ and $K_{\rm C}$ total ordinates in the uncollided and collided components, respectively.

3.4.1 One spatial dimension

We first examine the convergence behavior of the hybrid discrete ordinates methods in one-dimensional slab geometry using Gauss-Legendre quadratures. With $\sigma_s = \sigma_t$, q = 0, and under the asymptotic scaling (1.4), the slab-geometry form of the transport equation (1.2) can be written as

$$\varepsilon \partial_t \psi + \mu \partial_z \psi + \frac{\sigma_t}{\varepsilon} \psi = \frac{\sigma_t}{2\varepsilon} \langle \psi \rangle. \tag{3.93}$$

A sequence of convergence tests are performed for decreasing ε on a spatial domain of $z \in [-4,4]$ with vacuum boundary conditions (zero inflow) and $\sigma_t \equiv 1$ for all z. Convergence tests are conducted for $\varepsilon = 1\text{E}+0$, 1E-1, 1E-2, and 1E-3. The simulations are run to a final time of t=0.5 and refinement of the space-time mesh is performed by successively doubling the number of spatial cells from 2^7 to 2^{17} under a constant CFL of 8. The initial condition is set from the isotropic distribution $\psi(z,\mu,0) = \eta_T(z)/2$, where η_T is a compactly supported C^∞ function defined by

$$\eta_{r}(z) = \begin{cases}
C \exp\left[\left(\frac{|z|^{2}}{r^{2}} - 1\right)^{-1}\right], & |z|^{2} < r^{2} \\
0, & \text{otherwise.}
\end{cases}$$
(3.94)

Here r is the radius of the support of η_r and C > 0 is a constant such that $||\eta_r||_{L^1} = 1$. Each convergence test is performed with a radius of r = 1. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} standard discrete ordinates angular approximation.

The results of the convergence tests for DIRK integrators are shown in figures 3.3 and 3.4, for STDG integrators in figure 3.5, and for IDC5 and LS-IDC5 integrators in figure 3.6. It is clear from these tests that the use of very coarse angular approximations in the collided component of the hybrid methods can significantly degrade the overall accuracy of the methods. We observe two distinct phenomena that contribute to the error profiles of the hybrid methods: convergence saturation and order reduction. These effects appear to be governed by a combination of (i) the resolution of the angular approximation that is used in the collided component, (ii) the timestep size of the integrator, and (iii) the particle mean-free-time.

When the particle mean-free-time is smaller than the timestep size, the convergence of the hybrid methods saturates at the error level of the low-resolution quadrature that is used in the collided flux. This is a result of the fact that the collided component dominates in highly collisional regimes where the vast majority of particles are collisional within each timestep. As the mesh is refined the timestep size eventually becomes small enough to resolve the particle mean-free-time and the solution begins to shift from the collided component into the uncollided component. This shift causes the hybrid methods to resume converging to the reference solution, but at a rate corresponding to the order of the low-resolution angular quadrature. Order reduction is therefore observed when this rate is less than the convergence rate of the time-integration scheme. This order reduction is unexpected, and a rigorous study of this phenomenon is currently underway. Preliminary results suggest that this phenomenon is a result of a truncation error with respect to the length of time between applications of the reconstruction procedure for mapping the collided flux into the uncollided flux, which in this case corresponds to the timestep size.

We note that there are two reasons for convergence saturation of the hybrid methods in these tests: (i) the conditioning of the linear systems that must be solved at each stage of the time integrators, and (ii) the accuracy of the angular approximation used in the collided component. As $\varepsilon \to 0$ the linear systems that must be solved at each stage become increasingly ill-conditioned, and hence the minimum error that one can expect to realize when solving such systems increases proportionally. Roughly speaking, the loss of accuracy in these tests due to the conditioning of the linear systems that must be solved is generally within one order of magnitude of $1/\varepsilon$. This yields minimum effective tolerances of approximately $1\varepsilon - 13$, $1\varepsilon - 12$, $1\varepsilon - 11$, and $1\varepsilon - 10$, depending on the time integrator, for the tests considered here in order of decreasing ε . We emphasize that any behavior observed in these convergence tests below

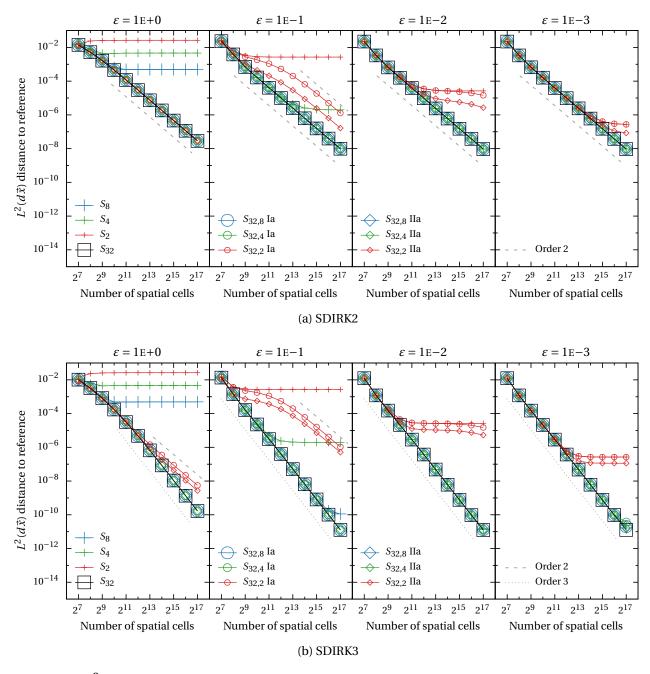


Figure 3.3: L^2 convergence studies of hybrid-Ia and hybrid-IIa methods based on SDIRK2 and SDIRK3 integrators for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

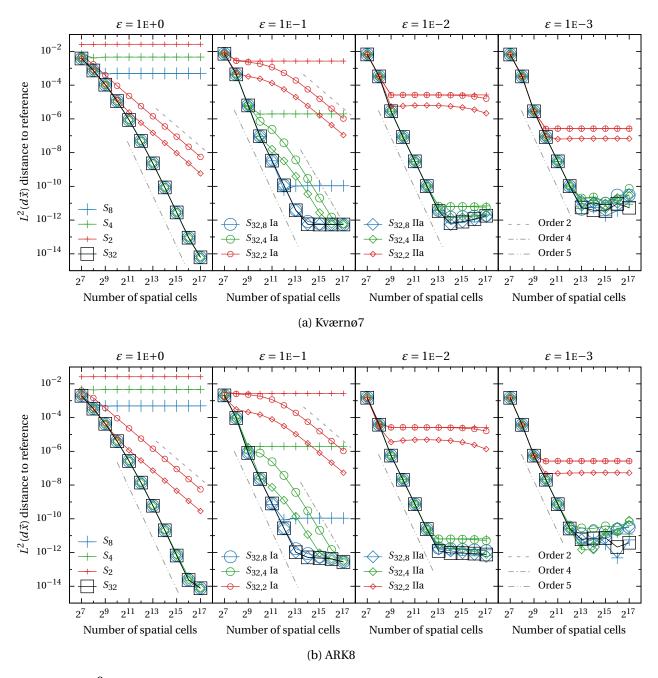


Figure 3.4: L^2 convergence studies of hybrid-Ia and hybrid-IIa methods based on Kværnø7 and ARK8 integrators for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

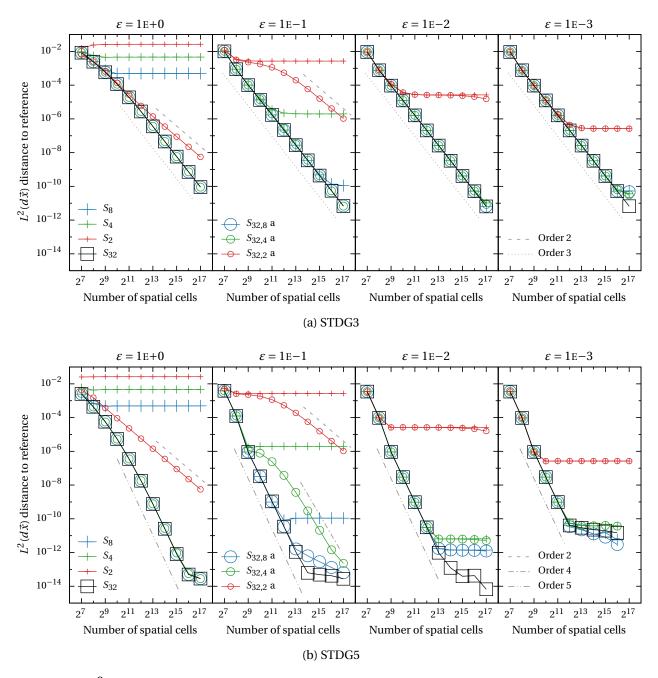


Figure 3.5: L^2 convergence studies of hybrid-a methods based on STDG integrators for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

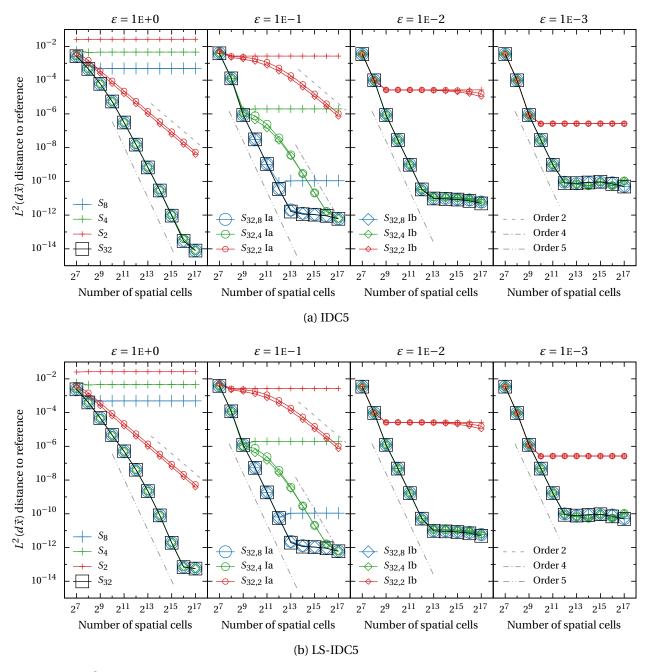


Figure 3.6: L^2 convergence studies of hybrid-I methods based on IDC5 and LS-IDC5 integrators for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

these tolerances should be considered suspect.

Even though convergence saturation and order reduction can be observed, the accuracy of many of the hybrid schemes is generally quite good. The $S_{32,2}$ methods display marked increases in error over the reference S_{32} approximations (in the form of saturation, order reduction, or both) in nearly every case, with more significant deviations in accuracy observed for higher-order integrators and moderate values of ε . However the $S_{32,4}$ methods only show significant increases in error for fifth-order integrators and the case of $\varepsilon = 1E-1$.

The behavior of the $S_{32,2}$ and $S_{32,4}$ solutions suggests that the splitting error of the hybrid-IIa DIRK methods is smaller than that of the hybrid-Ia DIRK methods. However, we make two observations concerning this phenomenon. First, the difference in accuracy appears to depend strongly on the time integration scheme that is applied: The difference in error between the hybrid-Ia and hybrid-IIa methods is nearly two orders of magnitude for ARK8 and approximately one order of magnitude for SDIRK2, but the difference for SDIRK3 is much smaller. Second, this difference in accuracy amounts to a reduction in error by a constant factor: All of the hybrid DIRK methods display similar levels of convergence saturation and order reduction below the saturation point, but at marginally different tolerances. At this time it is not clear what properties of the DIRK schemes might affect this behavior.

A similar difference in the splitting error is observed in the hybrid-Ia and hybrid-Ib methods using IDC and LS-IDC integrators. In this case the methods display similar saturation behavior and the difference in error appears only in the order-reduction regime. Additionally, the difference in accuracy between the hybrid-Ia and hybrid-Ib IDC and LS-IDC methods in each test is very small. The fact that the hybrid-Ia methods apply the reconstruction procedure only once per timestep suggests that they may offer an advantage over the hybrid-Ib methods if the reconstruction procedure that is used is very expensive to apply. On the other hand, the increased frequency with which the hybrid-Ib methods apply the reconstruction procedure may be more advantageous for problems in multidimensional geometries that contain complex material interfaces.

3.4.2 Two spatial dimensions

To compare the efficiency of the hybrid discrete ordinates methods described in section 3.2 with classical discrete ordinates methods we consider the following test problem in two-dimensional *xy*-geometry,

which we refer to as *the hohlraum problem*. The problem is a simplified mock-up of standard hohlraum configurations used in nuclear fusion devices. The configuration used here is a modification of that originally introduced in [31], where models of thermal radiative transfer were considered that couple the transport equation (1.2) to an energy equation for the background material. Since this type of coupling is not considered in this work, we use a modified configuration in which particle scattering is used to emulate the effects of thermal absorption and re-emission of particles. The layout of the problem is shown in figure 3.7. We note that the cross-section values of the material are not rigorously selected to reproduce the effects of the material energy coupling, but instead are chosen such that the behavior of the system is qualitatively similar to that of the coupled problem for the time scales selected for these tests. The problem is simulated to a final time of 2.6 using a CFL of 104 using hybrid and standard discrete ordinates methods with various angular resolutions. The number of cells in the spatial mesh is selected based on the order of the method: second-order methods use 312 cells per dimension, third-order methods use 208 cells per dimension, and fifth-order methods use 104 cells per dimension.

An approximation of the error in each solution is obtained by comparing the scalar flux distribution of each solution at the final time to a particular reference solution. The reference solution that is used for each approximate solution is computed using the same space-time discretization as the approximate solution and a highly-refined standard discrete ordinates angular approximation. Since each approximate solution is compared to a reference solution using the same space-time discretization, the computed error approximation contains only the angular component of the overall solution error. In particular, this approach ignores errors in the solution due to approximations in space and time. As a result of this, we emphasize that these tests cannot be used to make efficiency comparisons between different space-time discretizations themselves because these components of the error are intentionally suppressed.

We present results for both multidimensional angular quadratures described in section 3.3.2; that is, we present results for CL quadratures using reconstruction procedures based on hyperinterpolation and for T_N quadratures using reconstruction procedures based on piecewise-constant interpolation. Reference solutions for methods using CL quadratures are computed using a CL quadrature of order 134 and reference solutions for methods using T_N quadratures are computed using a T_N quadrature of order 67. The order of these quadratures is chosen based on prime factors to avoid the possibility of aliasing in the quadrature directions between approximate and reference solutions from impacting the error estimates for each approximate solution.

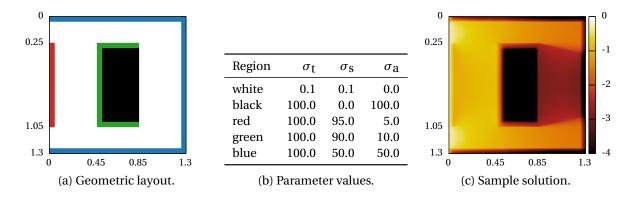


Figure 3.7: Layout and sample solution for the hohlraum test problem. Color scale of solution is log_{10} .

We present results for the hohlraum test problem in three forms. First, we show plots comparing the scalar flux density and the error distribution between hybrid and standard discrete ordinates methods using a selection of angular quadrature resolutions. For CL quadratures, standard discrete ordinates methods using an angular quadrature of order 32 are compared with hybrid discrete ordinates methods using an angular quadrature of order 64 in the uncollided component and angular quadratures of various resolutions in the collided component. For T_N quadratures, standard discrete ordinates methods using an angular quadrature of order 16 are compared with hybrid discrete ordinates methods using an angular quadrature of order 32 in the uncollided component and angular quadratures of various resolutions in the collided component. Results for DIRK methods are shown in figures 3.8 to 3.15, results for STDG methods are shown in figures 3.16 and 3.17, results for IDC methods are shown in figures 3.18 to 3.21, and results for LS-IDC methods are shown in figures 3.22 to 3.25. The solve time (in seconds) and an approximation of the L^∞ error are included for each solution.

We observe that the hybrid methods require significantly smaller solve times than the standard discrete ordinates methods that they are compared to. The $S_{64,2}$ and $S_{32,1}$ approximations require between one-fourth and one-fifth of the solve time of the S_{32} and S_{16} approximations, respectively, and the $S_{64,16}$ and $S_{32,8}$ approximations require between one-half and one-third of the solve time of the S_{32} and S_{16} approximations, respectively. However this reduction in solve time generally comes at the expense of accuracy: Using very coarse angular approximations in the collided component of the hybrid methods increases solution error by as much as an order of magnitude or more, depending on the quadrature. An examination of the plots of the scalar flux distribution of each solution shows that the increased error is typically accompanied by an increased presence of ray effects in the solution. The fact that these

ray effects are more pronounced in the solutions of higher-order methods is a result of the fact that each test is run with the same CFL number, so that higher-order methods use a larger timestep size than lower-order methods. This illustrates the fact that even though the use of a reconstruction procedure between timesteps to map the collided flux to the higher-resolution quadrature improves the accuracy of the collision-based splitting [97], increasing the timestep size significantly can decrease the effectiveness of this approach.

For CL quadratures, the $S_{64,16}$ hybrid discrete ordinates methods tend to yield an L^{∞} error that is as much as half that of the corresponding S_{32} standard discrete ordinates method in approximately half the time. For T_N quadratures, the $S_{32,8}$ hybrid discrete ordinates methods tend to yield an L^{∞} error that is comparable to that of the corresponding S_{16} standard discrete ordinates method in approximately half the time. Overall, there is a pronounced qualitative difference between the S_{32} and S_{16} discrete ordinates solutions and the $S_{64,16}$ and $S_{32,8}$ hybrid discrete ordinates solutions for each time integration scheme. The $S_{64,16}$ and $S_{32,8}$ methods tend to produce comparable or smaller errors that are more evenly distributed over the spatial domain. The difference between the hybrid and non-hybrid solutions is most apparent in regions with low collision probabilities. By using a higher-resolution quadrature in the uncollided component than the standard discrete ordinates methods to which they are compared, the hybrid methods that satisfactorily balance the quadrature resolutions that are applied to each component tend to produce solutions in which ray effects are less pronounced.

For DIRK integrators, the hybrid-IIa methods produce significantly more accurate solutions than the hybrid-Ia methods. In many cases the hybrid-IIa methods yield errors that are approximately half that of the corresponding hybrid-Ia methods. However, similar behavior is observed here as was the case for the convergence tests in the prior section: The level of improvement offered by the hybrid-IIa methods over the hybrid-Ia methods is not consistent across the different DIRK schemes. It is interesting to note that for the ARK8 scheme, the hybrid-IIa $S_{32,4}$ and $S_{32,8}$ methods using T_N quadratures yield larger L^{∞} errors than the corresponding hybrid-Ia methods, even though the hybrid-IIa solutions show clear reductions in error over most of the spatial domain. Again, it is not clear at this time what properties of the methods may impact the expression of this phenomenon.

For IDC and LS-IDC integrators, the hybrid-Ib methods yield solutions with errors that are 20 to 40 percent smaller than the corresponding hybrid-Ia methods while increasing solve time by at most a few percent. We therefore see no reason that the hybrid-Ib methods should not be preferred over the hybrid-

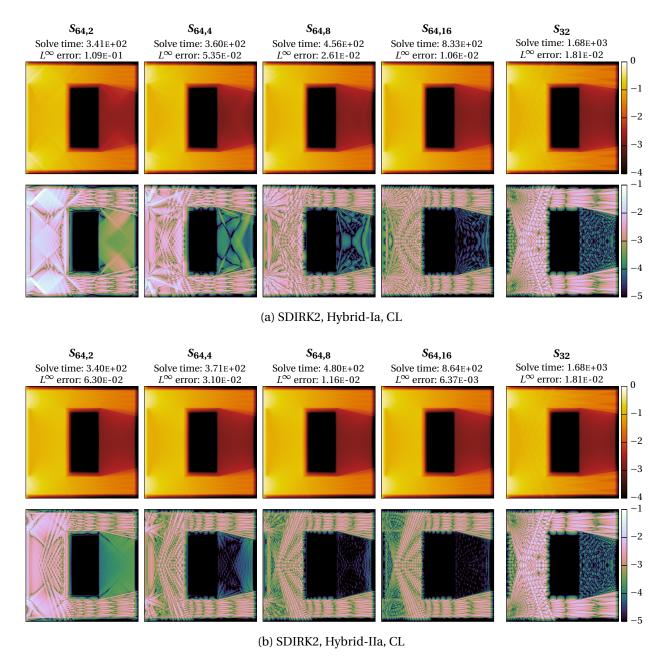


Figure 3.8: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK2 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

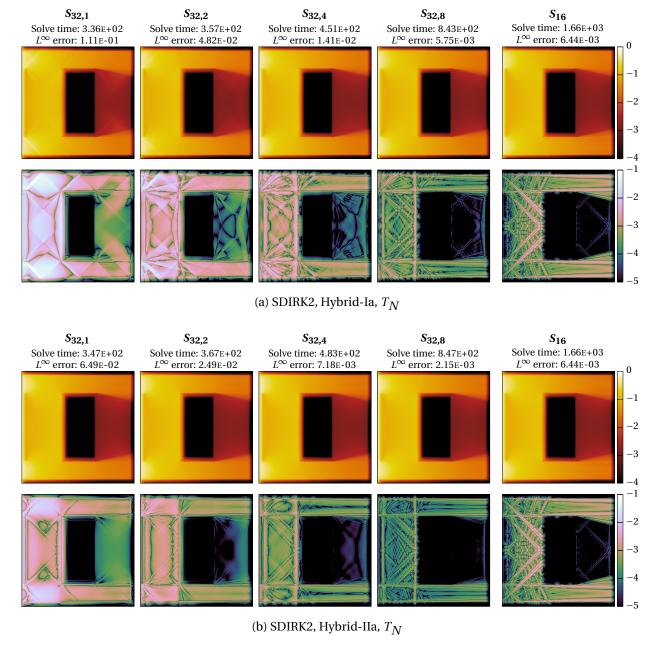


Figure 3.9: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK2 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

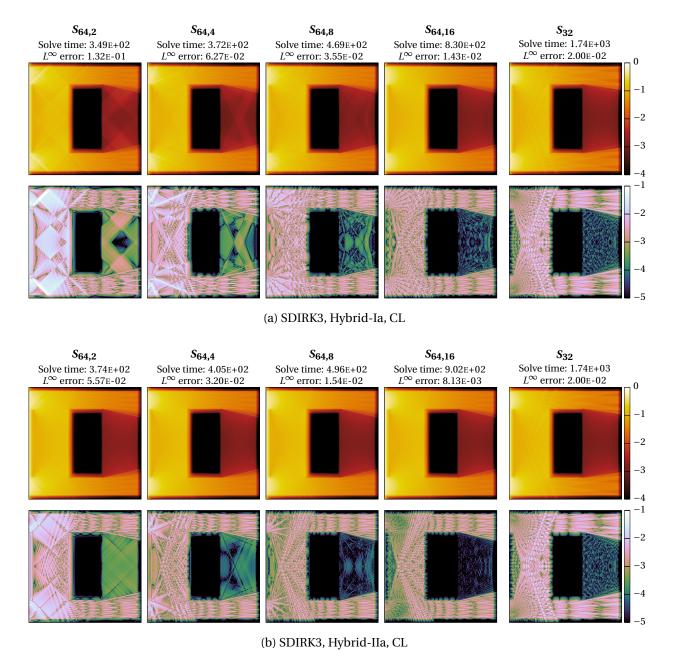


Figure 3.10: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK3 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

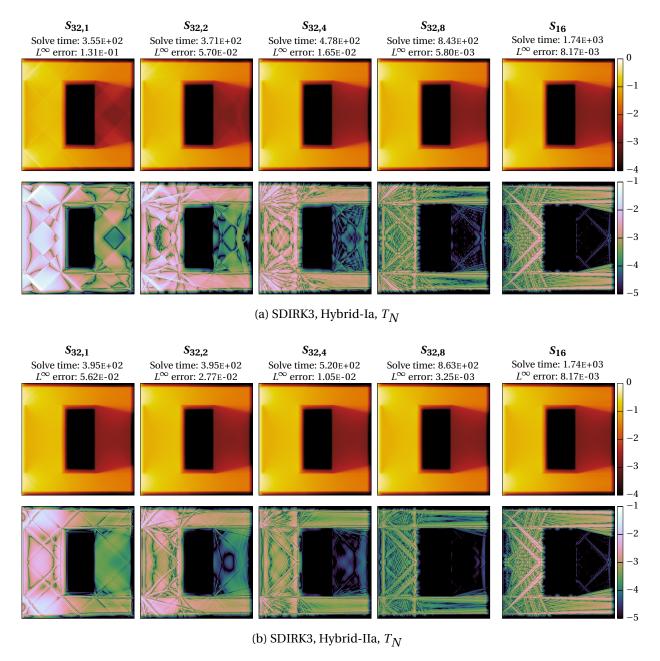
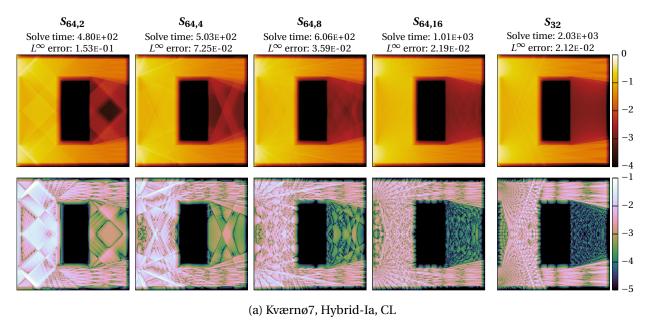
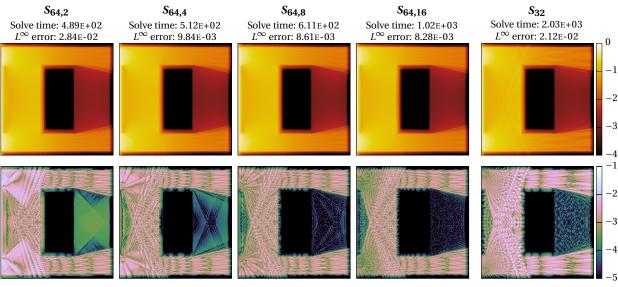


Figure 3.11: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK3 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.





(b) Kværnø7, Hybrid-IIa, CL

Figure 3.12: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid Kværnø7 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

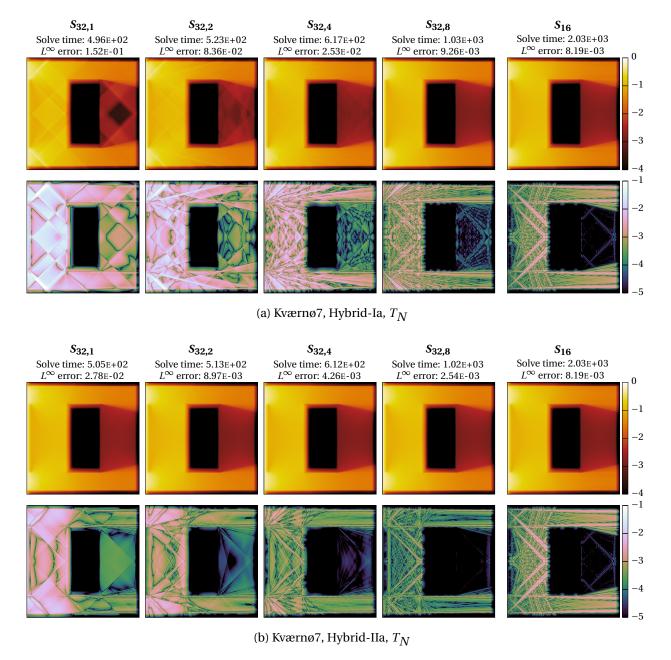


Figure 3.13: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid Kværnø7 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution.

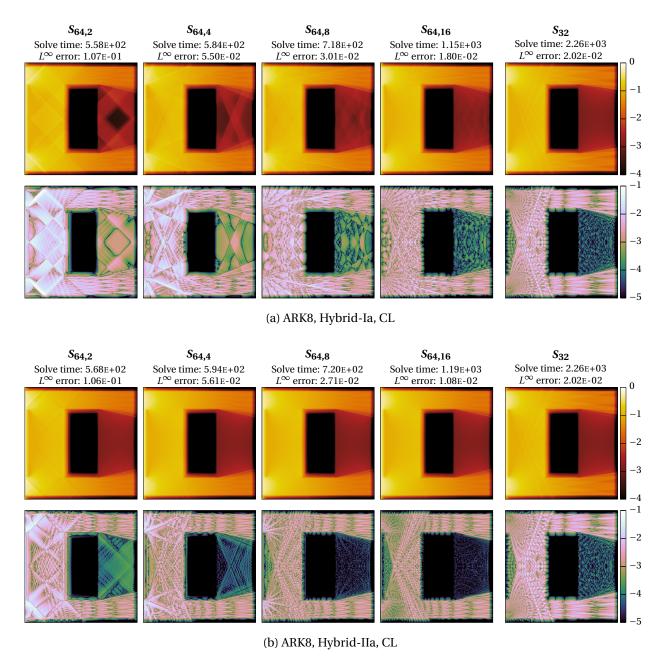


Figure 3.14: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid ARK8 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

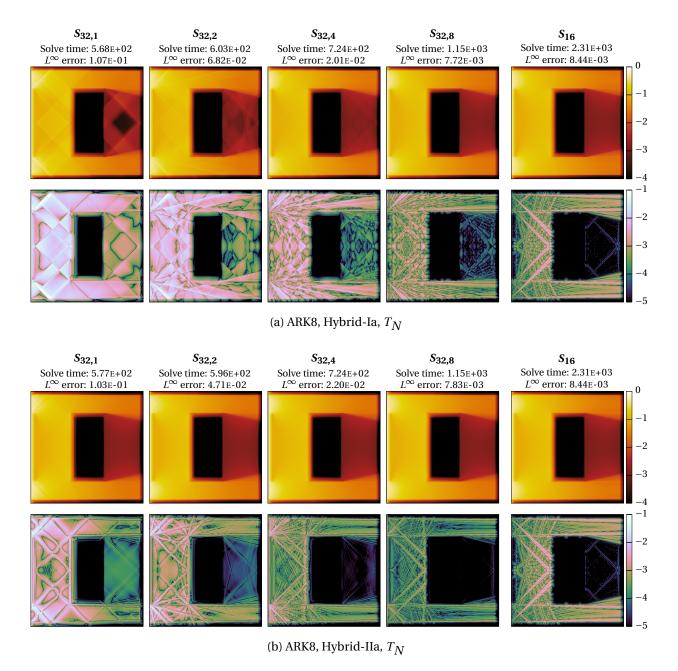


Figure 3.15: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid ARK8 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution.

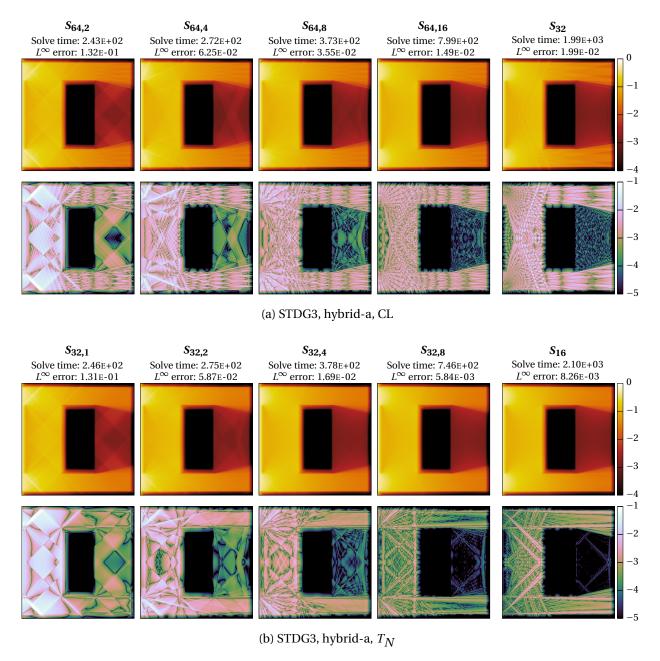


Figure 3.16: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid STDG3 methods using CL and T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

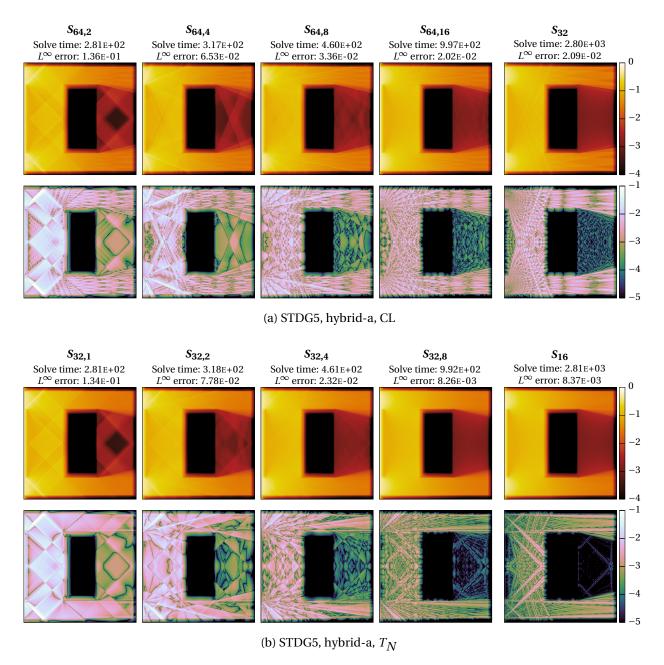


Figure 3.17: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid STDG5 methods using CL and T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

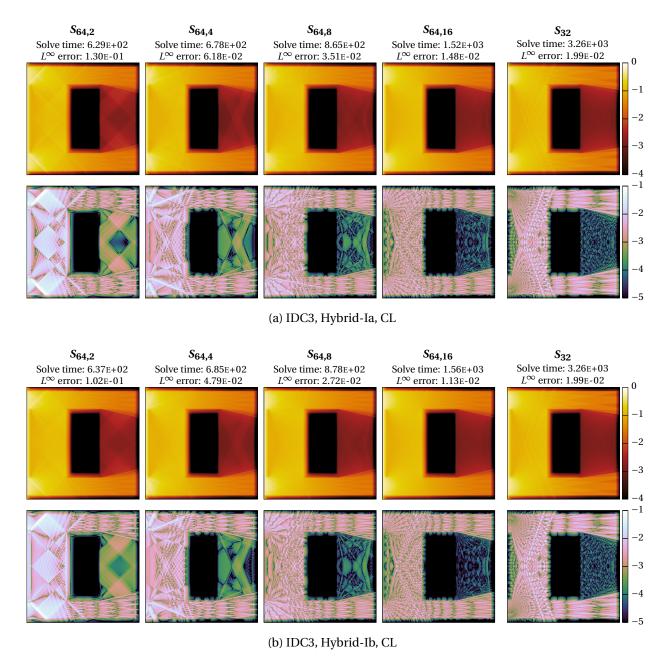


Figure 3.18: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid IDC3 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

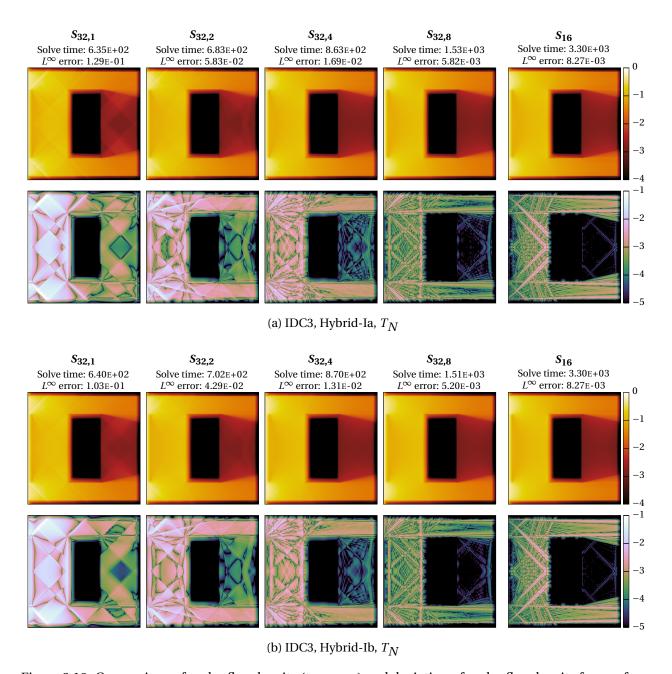


Figure 3.19: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid IDC3 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

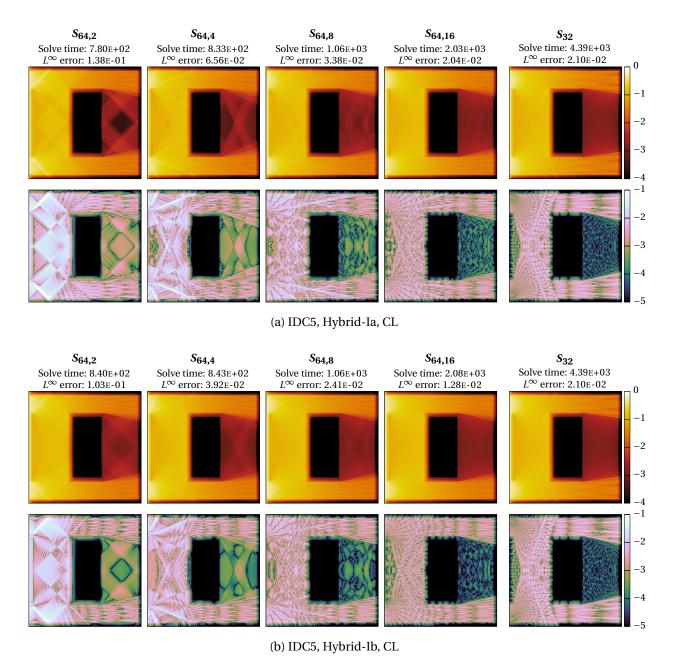


Figure 3.20: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid IDC5 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

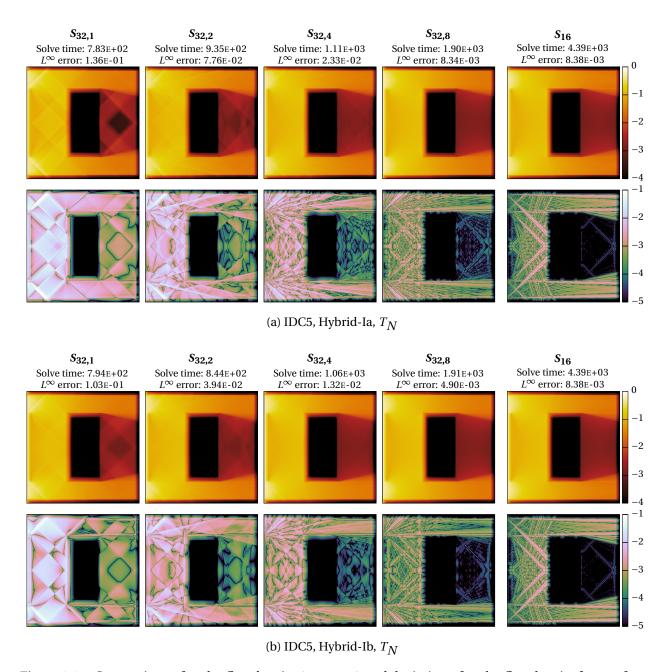


Figure 3.21: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid IDC5 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

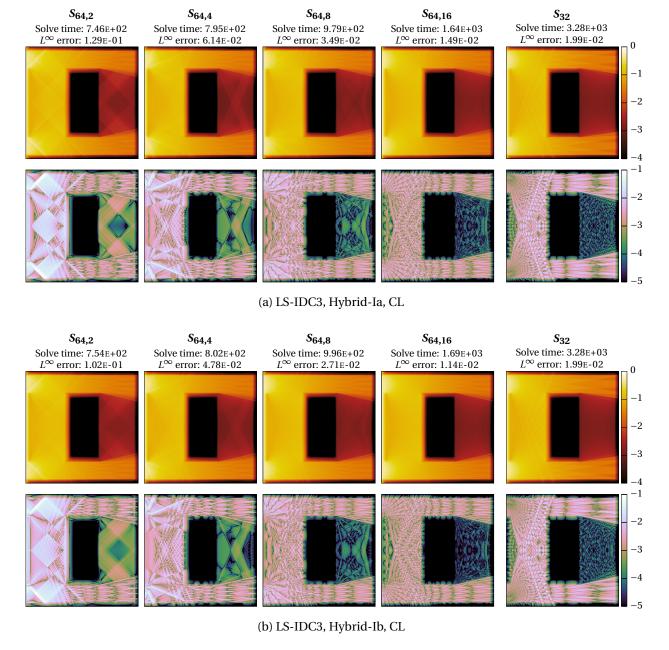


Figure 3.22: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid LS-IDC3 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

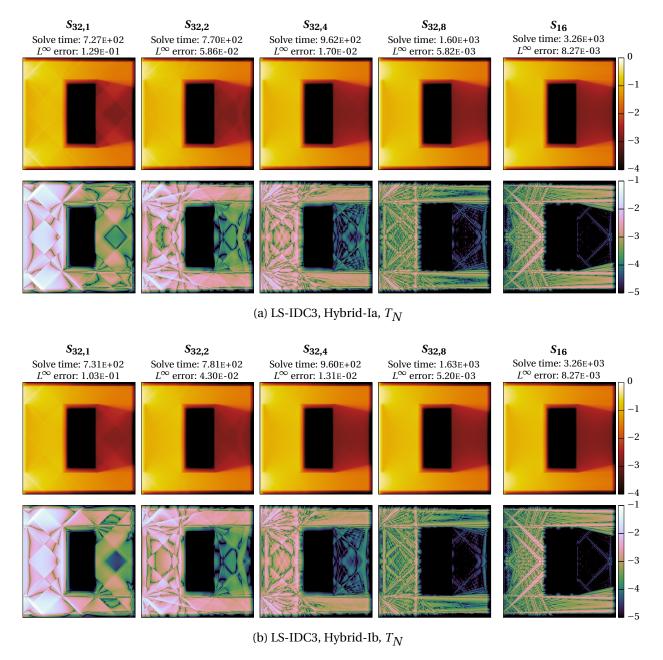


Figure 3.23: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid LS-IDC3 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution.

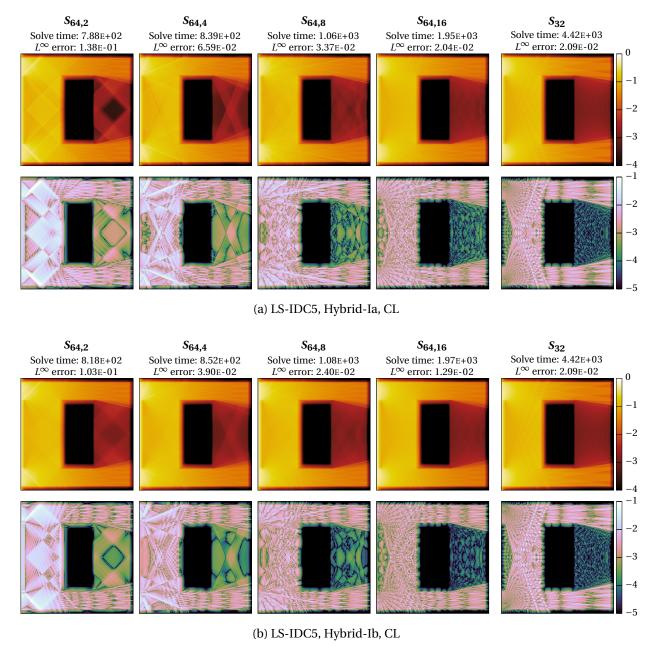


Figure 3.24: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid LS-IDC5 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

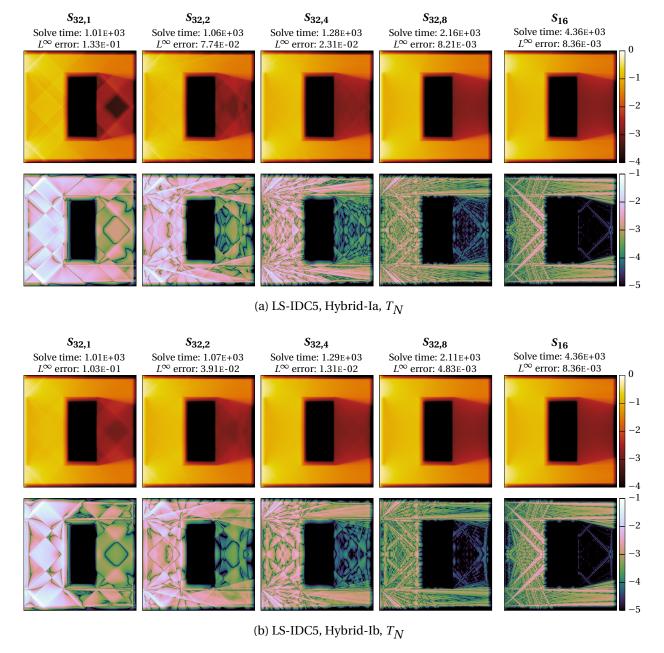


Figure 3.25: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid LS-IDC5 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

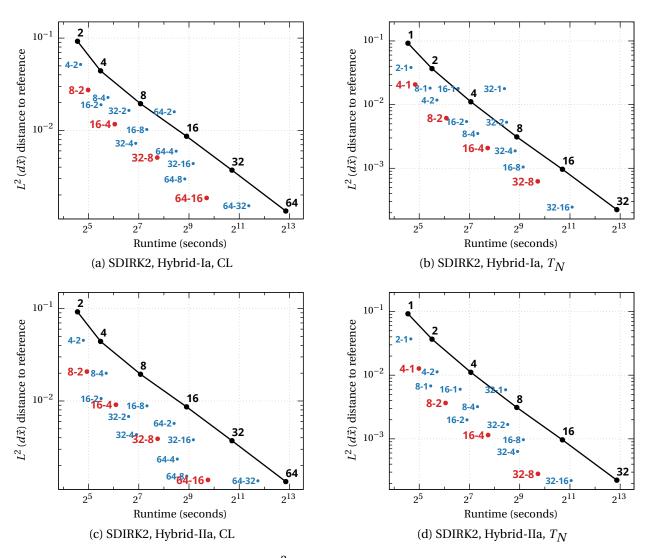


Figure 3.26: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the SDIRK2 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

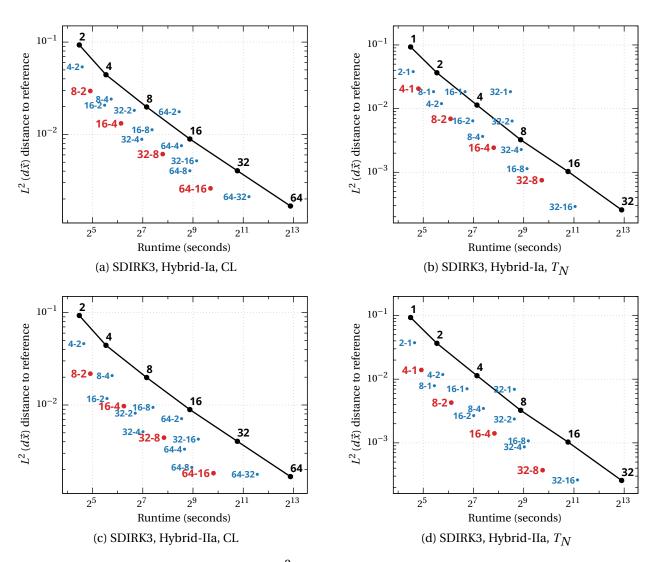


Figure 3.27: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the SDIRK3 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

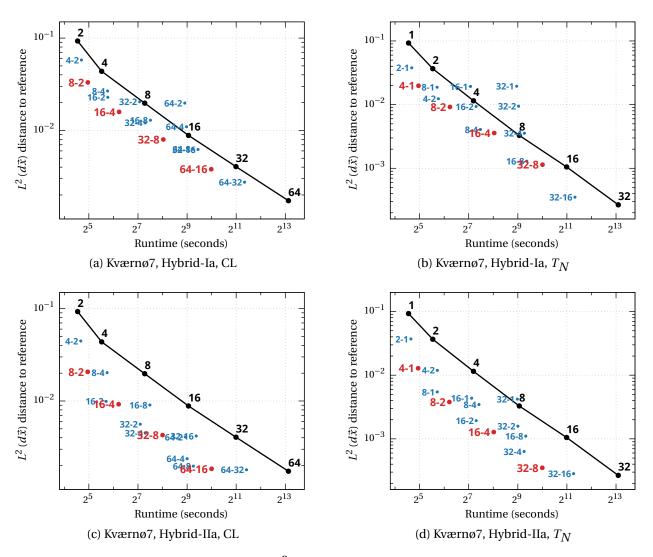


Figure 3.28: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the Kværnø7 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

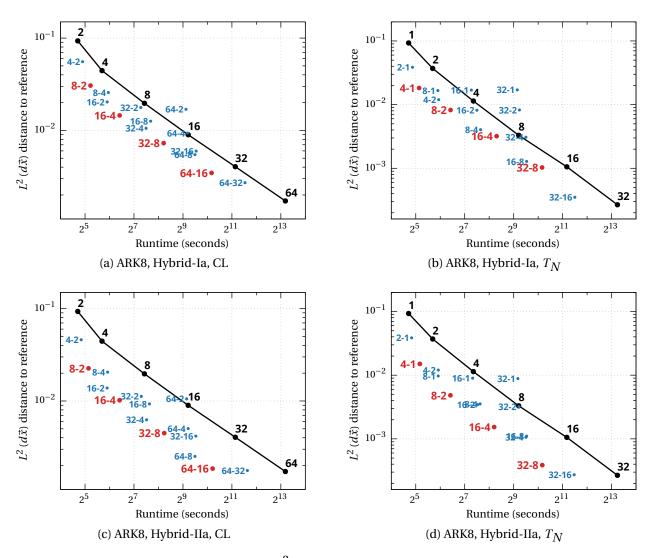


Figure 3.29: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the ARK8 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

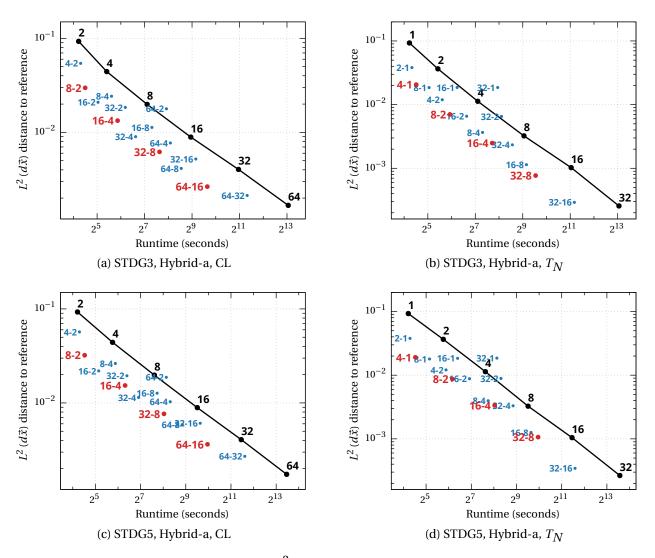


Figure 3.30: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the STDG3 and STDG5 integrators applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

Ia methods for the reconstruction procedures considered here. However, the use of other reconstruction procedures that are significantly more expensive than these may impact the relatively small efficiency advantage observed here for the hybrid-Ib methods.

In order for the hybrid approach to be effective in realistic scenarios, it is important to determine how the resolutions of the quadratures applied to the collided and uncollided components should be chosen to balance reductions in solve time with increases in solution error. In an attempt to characterize the computational efficiency of the hybrid methods, the second set of results presented here plots the solve time required by each method against the error of the solution as measured against a reference solution. Here errors are computed for the scalar flux in the L^2 norm. If method A lies below and to the left of method B, then method A requires less solve time and yields a smaller error than method B: in this case we say that method A is *more efficient* than method B. For each integrator, hybrid method, and angular quadrature type we survey a wide variety of angular quadrature resolutions, considering only cases where $N_C < N_U$. Results for DIRK methods are shown in figures 3.26 to 3.29, results for STDG methods are shown in figure 3.30, results for IDC methods are shown in figures 3.31 and 3.32, and results for LS-IDC methods are shown in figures 3.33 and 3.34.

We make two primary observations: (i) hybrid methods using angular quadratures that differ in order by a factor of 4 tend to consistently produce good computational efficiency relative to standard discrete ordinates methods and other hybrid methods; and (ii) in the vast majority of cases hybrid methods using angular quadratures of order 2N and N/2 for the uncollided and collided components, respectively, yield solutions with comparable or smaller errors in less time that a standard discrete ordinates method using an angular quadrature of order N. In most cases, the aforementioned hybrid methods reduce solve time by a factor of between two and three for the hohlraum test problem compared to standard discrete ordinates methods. The efficiency gains offered by the hybrid methods depends strongly on the specific form of the hybrid method, the time integrator and timestep size used, and the type of angular quadrature. As noted before, the hybrid-IIa DIRK methods again yield consistently smaller errors than the corresponding hybrid-Ia DIRK methods for all DIRK schemes considered. Similarly, the hybrid-Ib IDC methods yield marginally smaller errors than the corresponding hybrid-Ia IDC methods.

Larger efficiency gains are observed for hybrid methods based on lower-order time integrators and CL quadratures, and smaller efficiency gains are observed for hybrid methods based on higher-order time integrators and T_N quadratures, particularly when higher-resolution angular quadratures are used.

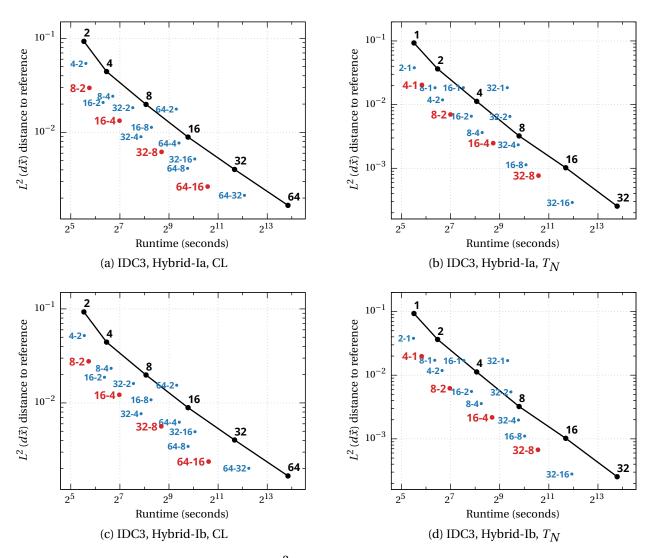


Figure 3.31: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the IDC3 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

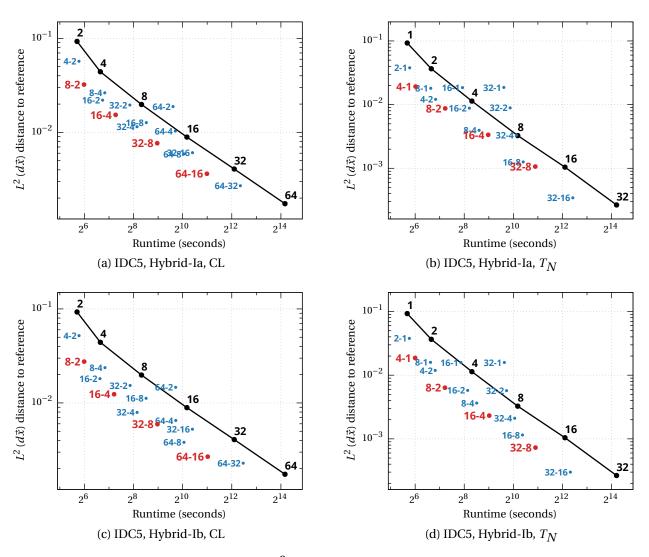


Figure 3.32: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the IDC5 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

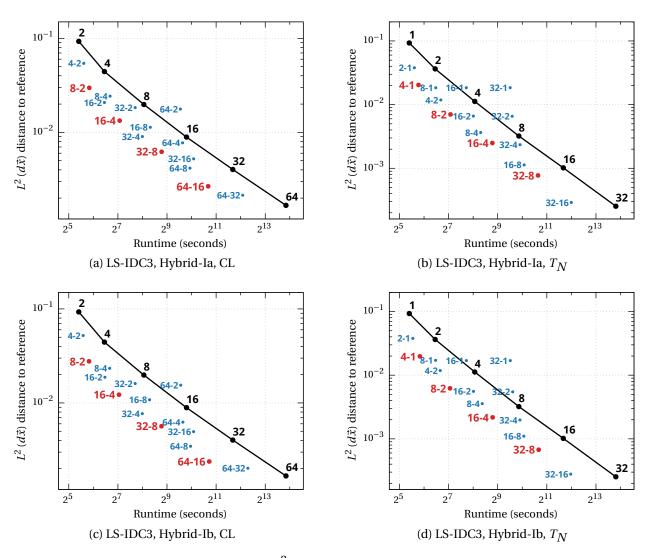


Figure 3.33: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the LS-IDC3 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

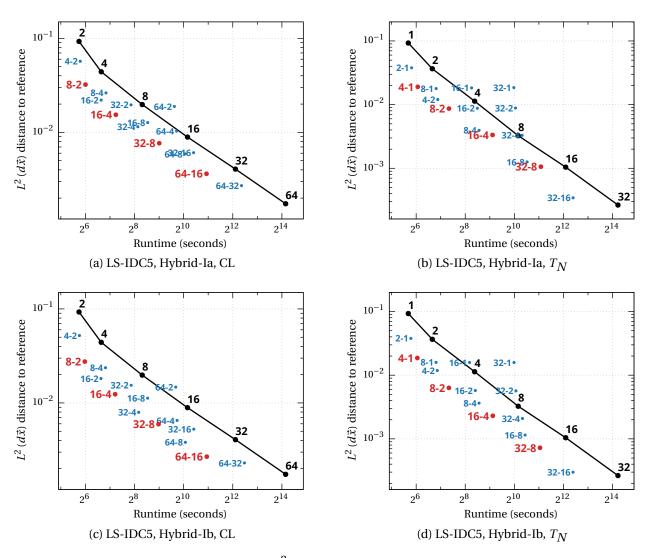


Figure 3.34: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the LS-IDC5 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

We believe that this disparity is caused mainly by a combination of two factors. First, while it is observed that efficiency gains are inversely correlated with the order of the time integrator that is used, we believe that the important factor here is the timestep size, not any property of the integrator such as its order of accuracy. Recall that all of the tests conducted here are performed using a constant CFL for all methods, so that higher-order methods use larger timestep sizes than lower-order methods. A larger timestep permits a larger proportion of particles within the system to experience collision events during each timestep, pushing more of the solution into the collided component.

The second contributing factor is the efficiency of the angular quadratures themselves. Figures 3.26 to 3.34 compare hybrid discrete ordinates methods to standard discrete ordinates methods using the same type of angular quadrature. This comparison implicitly ignores the fact that the angular quadratures considered here (CL and T_N) have radically different properties and hence behave differently when applied to certain problems. Comparisons between standard discrete ordinates methods using CL and T_N quadratures and hybrid discrete ordinates methods using CL and T_N quadratures that differ by a factor of 4 are shown in figures 3.35 to 3.39. Since the distribution of quadrature nodes varies significantly between the two quadrature types, T_N quadratures require many fewer nodes to attain a certain error tolerance than CL quadratures when applied to the hohlraum problem. This has two effects: First, the tolerances at which the hybrid and nonhybrid methods are compared is different for \mathcal{T}_N and CL quadratures. Second, the difference in accuracy between angular quadratures that differ by a factor of 4 is different for T_N and CL quadratures. We expect that this second fact contributes to the disparity in the efficiency gains observed fro different quadratures. While choosing angular quadratures that differ in order by a factor of 4 seems to produce acceptable results here, it may be the case that a different choice of angular quadrature resolutions, potentially dependent on the type of quadrature used and the level of refinement, may yield better results than what is observed here.

One important observation to be made from figures 3.35 to 3.39 is that any efficiency advantage derived from the hybrid approach is dependent on the angular quadrature that is used. Just as the choice of angular quadrature can significantly impact the accuracy of standard discrete ordinates methods, making a prudent choice of angular quadratures for the collided and uncollided components of the hybrid methods remains crucial to the overall efficiency of the methods. While hybrid methods using CL quadratures can be more computationally efficient that standard discrete ordinates methods using T_N quadratures can be more computationally

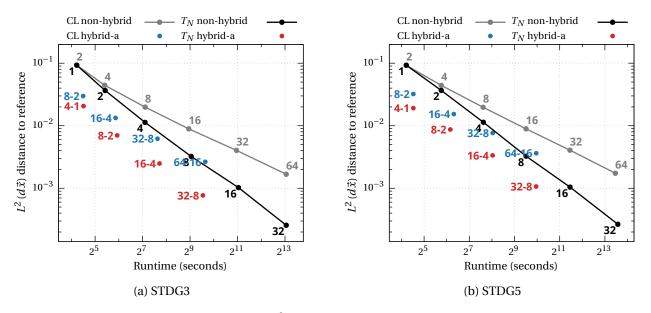


Figure 3.35: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray) and hybrid (red and blue) discrete ordinates methods based on STDG integrators and various angular resolutions applied to the hohlraum test problem. CL (gray and blue) and T_N (black and red) angular quadratures are considered.

efficient than hybrid methods using CL quadratures.

Overall, the results presented here suggest that the hybrid approach can be an effective technique for reducing the computation time of standard discrete ordinates methods. However this comes at the cost of increased memory consumption. The recommendation that a standard discrete ordinates method using an angular quadrature of order N be replaced by a hybrid discrete ordinates method using angular quadratures of order 2N and N/2 in the uncollided and collided components, respectively, increases memory consumption in two ways. First, because the number of quadrature nodes scales proportionally to the second power of the order of the angular quadrature, storing the uncollided components of the stages requires four times the memory of the corresponding nonhybrid method. Second, storing the collided components of the stages requires one-fourth the memory of the corresponding nonhybrid method, but this is in addition to the storage of the uncollided component. The result is that, for each time integrator, replacing a standard discrete ordinates method with a hybrid discrete ordinates method increases memory consumption by a factor of approximately 4.25. This value may be marginally lower for certain methods, such as the hybrid-IIa DIRK methods, where only one collided stage must be stored at any given time. These methods will nevertheless increase memory consumption by a factor of at least 4 over the corresponding standard discrete ordinates method. In cases where traditional discrete ordinates ordinates method.

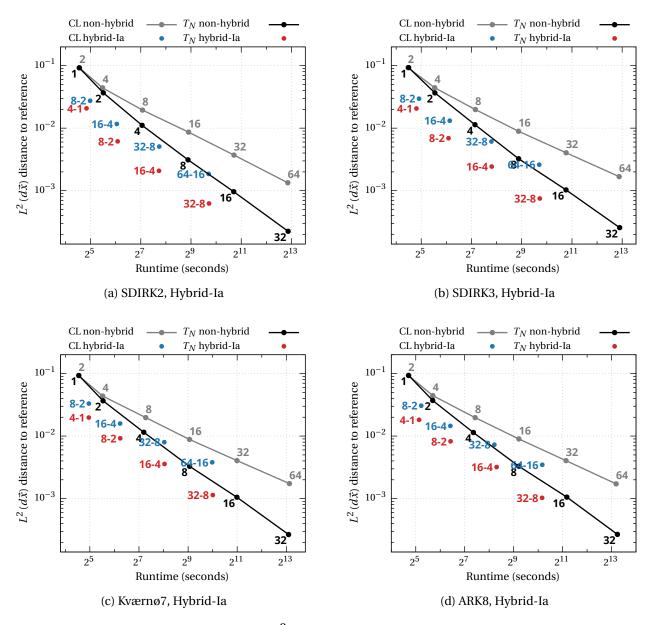


Figure 3.36: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray) and hybrid (red and blue) discrete ordinates methods based on DIRK integrators and various angular resolutions applied to the hohlraum test problem. CL (gray and blue) and T_N (black and red) angular quadratures are considered.

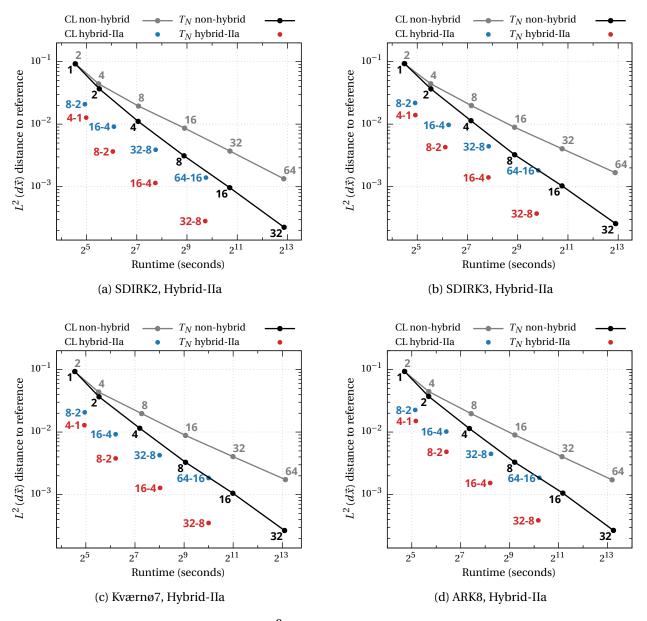


Figure 3.37: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray) and hybrid (red and blue) discrete ordinates methods based on DIRK integrators and various angular resolutions applied to the hohlraum test problem. CL (gray and blue) and T_N (black and red) angular quadratures are considered.

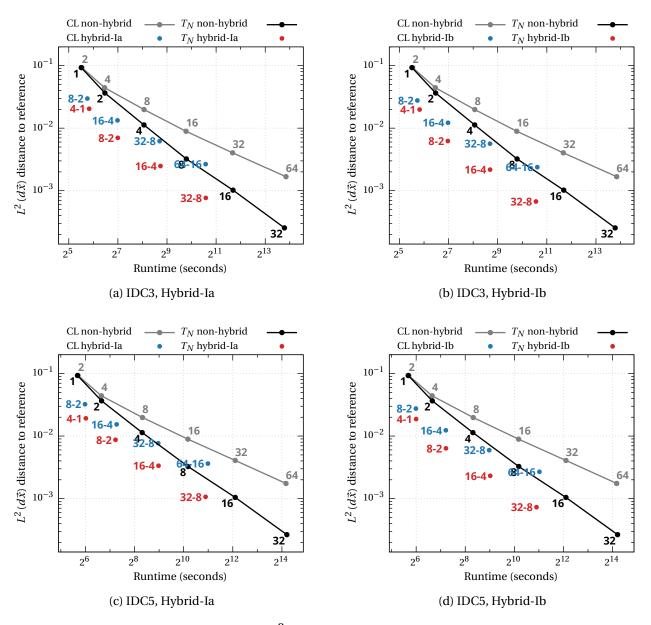


Figure 3.38: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray) and hybrid (red and blue) discrete ordinates methods based on IDC integrators and various angular resolutions applied to the hohlraum test problem. CL (gray and blue) and T_N (black and red) angular quadratures are considered.

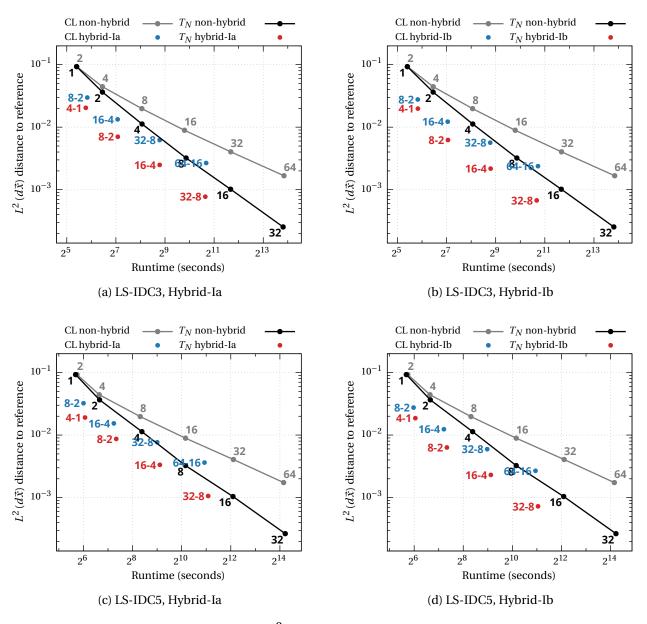


Figure 3.39: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray) and hybrid (red and blue) discrete ordinates methods based on LS-IDC integrators and various angular resolutions applied to the hohlraum test problem. CL (gray and blue) and T_N (black and red) angular quadratures are considered.

nates methods are already limited by available system memory, the savings in solve time that the hybrid methods discussed here provide may be overshadowed by the burded of increased memory consumption.

CHAPTER 4

NYSTRÖM RECONSTRUCTION AND TWO-GRID DEFECT CORRECTION FOR HYBRID METHODS

4.1 Introduction

At its most basic level, the discrete ordinates method is based on the principle of collocation. Many methods for the numerical solution of integral equations have been developed based on this principle. The purpose of this chapter is to examine several classes of hybrid methods derived from the collision-based splitting (3.16) that incorporate various approaches for extending the ideas behind collocation methods. We focus on methods related to two approaches in particular: Nyström interpolation and two-grid iterative methods.

We provide brief descriptions of Nyström interpolation and two-grid iterative methods below. These descriptions are based on those of [17, Section 4.1] and [17, Section 6.2], respectively, and we refer the interested reader to the reference [17] for more detailed descriptions and analyses of these methods. Our descriptions are given in the context of a canonical integral equation of the second kind:

$$\lambda f(x) - \int_D K(x, y) f(y) \, dy = g(x), \qquad (x \in D). \tag{4.1}$$

Here $f: D \to \mathbb{R}$ is the unknown function whose values are sought, and the kernel $K: D \times D \to \mathbb{R}$ of the integral operator, the function $g: D \to \mathbb{R}$, and the scalar λ are assumed to be known. The set D is assumed to be a manifold embedded within a finite dimensional Euclidean space \mathbb{R}^M , e.g., the unit two-sphere \mathbb{S}^2 .

4.1.1 Nyström interpolation

Collocation methods use numerical quadrature to provide an approximation to the solution of an equation of the form (4.1) on a discrete set of points; specifically, at the nodes of the quadrature that is applied. If solution values are desired at other points, then either these points must be incorporated into the quadrature that is used or the solution values at the quadrature nodes must be extended to the desired points in some manner after the values of the collocation solution have been computed. The first approach can only be used if the points at which the solution is desired are known a priori. As we will discuss, implementing the second approach is not necessarily straightforward.

Let $\{x_{n,j},w_{n,j}\}_{j=1}^{Nn}\subset D\times\mathbb{R}$ be a sequence of quadratures (indexed here by n) such that

$$\sum_{j=1}^{N_n} w_{n,j} f(x_{n,j}) \to \int_D f(x) dx \tag{4.2}$$

as $n \to \infty$ for any f belonging to a suitable space of functions (e.g., the space C(D) of continuous functions). Approximations f_n to the solution f of (4.1), each corresponding to one of the quadratures in the sequence, can be defined by replacing the integral operator in (4.1) with an approximation using the quadrature:

$$\lambda f_n(x) - \sum_{j=1}^{N_n} w_{n,j} K(x, x_{n,j}) f_n(x_{n,j}) = g(x).$$
 (4.3)

There are two important observations to be made at this point. The first is that each function f_n is defined on the entire domain of integration D. The second is that each of the functions f_n can be defined uniquely through its values at the associated quadrature nodes. More specifically, for each n the function f_n is defined uniquely by the set of values

$$\left\{ f_n(x_{n,j}) \right\}_{j=1}^{N_n} \tag{4.4}$$

satisfying

$$\lambda f_n(x_{n,i}) - \sum_{i=1}^{N_n} w_{n,j} K(x_{n,i}, x_{n,j}) f_n(x_{n,j}) = g(x_{n,i}), \qquad (i = 1, ..., N_n).$$
(4.5)

Since the nodal values (4.4) satisfy (4.5) exactly, these values correspond to those of the collocation solution. The question at this point is: once the discrete values $\{f_n(x_{n,j})\}_{j=1}^{N_n}$ are known, how does one extend these values to obtain an approximate solution value for an arbitrary point in D?

Extending a discrete set of values to other points within some space is a problem that is commonly solved using various forms of interpolation or extrapolation. A general interpolation problem may be roughly stated as follows:

Given a set of values $\{y_j\}_{j=1}^N \subset \mathbb{R}$ on a set of points $\{x_j\}_{j=1}^N \subset X \subset \mathbb{R}^n$ find $f: X \to \mathbb{R}$ belonging to some set of functions F such that $f(x_j) = y_j$ for all $j = 1, \ldots, N$.

The space X and the set of functions F to which f should belong will depend on the problem. However, it is well-known that general forms of interpolation are not necessarily well-behaved in all situations (c.f.,

for example, [67, Chapter 7]). The situation is particularly bleak for interpolation problems on higher dimensional manifolds such as \mathbb{S}^2 , where solutions may not exist [170].

This generic approach to interpolation was used to define the reconstruction procedures in chapter 3. The nature of these reconstruction procedures is a result of the difficulties inherent to interpolation problems on spaces such as \mathbb{S}^2 . In addition to the fact that the reconstruction will typically need to be tailored to the properties of the specific quadrature that is being used, other sacrifices must often be made: Not all methods will preserve properties of the solution such as non-negativity, conserve quantities of the solution such as the current density, or actually interpolate the discrete values.

The key observation to make is that there is more information to work with in the case of collocation methods for integral equations than there is for a generic interpolation problem. More specifically, we know that the set of points (4.4) that we would like to interpolate satisfies the system of equations (4.5) and, in particular, that the approximate solution f_n satisfying (4.3) is defined uniquely by the point values (4.4) obtained from (4.5). Hence using the discrete values (4.4) in (4.3) yields an expression for any value of f_n in D:

$$f_n(x) = \frac{1}{\lambda} \left[g(x) + \sum_{j=1}^{N_n} w_{n,j} K(x, x_{n,j}) f_n(x_{n,j}) \right]. \tag{4.6}$$

The formula (4.6), originally presented in [151], is called the *Nyström interpolation formula*.

4.1.2 Two-grid iteration

The motivation for discussing collocation methods in the context of a sequence of quadratures satisfying the property (4.2) is that the approximate solutions f_n satisfying (4.3) converge (in some sense) to the true solution f of (4.1) as $n \to \infty$. In theory, this implies that it is possible to obtain an approximate solution to (4.1) with any level of accuracy by simply refining the quadrature that is used until the desired level of accuracy has been attained. However a simple brute-force approach is not necessarily effective in practice.

Computing each approximate solution f_n requires solving a linear system of equations given by (4.5), whose dimension goes to infinity as n does. This means that computing the approximate solutions becomes significantly more expensive as the number of quadrature nodes used increases. There are two factors that can make this effect particularly problematic. The first is that the rate at which the number of quadrature nodes increases is related to the dimension of the domain of integration D: when the di-

mension of D is large, the number of quadrature nodes will generally increase very rapidly. Additionally, if the kernel K is very complicated then the linear system (4.5) may be dense and lack a clear sense of structure that can be exploited to reduce the computational cost of each solve. Compounding the scaling of the number of quadrature nodes with the scaling of the computational cost associated with dense linear algebra solves can render many problems exceedingly difficult to solve to any meaningful degree of accuracy.

Two-grid methods seek to reduce the significant computational cost associated with solving the linear system (4.5) when the number of quadrature nodes is large. This is done using a residual-correction approach that is similar to that used by IDC and other defect correction methods. Whereas IDC methods use high- and low-order approximations on the same quadrature grid, two-grid methods vary the quadrature nodes themselves. The effect is nevertheless the same in both cases: the higher-resolution approximation is only applied when computing the residual, and cost savings are obtained by inverting a simpler, lower-resolution approximation at each iteration.

The two-grid method that we will consider for (4.1) can be derived as follows. For each n, define the numerical integration operator \mathcal{K}_n by

$$(\mathcal{K}_n f)(x) = \sum_{i=1}^{N_n} w_{n,j} K(x, x_{n,j}) f(x_{n,j}). \tag{4.7}$$

Then (4.3) can be written as

$$(\lambda - \mathcal{K}_n) f_n = g. \tag{4.8}$$

Suppose that $f_n^{(k)}$ is some approximation to the solution f_n of (4.8). Here $f_n^{(k)}$ is some initial guess or, as the superscript would suggest, a previous iterate of the two-grid process. If the residual and error are defined by

$$r_n^{(k)} = g - (\lambda - \mathcal{K}_n) f_n^{(k)} \tag{4.9}$$

and

$$e_n^{(k)} = f_n - f_n^{(k)}, (4.10)$$

respectively, then the error satisfies the equation

$$(\lambda - \mathcal{K}_n)e_n^{(k)} = r_n^{(k)}. (4.11)$$

If one solves (4.11) using a standard collocation approach, then the desired solution f_n is obtained exactly through (4.10). But this would provide no reduction in computational cost relative to solving

(4.8) directly because the dimension of the linear systems that must be solved in each case are the same. Instead, K_n is replaced by K_m with m < n to obtain

$$(\lambda - \mathcal{K}_m)e_m^{(k)} = r_n^{(k)},\tag{4.12}$$

which is solved using a collocation approximation in terms of the lower-resolution quadrature. This yields a lower-resolution approximation $e_m^{(k)}$ to $e_n^{(k)}$ satisfying

$$\lambda e_m^{(k)}(x_{m,i}) - \sum_{j=1}^{N_m} w_{m,j} K(x_{m,i}, x_{m,j}) e_m^{(k)}(x_{m,j}) = r_n^{(k)}(x_{m,i}), \tag{4.13}$$

for $i=1,\ldots,N_m$. Once the point values $\{e_m^{(k)}(x_{m,j})\}_{j=1}^{Nm}$ have been extended to a representation for $e_m^{(k)}$ on all of D, the approximate solution $f_n^{(k)}$ can be updated using

$$f_n^{(k+1)} = f_n^{(k)} + e_m^{(k)}. (4.14)$$

This process is then repeated until some convergence tolerance is satisfied.

The values of the function $e_m^{(k)}$ can be recovered from the point values $\{e_m^{(k)}(x_{m,j})\}_{j=1}^{Nm}$ using the Nyström interpolation formula (4.6). However, our initial efforts at constructing defect correction methods for the hybrid splitting (3.16) used the reconstruction procedures discussed in chapter 3. Results for these defect correction methods are included here for completeness. While we did not observe any behavior in the numerical tests that were conducted that we believe would suggest a lack of convergence, we note that it is not clear which of the reconstruction methods discussed previously would be guaranteed to produce convergent iterations in all situations. For standard two-grid methods as outlined above using Nyström interpolation, the following can be said:

Theorem 1. For a sequence of quadratures that are convergent in the sense of (4.2), the two-grid iteration is convergent, i.e., $f_n^{(k)} \to f_n$ as $k \to \infty$, for m sufficiently large.

While theorem 1 states that the two-grid method defined by (4.9), (4.12), and (4.14) converges for m sufficiently large, this theorem says little about the nature of this convergence. An examination of the proof given in [17, Theorem 6.2.1] reveals that this two-grid method does not converge monotonically: Instead the method is expected to follow an even-odd pattern of convergence where every other iterate

converges monotonically but the overall iteration does not. We will see that the two-grid IDC methods that are presented in this work display similar behavior in certain situations.

It should be noted that a different two-grid method can be devised whose convergence behavior is superior to that of the method considered here. A description of this alternative method can be found in [17, Section 6.2.2], along with a proof that the method converges monotonically with a geometric rate of convergence. The discussion of two-grid methods in the present work is not intended to be an exhaustive comparison of all methods that can be derived using different combinations of IDC and hybrid techniques. Instead, it is meant to serve as a proof-of-concept that motivates future investigations into methods in which defect correction approaches targeted towards different discretization errors are combined into a single iterative procedure. Therefore, while it is possible to construct two-grid IDC methods based on this alternative technique, we reserve the consideration of such methods to future extensions of this work.

4.2 Method descriptions

We now turn our attention to describing hybrid methods for the transport equation (1.2) using the ideas behind Nyström interpolation and two-grid methods in combination with several classes of time-integration schemes.

4.2.1 Nyström reconstruction

We first describe modifications of the hybrid methods introduced in chapter 3 that have a strong similarity to the approach used in Nyström interpolation. In this context it is not necessarily clear that the term "interpolation" is appropriate, so we use the term *Nyström reconstruction* when referring to these methods. These methods are assigned the subtype c to denote the fact that a Nyström reconstruction is used. Because Nyström methods rely on the structure of the equation that one is trying to solve, the form of these reconstruction procedures depends on the time integrator used and the manner in which the splitting is applied. However, the general idea behind each specific instance is the same.

First-order hybrid method with Nyström reconstruction

The collided and uncollided components of the first-order hybrid method with Nyström reconstruction are computed using (3.20) as before. However, instead of defining the reconstructed flux ψ_*^n explicitly through (3.21), we return to the system (3.11) that is solved for the standard discrete ordinates method with an implicit Euler discretization in time. For the sake of consistency of notation, we first re-write the system (3.11) in the form

$$\left(\mathcal{L} + \frac{1}{\Delta t} - \mathcal{S}\mathcal{P}\right) \psi^n = \frac{1}{\Delta t} \psi^{n-1} + q^n, \tag{4.15}$$

ignoring the spatial discretization. The Nyström reconstruction is performed by using the collided and uncollided components obtained from (3.20) to approximate the scattering source in (4.15):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t}\right)\boldsymbol{\psi}_{*}^{n} = \mathcal{S}_{\mathbf{u}}\left(\mathcal{P}_{\mathbf{u}}\boldsymbol{\psi}_{\mathbf{u}}^{n} + \mathcal{P}_{\mathbf{c}}\boldsymbol{\psi}_{\mathbf{c}}^{n}\right) + \frac{1}{\Delta t}\boldsymbol{\psi}_{*}^{n-1} + \boldsymbol{q}^{n}.$$
(4.16)

The reconstructed angular flux distribution is defined implicitly by (4.16), which can be solved using a single transport sweep. We note that although the structure of (4.16) shares a significant degree of similarity with that of the source iteration method, (4.16) is not used to define an iteration procedure. Updating algorithm 3.1 to include the use of Nyström reconstruction yields algorithm 4.1.

The Nyström reconstruction procedure has a number of advantages. The first advantage is the flexibility of the method: Because this procedure is defined entirely through the operators of which the transport equation is composed, there is no explicit dependence on the problem geometry or choice of angular quadratures. This also means that hybrid methods leveraging this reconstruction can be used in existing discrete ordinates implementations without requiring significant modifications. Second, the accuracy of this procedure tends to be superior to that of the relabeling procedures described in section 3.3.2 – see section 4.4 for comparisons. Finally, when the Nyström reconstruction is used, the hybrid method naturally reduces to the non-hybrid, standard discrete ordinates method when the quadratures applied to the collided and uncollided components are the same. We emphasize two facts here: (i) this result is independent of the specific quadrature that is applied, and (ii) not all reconstruction procedures satisfy the above property, including the hyperinterpolation method used for CL quadratures.

Proposition 5. Suppose that the same quadrature is applied to the collided and uncollided components of the hybrid method (3.20) and that $\boldsymbol{\psi}_*^n$ is computed through (4.16). Then $\boldsymbol{\psi}_*^n = \boldsymbol{\psi}_{\mathrm{u}}^n + \boldsymbol{\psi}_{\mathrm{c}}^n$ and $\boldsymbol{\psi}_*^n = \boldsymbol{\psi}^n$ where $\boldsymbol{\psi}^n$ is the solution of the standard discrete ordinates system (4.15).

Algorithm 4.1 Hybrid Euler timestep with Nyström reconstruction.

input: Initial condition ψ_*^{n-1} .

1: Compute ψ_{11}^{n} by solving (3.20a):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n} = \frac{1}{\Delta t} \boldsymbol{\psi}_{*}^{n-1} + \boldsymbol{q}^{n}.$$

2: Compute $\psi_{\rm C}^n$ by solving (3.20b):

$$\left(\mathcal{L}_{\mathbf{C}} + \frac{1}{\Delta t} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{n} = \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{n}.$$

3: switch reconstruction method do

4: **case** Use a reconstruction operator \mathcal{R} :

5: $\psi_*^n \leftarrow \psi_u^n + \mathcal{R}\psi_c^n.$

6: end case

7: **case** Use the Nyström reconstruction (4.16) by solving:

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t}\right) \boldsymbol{\psi}_{*}^{n} = \mathcal{S}_{\mathbf{u}} \left(\mathcal{P}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{n} + \mathcal{P}_{\mathbf{c}} \boldsymbol{\psi}_{\mathbf{c}}^{n}\right) + \frac{1}{\Delta t} \boldsymbol{\psi}_{*}^{n-1} + \boldsymbol{q}^{n}.$$

8: end case

9: end switch

return: ψ_*^n

Proof. Since the angular quadrature is assumed to be the same for both the collided and uncollided components, we drop the subscripts used on the operators in the hybrid equations with the understanding that $\mathcal{L}_u = \mathcal{L}_c = \mathcal{L}$, and similarly for \mathcal{S} and \mathcal{P} . For simplicity, we do not include the terms due to the inflow of radiation at the boundary, in which case the operator \mathcal{L} is linear. Using the equations in (3.20) to replace the terms on the right side of (4.16) implies that

$$\left(\mathcal{L} + \frac{1}{\Delta t}\right)\boldsymbol{\psi}_{*}^{n} = \left(\mathcal{L} + \frac{1}{\Delta t}\right)\boldsymbol{\psi}_{c}^{n} + \left(\mathcal{L} + \frac{1}{\Delta t}\right)\boldsymbol{\psi}_{u}^{n} = \left(\mathcal{L} + \frac{1}{\Delta t}\right)\left(\boldsymbol{\psi}_{u}^{n} + \boldsymbol{\psi}_{c}^{n}\right), \tag{4.17}$$

from which the result $\psi_*^n = \psi_u^n + \psi_c^n$ follows through inversion of the streaming operator. The fact that $\psi_*^n = \psi^n$ follows by simply combining the terms of the scattering source in (4.16) and observing that the resulting equation for ψ_*^n is the same as that of equation (4.15) for ψ^n .

Results similar to proposition 5 can be stated for the higher-order methods discussed in the remainder of this section. These results are essentially a side-effect of the fact that the splitting (3.16) is exact on the continuum level, and hence the hybrid approach using identical quadratures and Nyström reconstruction yields a matrix inversion identity for each method.

Hybrid-Ic DIRK method

The collided and uncollided components of the hybrid-Ic DIRK method are computed using the same equations (3.25) as the hybrid-Ia DIRK method. The Nyström reconstruction procedure is performed at each stage using the following modified form of (3.24b):

$$\boldsymbol{\psi}_{*}^{(i)} = \boldsymbol{\psi}_{*}^{n-1} - \Delta t a_{ii} \left[\mathcal{L}_{u} \boldsymbol{\psi}_{*}^{(i)} - \mathcal{S}_{u} \left(\mathcal{P}_{u} \boldsymbol{\psi}_{u}^{(i)} + \mathcal{P}_{c} \boldsymbol{\psi}_{c}^{(i)} \right) - \boldsymbol{q}^{(i)} \right]$$

$$- \Delta t \sum_{j=1}^{i-1} a_{ij} \left[\left(\mathcal{L}_{u} - \mathcal{S}_{u} \mathcal{P}_{u} \right) \boldsymbol{\psi}_{*}^{(j)} - \boldsymbol{q}^{(j)} \right],$$

$$(4.18)$$

for $i=1,\ldots,s$. Note that, in contrast to the hybrid-Ia DIRK method, the reconstruction must be performed for each stage in sequence. Once the reconstructed stage values $\{\psi_*^{(i)}\}_{i=1}^s$ are known, the solution for the timestep is computed using (3.28) or is obtained directly from $\psi_*^{(s)}$ if the DIRK scheme is stiffly accurate. A result similar to that of proposition 5 for the hybrid-Ic DIRK method is stated in proposition 6. Updating algorithm 3.3 to include the use of Nyström reconstruction yields algorithm 4.2.

Proposition 6. Suppose that the same quadrature is applied to the collided and uncollided components of the hybrid-I DIRK method (3.25) and that the reconstructed stage values $\boldsymbol{\psi}_*^{(i)}$ for $i=1,\ldots,s$ are computed through (4.18). Then the reconstructed stage values satisfy $\boldsymbol{\psi}_*^{(i)} = \boldsymbol{\psi}_u^{(i)} + \boldsymbol{\psi}_c^{(i)}$ and $\boldsymbol{\psi}_*^{(i)} = \boldsymbol{\psi}_u^{(i)}$ for $i=1,\ldots,s$, where $\boldsymbol{\psi}_*^{(i)}$ are the stages for the standard discrete ordinates method (3.24b).

Proof. As in the proof of proposition 5, we drop the subscripts used on the operators of the hybrid equations and omit the terms due to the inflow of radiation at the boundary. We proceed by induction on the stage index i: Suppose that $i \ge 1$ and that proposition 6 is true for each stage with index less than i. From (3.25c) and (3.25d) it follows that

$$\frac{1}{\Delta t a_{ii}} \boldsymbol{\psi}_{*}^{n-1} + \boldsymbol{q}^{(i)} = \left(\boldsymbol{\mathcal{L}} + \frac{1}{\Delta t a_{ii}} \right) \boldsymbol{\psi}_{u}^{(i)} + \sum_{i=1}^{i-1} \frac{a_{ij}}{a_{ii}} \left[\boldsymbol{\mathcal{L}} \boldsymbol{\psi}_{u}^{(j)} - \boldsymbol{q}^{(j)} \right], \tag{4.19a}$$

$$\mathcal{SP}\left(\boldsymbol{\psi}_{\mathrm{u}}^{(i)} + \boldsymbol{\psi}_{\mathrm{c}}^{(i)}\right) = \left(\mathcal{L} + \frac{1}{\Delta t a_{ii}}\right)\boldsymbol{\psi}_{\mathrm{c}}^{(i)} + \sum_{j=1}^{i-1} \frac{a_{ij}}{a_{ii}}\left[\mathcal{L}\boldsymbol{\psi}_{\mathrm{c}}^{(j)} - \mathcal{SP}\left(\boldsymbol{\psi}_{\mathrm{c}}^{(j)} + \boldsymbol{\psi}_{\mathrm{u}}^{(j)}\right)\right]. \tag{4.19b}$$

Replacing the terms on the right side of (4.18) using (4.19) yields

$$\left(\mathcal{L} + \frac{1}{\Delta t a_{ii}}\right)\boldsymbol{\psi}_{*}^{(i)} = \left(\mathcal{L} + \frac{1}{\Delta t a_{ii}}\right)\left(\boldsymbol{\psi}_{u}^{(i)} + \boldsymbol{\psi}_{c}^{(i)}\right) + \sum_{j=1}^{i-1} \frac{a_{ij}}{a_{ii}}\left(\mathcal{L} - \mathcal{SP}\right)\left(\boldsymbol{\psi}_{u}^{(j)} + \boldsymbol{\psi}_{c}^{(j)} - \boldsymbol{\psi}_{*}^{(j)}\right)$$
(4.20a)

$$= \left(\mathcal{L} + \frac{1}{\Delta t a_{ii}}\right) \left(\boldsymbol{\psi}_{\mathbf{u}}^{(i)} + \boldsymbol{\psi}_{\mathbf{c}}^{(i)}\right), \tag{4.20b}$$

Algorithm 4.2 Hybrid-I DIRK timestep with Nyström reconstruction.

input: Initial condition ψ_*^{n-1} .

1: **for** i = 1, ..., s **do**

2: Compute
$$\boldsymbol{\psi}_{\mathrm{u}}^{(i)}$$
 by solving (3.25c):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t a_{ii}}\right) \boldsymbol{\psi}_{\mathbf{u}}^{(i)} = \frac{1}{\Delta t a_{ii}} \boldsymbol{\psi}_{*}^{n-1} - \sum_{j=1}^{i-1} \frac{a_{ij}}{a_{ii}} \left[\mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{(j)} - \boldsymbol{q}^{(j)}\right] + \boldsymbol{q}^{(i)}.$$

3: Compute $\psi_c^{(i)}$ by solving (3.25d):

$$\left(\mathcal{L}_{\mathbf{C}} + \frac{1}{\Delta t a_{ii}} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{(i)} = -\sum_{j=1}^{i-1} \frac{a_{ij}}{a_{ii}} \left[\left(\mathcal{L}_{\mathbf{C}} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{(j)} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{(j)} \right] + \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{(i)}.$$

4: end for

return: ψ_*^n

```
5: switch hybrid method do
                     case hybrid-Ia:
                               if the DIRK scheme is stiffly accurate then
  7:
                              \begin{aligned} \boldsymbol{\psi}_{\mathbf{u}}^{n} &\leftarrow \boldsymbol{\psi}_{\mathbf{u}}^{(s)}.\\ \boldsymbol{\psi}_{\mathbf{c}}^{n} &\leftarrow \boldsymbol{\psi}_{\mathbf{c}}^{(s)}.\\ \end{aligned} \mathbf{else}
  8:
  9:
10:
                                         \boldsymbol{\psi}_{\mathbf{u}}^{n} is defined by (3.25a): \boldsymbol{\psi}_{\mathbf{u}}^{n} \leftarrow \boldsymbol{\psi}_{*}^{n-1} - \Delta t \sum_{i=1}^{S} b_{i} \left[ \mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{(i)} - \boldsymbol{q}^{(i)} \right].
11:
                                         \boldsymbol{\psi}_{\mathrm{C}}^{n} is defined by (3.25b): \boldsymbol{\psi}_{\mathrm{C}}^{n} \leftarrow -\Delta t \sum_{i=1}^{S} b_{i} \left[ \left( \mathcal{L}_{\mathrm{C}} - \mathcal{S}_{\mathrm{C}} \mathcal{P}_{\mathrm{C}} \right) \boldsymbol{\psi}_{\mathrm{C}}^{(i)} - \mathcal{S}_{\mathrm{C}} \mathcal{P}_{\mathrm{U}} \boldsymbol{\psi}_{\mathrm{U}}^{(i)} \right].
12:
                               end if
13:
                    \psi_*^n \leftarrow \psi_u^n + \mathcal{R}\psi_c^n. end case
14:
15:
                     case hybrid-Ic:
16:
                               for i = 1, ..., s do
17:
                                         Compute \psi_*^{(i)} by solving (4.18):
18:
                \left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t a_{ij}}\right) \boldsymbol{\psi}_{*}^{(i)} = \frac{1}{\Delta t a_{ij}} \boldsymbol{\psi}_{*}^{n-1} + \mathcal{S}_{\mathbf{u}} \left(\mathcal{P}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{(i)} + \mathcal{P}_{\mathbf{c}} \boldsymbol{\psi}_{\mathbf{c}}^{(i)}\right) + \boldsymbol{q}^{(i)} - \sum_{i=1}^{i-1} \frac{a_{ij}}{a_{ij}} \left[ \left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}}\right) \boldsymbol{\psi}_{*}^{(j)} - \boldsymbol{q}^{(j)} \right].
19:
                              end for
                               if the DIRK scheme is stiffly accurate then
20:
21:
22:
                                         \boldsymbol{\psi}_{*}^{n} is defined by (3.28): \boldsymbol{\psi}_{*}^{n} \leftarrow \boldsymbol{\psi}_{*}^{n-1} - \Delta t \sum_{i=1}^{S} b_{i} \left[ \left( \mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}} \right) \boldsymbol{\psi}_{*}^{(i)} - \boldsymbol{q}^{(i)} \right].
23:
24:
                     end case
25:
26: end switch
```

Algorithm 4.3 Hybrid-II DIRK timestep with Nyström reconstruction.

input: Initial condition ψ_*^{n-1} .

- 1: **for** i = 1, ..., s **do**
- 2: Compute $oldsymbol{\psi}_{\mathrm{II}}^{(i)}$ by solving (3.26a):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t a_{ii}}\right) \boldsymbol{\psi}_{\mathbf{u}}^{(i)} = \frac{1}{\Delta t a_{ii}} \boldsymbol{\psi}_{*}^{n-1} - \sum_{j=1}^{i-1} \frac{a_{ij}}{a_{ii}} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}}\right) \boldsymbol{\psi}_{*}^{(j)} - \boldsymbol{q}^{(j)} \right] + \boldsymbol{q}^{(i)}.$$

3: Compute $\psi_{\rm C}^{(i)}$ by solving (3.26b):

$$\left(\mathcal{L}_{C} + \frac{1}{\Delta t a_{ii}} - \mathcal{S}_{C} \mathcal{P}_{C}\right) \psi_{C}^{(i)} = \mathcal{S}_{C} \mathcal{P}_{U} \psi_{U}^{(i)}.$$

- 4: **switch** hybrid method **do** 5: **case** hybrid-IIa:
- 6: $\psi_*^{(i)} \leftarrow \psi_u^{(i)} + \mathcal{R}\psi_c^{(i)}$.
- 7: end case
- 8: **case** hybrid-IIc:
- 9: Compute $\boldsymbol{\psi}_{*}^{(i)}$ by solving (4.18):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{\Delta t a_{ii}}\right) \boldsymbol{\psi}_{*}^{(i)} = \frac{1}{\Delta t a_{ii}} \boldsymbol{\psi}_{*}^{n-1} + \mathcal{S}_{\mathbf{u}} \left(\mathcal{P}_{\mathbf{u}} \boldsymbol{\psi}_{\mathbf{u}}^{(i)} + \mathcal{P}_{\mathbf{c}} \boldsymbol{\psi}_{\mathbf{c}}^{(i)}\right) + \boldsymbol{q}^{(i)} - \sum_{j=1}^{i-1} \frac{a_{ij}}{a_{ii}} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}}\right) \boldsymbol{\psi}_{*}^{(j)} - \boldsymbol{q}^{(j)} \right].$$

- 10: end case
- 11: end switch
- **12: end for**
- 13: if the DIRK scheme is stiffly accurate then
- 14: $\boldsymbol{\psi}^{n} \leftarrow \boldsymbol{\psi}^{(S)}$
- 15: **else**
- 16: $\boldsymbol{\psi}_{*}^{n}$ is defined by (3.28): $\boldsymbol{\psi}_{*}^{n} \leftarrow \boldsymbol{\psi}_{*}^{n-1} \Delta t \sum_{i=1}^{s} b_{i} \left[\left(\mathcal{L}_{u} \mathcal{S}_{u} \mathcal{P}_{u} \right) \boldsymbol{\psi}_{*}^{(i)} \boldsymbol{q}^{(i)} \right].$
- 17: end if

return: ψ_*^n

where the second equality follows from the inductive hypothesis, or, for the case when i=1, from the fact that the sum contains zero terms. Hence the result $\boldsymbol{\psi}_*^{(i)} = \boldsymbol{\psi}_u^{(i)} + \boldsymbol{\psi}_c^{(i)}$ follows from the inversion of the streaming operator. The fact that $\boldsymbol{\psi}_*^{(i)} = \boldsymbol{\psi}^{(i)}$ follows by simply combining the collided and uncollided terms of the scattering source in (4.18) and observing that the resulting equation for $\boldsymbol{\psi}_*^{(i)}$ is the same as that of the equation (3.24b) for $\boldsymbol{\psi}^{(i)}$.

Hybrid-IIc DIRK method

The collided and uncollided components of the hybrid-IIc DIRK method are computed using the same equations (3.26) as the hybrid-IIa DIRK method. The reconstructed stage values $\psi_*^{(i)}$ are defined implicitly using the same equation (4.18) as the reconstructed stage values of the hybrid-Ic DIRK method. Equation (3.28) may also be used if the DIRK scheme is not stiffly-accurate. A result similar to that of proposition 5 for the hybrid-IIc DIRK method is stated in proposition 7. Updating algorithm 3.4 to include the use of Nyström reconstruction yields algorithm 4.3.

Proposition 7. Suppose that the same quadrature is applied to the collided and uncollided components of the hybrid-II DIRK method (3.26) and that the reconstructed stage values $\boldsymbol{\psi}_*^{(i)}$ for $i=1,\ldots,s$ are computed through (4.18). Then the reconstructed stage values satisfy $\boldsymbol{\psi}_*^{(i)} = \boldsymbol{\psi}_u^{(i)} + \boldsymbol{\psi}_c^{(i)}$ and $\boldsymbol{\psi}_*^{(i)} = \boldsymbol{\psi}^{(i)}$ for $i=1,\ldots,s$, where $\boldsymbol{\psi}^{(i)}$ are the stages for the standard discrete ordinates method (3.24b).

Proof. The proof proceeds in the same manner as proposition 6 with the relations (4.19) replaced by

$$\frac{1}{\Delta t a_{ii}} \boldsymbol{\psi}_{*}^{n-1} + \boldsymbol{q}^{(i)} = \left(\mathcal{L} + \frac{1}{\Delta t a_{ii}} \right) \boldsymbol{\psi}_{u}^{(i)} + \sum_{i=1}^{i-1} \frac{a_{ij}}{a_{ii}} \left[\left(\mathcal{L} - \mathcal{SP} \right) \boldsymbol{\psi}_{*}^{(i)} - \boldsymbol{q}^{(j)} \right], \tag{4.21a}$$

$$\mathcal{SP}\left(\boldsymbol{\psi}_{\mathrm{u}}^{(i)} + \boldsymbol{\psi}_{\mathrm{c}}^{(i)}\right) = \left(\mathcal{L} + \frac{1}{\Delta t a_{ii}}\right)\boldsymbol{\psi}_{\mathrm{c}}^{(i)}.$$
(4.21b)

Hybrid-c STDG method

The collided and uncollided components of the hybrid-c STDG method are computed using the same equations (3.46) as the hybrid-a STDG method. The coefficients of the reconstructed approximation $\boldsymbol{\psi}_{*}^{h,n}$ are defined implicitly through

$$\mathfrak{L}_{\mathbf{u}}^{h,n} \Psi_{*}^{h,n} = \mathfrak{S}_{\mathbf{u}}^{h,n} \left(\mathfrak{P}_{\mathbf{u}} \Psi_{\mathbf{u}}^{h,n} + \mathfrak{P}_{\mathbf{c}} \Psi_{\mathbf{c}}^{h,n} \right) + \mathbf{Q}^{h,n}. \tag{4.22}$$

The approximate solution at the time t_{n+1} is obtained by evaluating $\Psi_*^{h,n}$ at t_{n+1} . A result analogous to that of proposition 5 for the hybrid-c STDG method is stated in proposition 8. Updating algorithm 3.5 to include the use of Nyström reconstruction yields algorithm 4.4.

Algorithm 4.4 Hybrid STDG timestep with Nyström reconstruction.

input: Initial condition ψ_*^n .

1: Compute $\Psi_{\mathbf{u}}^{h,n}$ by solving (3.46a):

$$\mathfrak{L}_{\mathbf{u}}^{h,n}\Psi_{\mathbf{u}}^{h,n}=\mathbf{Q}^{h,n}.$$

2: Compute $\Psi^{h,n}_{\mathbb{C}}$ by solving (3.46b):

$$\left(\mathfrak{L}_{\mathsf{c}}^{h,n}-\mathfrak{S}_{\mathsf{c}}^{h,n}\mathfrak{P}_{\mathsf{c}}^{h,n}\right)\Psi_{\mathsf{c}}^{h,n}=\mathfrak{S}_{\mathsf{c}}^{h,n}\mathfrak{P}_{\mathsf{c}}^{h,n}\Psi_{\mathsf{u}}^{h,n}.$$

```
3: switch hybrid method do
4: case hybrid-a:
5: \Psi_*^{h,n} \leftarrow \Psi_{\mathbf{u}}^{h,n} + \Re \Psi_{\mathbf{c}}^{h,n}.
6: end case
7: case hybrid-c:
8: Compute \Psi_*^{h,n} by solving (4.22):
\mathfrak{L}_{\mathbf{u}}^{h,n} \Psi_*^{h,n} = \mathfrak{S}_{\mathbf{u}}^{h,n} \left( \mathcal{P}_{\mathbf{u}} \Psi_{\mathbf{u}}^{h,n} + \mathfrak{P}_{\mathbf{c}} \Psi_{\mathbf{c}}^{h,n} \right) + \mathbf{Q}^{h,n}.
9: end case
10: end switch
\mathbf{return:} \ \psi_*^{h,n} \Big|_{t=t_{n+1}}
```

Proposition 8. Suppose that the same quadrature is applied to the collided and uncollided components of the hybrid STDG method (3.46) and that the coefficients of the reconstructed approximation $\Psi_*^{h,n}$ are computed through (4.22). Then the coefficients of the reconstructed approximation satisfy $\Psi_*^{h,n} = \Psi_u^{h,n} + \Psi_c^{h,n}$ and $\Psi_*^{h,n} = \Psi_u^{h,n}$ where $\Psi_*^{h,n}$ are the coefficients for the standard discrete ordinates method (3.35).

Proof. The proof follows the same form as that of proposition 5.

Hybrid-Ic IDC method

The construction of a Nyström reconstruction procedure for hybrid-I IDC methods is less obvious. Nyström reconstruction procedures were constructed for implicit Euler, DIRK, and STDG methods by using the approximations of the collided and uncollided stage values obtained by solving the split system to approximate the scattering sources in the equations for the stages of the corresponding (un-split or non-hybrid) discrete ordinates system. The approximations of the collided and uncollided stage values for the hybrid-I IDC methods are computed by solving (3.53) and the corresponding expression for the

stage values of the non-hybrid IDC method is (3.51). At first glance, one might therefore suggest that a Nyström reconstruction procedure for hybrid-I IDC methods should be constructed based on (3.51).

Upon closer inspection, it becomes clear that the collided and uncollided fluxes obtained from (3.53) are incompatible with the form of (3.51) and that another approach must be taken. The reason that the collided and uncollided components obtained from solving (3.53) are inconsistent with the form of (3.51) is due to the fact that the operator \mathcal{L} is applied to the provisional solutions from the previous correction level; i.e., $\{\psi^{\ell,[p-1]}\}_{\ell=1}^N$. This inconsistency can be avoided by defining a Nyström reconstruction procedure for hybrid-I IDC methods using the non-hybrid collocation scheme. The justification for this is the observation that the provisional solutions computed by the IDC method converge to the stage values of the associated collocation scheme.

Applying the collocation approximation (2.8) to the discrete ordinates system (3.9) yields

$$\widehat{\boldsymbol{\psi}} = \widehat{\boldsymbol{\psi}}^{0} - \Delta t \left(\Theta \otimes (\mathcal{L} - \mathcal{SP})\right) \widehat{\boldsymbol{\psi}} + \Delta t \left(\Theta \otimes \mathcal{I}\right) \widehat{\boldsymbol{q}}, \tag{4.23}$$

where

$$\hat{\boldsymbol{\psi}}^T = \left[\left(\boldsymbol{\psi}^1 \right)^T, \dots, \left(\boldsymbol{\psi}^N \right)^T \right] \quad \text{and} \quad \hat{\boldsymbol{q}}^T = \left[\left(\boldsymbol{q}^1 \right)^T, \dots, \left(\boldsymbol{q}^N \right)^T \right]$$
 (4.24)

are the vectors of stage values and sources, respectively, with $\boldsymbol{q}^n = \boldsymbol{q}(t_n)$, $\hat{\boldsymbol{\psi}}^0 = \mathbb{1}_N \otimes \boldsymbol{\psi}^0$, and $\mathbb{1}_N$ is the N-dimensional column vector of ones. Let $\hat{\boldsymbol{\psi}}_u$ and $\hat{\boldsymbol{\psi}}_c$ denote the stacked vectors of provisional solutions obtained from (3.53) after p correction iterations; i.e.,

$$\left(\widehat{\boldsymbol{\psi}}_{\mathbf{u}}^{[p]}\right)^{T} = \left[\left(\boldsymbol{\psi}_{\mathbf{u}}^{1,[p]}\right)^{T}, \dots, \left(\boldsymbol{\psi}_{\mathbf{u}}^{N,[p]}\right)^{T}\right],\tag{4.25}$$

and similarly for $\widehat{\pmb{\psi}}_{\text{C}}.$ The reconstructed stages

$$\left(\widehat{\boldsymbol{\psi}}_{*}^{[p]}\right)^{T} = \left[\left(\widehat{\boldsymbol{\psi}}_{*}^{1,[p]}\right)^{T}, \dots, \left(\widehat{\boldsymbol{\psi}}_{*}^{N,[p]}\right)^{T}\right] \tag{4.26}$$

for the hybrid-Ic IDC method are defined implicitly through the following modified form of (4.23):

$$\left[\left(\mathcal{I} \otimes \mathcal{I}_{\mathbf{u}} \right) + \Delta t \left(\Theta \otimes \mathcal{L}_{\mathbf{u}} \right) \right] \widehat{\boldsymbol{\psi}}_{*}^{[p]} = \widehat{\boldsymbol{\psi}}_{*}^{0} + \Delta t \left(\Theta \otimes \mathcal{S}_{\mathbf{u}} \right) \left[\left(\mathcal{I} \otimes \mathcal{P}_{\mathbf{u}} \right) \widehat{\boldsymbol{\psi}}_{\mathbf{u}}^{[p]} + \left(\mathcal{I} \otimes \mathcal{P}_{\mathbf{c}} \right) \widehat{\boldsymbol{\psi}}_{\mathbf{c}}^{[p]} \right] + \Delta t \left(\Theta \otimes \mathcal{I}_{\mathbf{u}} \right) \widehat{\boldsymbol{q}}, \quad (4.27)$$

where $\hat{\psi}^0_* = \mathbb{1}_N \otimes \psi^0_*$. As before, the operator on the left of (4.27) can be inverted using a single transport sweep. Updating algorithm 3.7 to include the use of the Nyström reconstruction (4.27) yields algorithm 4.5.

Algorithm 4.5 Hybrid-Ia and hybrid-Ic IDC timestep.

input: Initial condition ψ^0_* , number of correction iterations *P*.

- 1: **for** p = 0, ..., P **do**
- $\psi_{\mathbf{u}}^{0,[p]} \leftarrow \psi_{*}^{0}.$ $\psi_{\mathbf{c}}^{0,[p]} \leftarrow 0.$

- 5: **for** n = 1, ..., N **do**

▶ Prediction steps.

⊳ Set initial condition.

Compute $\psi_{\mathrm{u}}^{n,[0]}$ by solving (3.52a):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{h_n \Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n,[0]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}_{\mathbf{u}}^{n-1,[0]} + \boldsymbol{q}^n.$$

Compute $\psi_{C}^{n,[0]}$ by solving (3.52b):

$$\left(\mathcal{L}_{\mathbf{C}} + \frac{1}{h_{n}\Delta t} - \mathcal{S}_{\mathbf{C}}\mathcal{P}_{\mathbf{C}}\right)\boldsymbol{\psi}_{\mathbf{C}}^{n,[0]} = \frac{1}{h_{n}\Delta t}\boldsymbol{\psi}_{\mathbf{U}}^{n-1,[0]} + \mathcal{S}_{\mathbf{C}}\mathcal{P}_{\mathbf{U}}\boldsymbol{\psi}_{\mathbf{U}}^{n,[0]}.$$

- 8: end for
- 9: **for** p = 1, ..., P **do**

▶ Correction iterations.

- **for** n = 1, ..., N **do** 10:
- Compute $\psi_{\rm u}^{n,[p]}$ by solving (3.53a): 11:

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{h_{n}\Delta t}\right)\boldsymbol{\psi}_{\mathbf{u}}^{n,[p]} = \frac{1}{h_{n}\Delta t}\boldsymbol{\psi}_{\mathbf{u}}^{n-1,[p]} + \sum_{\ell=1}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_{n}}\right)\mathcal{L}_{\mathbf{u}}\boldsymbol{\psi}_{\mathbf{u}}^{\ell,[p-1]} + \sum_{\ell=1}^{N} \frac{\gamma_{n,\ell}}{h_{n}}\boldsymbol{q}^{\ell}.$$

Compute $\psi_{C}^{n,[p]}$ by solving (3.53b): 12:

$$\begin{split} \left(\mathcal{L}_{\mathbf{C}} + \frac{1}{h_{n}\Delta t} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}}\right) \boldsymbol{\psi}_{\mathbf{C}}^{n,[p]} &= \frac{1}{h_{n}\Delta t} \boldsymbol{\psi}_{\mathbf{C}}^{n-1,[p]} + \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{n,[p]} \\ &+ \sum_{\ell=1}^{N} \left(\delta_{n,\ell} - \frac{\gamma_{n,\ell}}{h_{n}} \right) \left[\left(\mathcal{L}_{\mathbf{C}} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{C}} \right) \boldsymbol{\psi}_{\mathbf{C}}^{\ell,[p-1]} - \mathcal{S}_{\mathbf{C}} \mathcal{P}_{\mathbf{U}} \boldsymbol{\psi}_{\mathbf{U}}^{\ell,[p-1]} \right]. \end{split}$$

- end for 13:
- 14: end for
- 15: switch hybrid method do

▶ Reconstruction for final solution.

- 16:
- $\boldsymbol{\psi}_{*}^{N,[P]} \leftarrow \boldsymbol{\psi}_{\mathrm{u}}^{N,[P]} + \mathcal{R}\boldsymbol{\psi}_{\mathrm{c}}^{N,[P]}.$ 17:
- 18:
- case hybrid-Ic: 19:
- Compute $\hat{\boldsymbol{\psi}}_{*}^{[P]}$ by solving (4.27): 20:

$$\left[\left(\mathcal{I} \otimes \mathcal{I}_{u} \right) + \Delta t \left(\Theta \otimes \mathcal{L}_{u} \right) \right] \widehat{\boldsymbol{\psi}}_{*}^{[P]} = \widehat{\boldsymbol{\psi}}_{*}^{0} + \Delta t \left(\Theta \otimes \mathcal{S}_{u} \right) \left[\left(\mathcal{I} \otimes \mathcal{P}_{u} \right) \widehat{\boldsymbol{\psi}}_{u}^{[P]} + \left(\mathcal{I} \otimes \mathcal{P}_{c} \right) \widehat{\boldsymbol{\psi}}_{c}^{[P]} \right] + \Delta t \left(\Theta \otimes \mathcal{I}_{u} \right) \widehat{\boldsymbol{q}}.$$

- end case 21:
- 22: end switch

return: $\psi_*^{N,[P]}$

The hybrid-Ic IDC method has not been implemented at this time, and we include the description of this method here only for the sake of completeness. We anticipate that the accuracy of the hybrid-Ic IDC method would be similar to that of other integrators such as DIRK and STDG methods that leverage a Nyström reconstruction. The extension of proposition 5 to the hybrid-Ic IDC method requires a somewhat weaker statement: One would only be able to say that the non-hybrid and hybrid-Ic IDC methods produce identical results when the same quadrature is used if both IDC methods have been iterated to convergence with respect to the corresponding collocation solution¹. The disadvantage of the hybrid-Ic IDC method is that it effectively requires the application of two sweep algorithms: one for computing the collided and uncollided stages that operates on each of the temporal degrees of freedom independently, and one for computing the reconstructed stages which are fully coupled.

The coupling that is present in the time dimension in the reconstruction step is similar to that of STDG methods, and similar approaches can be applied to solve the reconstruction system (4.27). Our current implementation of the methods described in this work leverages separate solvers for Runge-Kutta methods and STDG methods. While we acknowledge that it is possible that this speaks to a design deficiency of our implementation, we have generally found that the performance of the RK and STDG solvers is maximized when different data layouts are used for each implementation that are tailored towards their particular memory access patterns. It is not clear how using a sub-optimal data layout for some parts of the hybrid-Ic IDC method would affect the overall performance of the method, or if it may be better to transpose the data layout between parts of the method. Because of this complexity, we have avoided the implementation of the hybrid-Ic IDC method in favor of pursuing work on the two-grid or hybrid-II IDC methods discussed in section 4.2.2, for which the implementation of a Nyström reconstruction procedure is much simpler.

4.2.2 Two-grid IDC methods

The hybrid-II IDC methods seek to leverage the defect correction properties of IDC methods to decrease the error in the hybrid approach due to the use of a lower-resolution angular discretization in the collided component. These methods have a strong similarity with standard two-grid methods for integral equations of the second kind (c.f. section 4.1.2). The construction of these methods is based on three ideas:

¹To be clear, the term *collocation solution* in this case refers to the collocation approximation (4.23) with respect to time, and not the discrete ordinates approximation in angle.

(i) apply the collision-based splitting (3.16) to the exact error equation (2.14c) for the discrete ordinates system (3.9); (ii) compute a set of reconstructed provisional solutions at each correction iteration using some reconstruction procedure; and (iii) apply the operators corresponding to the high-resolution discrete ordinates approximation to the reconstructed provisional solutions to compute the residual terms for each correction iteration. This last step is the fundamental mechanic of the defect correction approach: the residual should always be computed with respect to the operators corresponding to the higher-resolution or higher-order approximation that is ultimately desired. In this case, the residual is computed with respect to the operators obtained by applying a collocation approximation in time and a high-resolution discrete ordinates approximation in angle.

The derivation here alternates between elements of the derivation of the IDC method given in section 2.2.2 and the derivation of the hybrid discrete ordinates methods given in section 3.1.3, while attempting to follow both for some measure of consistency. We begin at the continuum level: Let $\psi(t)$ denote the exact solution of (3.2) and let $\psi_*^{(p-1)}(t)$ denote a provisional solution at correction level p-1. Let

$$e^{(p-1)}(t) = \psi(t) - \psi_*^{(p-1)}(t) \tag{4.28}$$

and

$$r^{(p-1)}(t) = \partial_t \psi_*^{(p-1)}(t) + (\mathcal{L} - \mathcal{SP}) \psi_*^{(p-1)}(t) - q(t)$$
(4.29)

denote the error and residual, respectively, of the provisional solution $\psi_*^{(p-1)}$. The exact error equation (2.14c) for $\psi_*^{(p-1)}$ takes the form

$$\partial_t e^{(p-1)}(t) = -(\widetilde{\mathcal{L}} - \mathcal{SP}) e^{(p-1)}(t) - \partial_t \int_{t_0}^t r^{(p-1)}(\tau) d\tau.$$
 (4.30)

Note that the inflow of radiation at the boundary is accounted for through the residual term: the streaming operator that is applied to the error $e^{(p-1)}(t)$ is that corresponding to the case of vacuum boundary conditions.

Applying the collision-based splitting (3.16) to (4.30) with $e^{(p-1)} = e_{\rm u}^{(p-1)} + e_{\rm c}^{(p-1)}$ and the residual term treated as a source term yields

$$\partial_t e_{\mathbf{u}}^{(p-1)} = -\widetilde{\mathcal{L}} e_{\mathbf{u}}^{(p-1)} - \partial_t \int_{t_0}^t r^{(p-1)}(\tau) d\tau, \tag{4.31a}$$

$$\partial_t e_{\mathbf{c}}^{(p-1)} = -\left(\tilde{\mathcal{L}} - \mathcal{SP}\right) e_{\mathbf{c}}^{(p-1)} + \mathcal{SP} e_{\mathbf{u}}^{(p-1)}. \tag{4.31b}$$

High- and low-resolution discrete ordinates approximations are now applied to the uncollided and collided components of (4.31), respectively, to obtain

$$\partial_t \boldsymbol{e}_{\mathrm{u}}^{(p-1)} = -\widetilde{\mathcal{L}}_{\mathrm{u}} \boldsymbol{e}_{\mathrm{u}}^{(p-1)} - \partial_t \int_{t_0}^t \boldsymbol{r}^{(p-1)}(\tau) d\tau, \tag{4.32a}$$

$$\partial_t \boldsymbol{e}_{\mathrm{C}}^{(p-1)} = -\left(\tilde{\mathcal{L}}_{\mathrm{C}} - \boldsymbol{\mathcal{S}}_{\mathrm{C}} \boldsymbol{\mathcal{P}}_{\mathrm{C}}\right) \boldsymbol{e}_{\mathrm{C}}^{(p-1)} + \boldsymbol{\mathcal{S}}_{\mathrm{C}} \boldsymbol{\mathcal{P}}_{\mathrm{U}} \boldsymbol{e}_{\mathrm{U}}^{(p-1)}. \tag{4.32b}$$

Note that the high-resolution discrete ordinates approximation used for the uncollided flux is applied to the residual term.

Two cases were considered for the hybrid-I IDC methods discussed in section 3.2.3: relabeling only the final step of the final correction iteration (hybrid-Ia) and relabeling after each substep of each correction iteration (hybrid-Ib). While relabeling only the final step can be advantageous when the relabeling procedure that is used is very computationally expensive to apply, the two-grid methods described here require a reconstruction to be performed for each stage of each correction iteration in order to compute the residuals of the provisional solutions. It is possible to defer the reconstruction of all substeps for a given correction level to the end of that correction iteration, but there is no performance advantage to be gained from doing so. However in this case, there are advantages to applying the reconstruction procedure between substeps. Namely, computing the reconstructed solution after each substep results in a simpler implementation and a reduction in the number of data vectors that must be stored, particularly for the LS-IDC methods discussed in section 4.2.3. For these reasons, we consider only two-grid methods wherein a reconstruction is applied after each substep.

When the reconstruction is applied between each substep, the initial set of provisional solutions are given by (3.57). Applying implicit Euler approximations to the time derivatives in (4.32) with the assumption that the initial conditions for the nth substep are given by $e_*^{n-1,[p-1]}$ and zero for the uncollided and collided components of the error, respectively, and expanding out the residual term yields

$$e_{\mathbf{u}}^{n,[p-1]} = -h_{n} \Delta t \widetilde{\mathcal{L}}_{\mathbf{u}} e_{\mathbf{u}}^{n,[p-1]} + \left(\boldsymbol{\psi}_{*}^{n-1,[p]} - \boldsymbol{\psi}_{*}^{n,[p-1]} \right)$$

$$-\Delta t \sum_{\ell=1}^{N} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}} \right) \boldsymbol{\psi}_{*}^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right],$$

$$(4.33a)$$

$$\boldsymbol{e}_{\mathrm{c}}^{n,[p-1]} = -h_{n} \Delta t \left[\left(\widetilde{\mathcal{L}}_{\mathrm{c}} - \boldsymbol{\mathcal{S}}_{\mathrm{c}} \boldsymbol{\mathcal{P}}_{\mathrm{c}} \right) \boldsymbol{e}_{\mathrm{c}}^{n,[p-1]} - \boldsymbol{\mathcal{S}}_{\mathrm{c}} \boldsymbol{\mathcal{P}}_{\mathrm{u}} \boldsymbol{e}_{\mathrm{u}}^{n,[p-1]} \right]. \tag{4.33b}$$

Correction iterations for the two-grid defect correction methods, which we also refer to as the hybrid-II IDC methods, are based on the error equations (4.33).

Algorithm 4.6 Hybrid-II IDC and LS-IDC prediction.

input: Initial condition ψ_*^0 , number of correction iterations P.

- 1: **for** p = 0,...,P **do**2: $\psi_*^{0,[p]} \leftarrow \psi_*^0$.
 3: **end for**4: **for** n = 1,...,N **do** \triangleright Set initial condition. \triangleright Prediction steps.
- 5: Compute $\psi_{11}^{n,[0]}$ by solving (3.57a):

$$\left(\mathcal{L}_{\mathbf{u}} + \frac{1}{h_n \Delta t}\right) \boldsymbol{\psi}_{\mathbf{u}}^{n,[0]} = \frac{1}{h_n \Delta t} \boldsymbol{\psi}_{*}^{n-1,[0]} + \boldsymbol{q}^{n}.$$

6: Compute $\psi_{\rm C}^{n,[0]}$ by solving (3.57b):

$$\left(\mathcal{L}_{C} + \frac{1}{h_{n}\Delta t} - \mathcal{S}_{C}\mathcal{P}_{C}\right)\boldsymbol{\psi}_{C}^{n,[0]} = \mathcal{S}_{C}\mathcal{P}_{U}\boldsymbol{\psi}_{U}^{n,[0]}.$$

- 7: **switch** hybrid method **do**8: **case** hybrid-IIb:
 9: $\boldsymbol{\psi}_{*}^{n,[0]} \leftarrow \boldsymbol{\psi}_{\mathrm{u}}^{n,[0]} + \mathcal{R}\boldsymbol{\psi}_{\mathrm{c}}^{n,[0]}$.
 10: **end case**11: **case** hybrid-IIc:
 12: Compute $\boldsymbol{\psi}_{*}^{n,[0]}$ by solving (4.16): $\left(\mathcal{L}_{\mathrm{u}} + \frac{1}{h_{n}\Delta t}\right)\boldsymbol{\psi}_{*}^{n,[0]} = \mathcal{S}_{\mathrm{u}}\left(\mathcal{P}_{\mathrm{u}}\boldsymbol{\psi}_{\mathrm{u}}^{n,[0]} + \mathcal{P}_{\mathrm{c}}\boldsymbol{\psi}_{\mathrm{c}}^{n,[0]}\right) + \frac{1}{h_{n}\Delta t}\boldsymbol{\psi}_{*}^{n-1,[0]} + \boldsymbol{q}^{n}.$
- 13: end case 14: end switch 15: end for

Two subtypes of hybrid-II IDC methods are considered: b and c. The hybrid-IIb IDC method employs a reconstruction operator \mathcal{R} that maps directly from the collided quadrature set to the uncollided quadrature set. The reconstruction operator is applied to the collided and uncollided components of the error approximation; that is, the reconstructed error approximation for each substep is given by

$$\boldsymbol{e}_{*}^{n,[p-1]} = \boldsymbol{e}_{11}^{n,[p-1]} + \mathcal{R}\boldsymbol{e}_{C}^{n,[p-1]}. \tag{4.34}$$

This approximation of the error is used to define the updated provisional solution

$$\boldsymbol{\psi}_{*}^{n,[p]} = \boldsymbol{\psi}_{*}^{n,[p-1]} + \boldsymbol{e}_{*}^{n,[p-1]}. \tag{4.35}$$

The hybrid-IIc IDC method leverages a Nyström reconstruction procedure for the initial set of provisional solutions that is analogous to that of (4.16). The correction iterations of the hybrid-IIc IDC method use a Nyström reconstruction procedure that is constructed using the error equation for the non-hybrid

IDC method discussed in section 3.2.3. The non-hybrid IDC method described in section 3.2.3 is written in a form that avoids explicitly constructing the error approximation; i.e., one solves (3.51) for the corrected solution directly. One can nevertheless write out the equation for the error approximation explicitly:

$$e^{n,[p-1]} = -h_n \Delta t (\widetilde{\mathcal{L}} - \mathcal{SP}) e^{n,[p-1]} + (\boldsymbol{\psi}^{n-1,[p]} - \boldsymbol{\psi}^{n,[p-1]})$$

$$-\Delta t \sum_{\ell=1}^{N} \gamma_{n,\ell} \left[(\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi}^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right].$$
(4.36)

The reconstruction equation for the hybrid-IIc IDC method is obtained by using the collided and uncollided components of the error from (4.33) to approximate the scattering source in (4.36):

$$\boldsymbol{e}_{*}^{n,[p-1]} = -h_{n} \Delta t \widetilde{\mathcal{L}}_{\mathbf{u}} \boldsymbol{e}_{*}^{n,[p-1]} + h_{n} \Delta t \mathcal{S}_{\mathbf{u}} \Big(\mathcal{P}_{\mathbf{u}} \boldsymbol{e}_{\mathbf{u}}^{n,[p-1]} + \mathcal{P}_{\mathbf{c}} \boldsymbol{e}_{\mathbf{c}}^{n,[p-1]} \Big)$$

$$+ \Big(\boldsymbol{\psi}_{*}^{n-1,[p]} - \boldsymbol{\psi}_{*}^{n,[p-1]} \Big) - \Delta t \sum_{\ell=1}^{N} \gamma_{n,\ell} \Big[\Big(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}} \Big) \boldsymbol{\psi}_{*}^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \Big].$$

$$(4.37)$$

Using (4.35) in (4.37) yields the following equation from which the updated solution $\psi_*^{n,[p]}$ can be obtained directly:

$$\boldsymbol{\psi}_{*}^{n,[p]} = -h_{n}\Delta t \mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{*}^{n,[p]} + h_{n}\Delta t \mathcal{S}_{\mathbf{u}} \left(\mathcal{P}_{\mathbf{u}} \boldsymbol{e}_{\mathbf{u}}^{n,[p-1]} + \mathcal{P}_{\mathbf{c}} \boldsymbol{e}_{\mathbf{c}}^{n,[p-1]} \right) + \boldsymbol{\psi}_{*}^{n-1,[p]}$$

$$+ h_{n}\Delta t \mathcal{L}_{\mathbf{u}} \boldsymbol{\psi}_{*}^{n,[p-1]} - \Delta t \sum_{\ell=1}^{N} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}} \right) \boldsymbol{\psi}_{*}^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right].$$

$$(4.38)$$

A result analogous to that of proposition 5 for the hybrid-IIc IDC method is stated in proposition 9. The algorithm for the prediction stage of the hybrid-IIb and hybrid-IIc IDC timesteps is outlined in algorithm 4.6, and the algorithm for the correction stages is outlined in algorithm 4.7.

Proposition 9. Suppose that the same quadrature is used to compute the initial provisional solutions $\psi_{\mathrm{u}}^{n,[0]}$ and $\psi_{\mathrm{c}}^{n,[0]}$ for the hybrid-IIIDC method and that the reconstructed provisional solutions $\psi_{*}^{n,[0]}$ are computed using the Nyström reconstruction formula (4.16). Suppose further that this quadrature is used to compute the collided and uncollided components of the error approximations of the hybrid-IIIDC method (4.33) and that the reconstructed error approximations $\mathbf{e}_{*}^{n,[p-1]}$ are computed through (4.37). Then the reconstructed error approximations satisfy $\mathbf{e}_{*}^{n,[p-1]} = \mathbf{e}_{\mathrm{u}}^{n,[p-1]} + \mathbf{e}_{\mathrm{c}}^{n,[p-1]}$ and $\mathbf{e}_{*}^{n,[p-1]} = \mathbf{e}_{*}^{n,[p-1]}$ for each substep $n = 1, \ldots, N$, where $\mathbf{e}^{n,[p-1]}$ is the error approximation for the standard IDC discrete ordinates method (4.36).

Proof. As before, we drop the subscripts used on the operators of the hybrid equations and omit the terms due to the inflow of radiation at the boundary. Since the initial provisional solutions are computed using the implicit Euler method, proposition 5 implies that the result holds for the initial set of provisional solutions.

The remainder of the proof proceeds by induction on the substep index n and the correction index p. Suppose that $n \ge 1$, p > 1, and that proposition 9 is true for each provisional solution with substep index less than p or correction index less than p. From (4.33) it follows that

$$(\mathcal{I} + h_n \Delta t \mathcal{L}) \boldsymbol{e}_{\mathbf{u}}^{n,[p-1]} = \boldsymbol{\psi}_*^{n-1,[p]} - \boldsymbol{\psi}_*^{n,[p-1]} - \Delta t \sum_{\ell=1}^N \gamma_{n,\ell} \left[(\mathcal{L} - \mathcal{SP}) \boldsymbol{\psi}_*^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right],$$
 (4.39a)

$$(\mathcal{I} + h_n \Delta t \mathcal{L}) \boldsymbol{e}_{\mathbf{c}}^{n,[p-1]} = h_n \Delta t \mathcal{S} \left(\mathcal{P} \boldsymbol{e}_{\mathbf{u}}^{n,[p-1]} + \mathcal{P} \boldsymbol{e}_{\mathbf{c}}^{n,[p-1]} \right). \tag{4.39b}$$

Replacing the terms on the right side of (4.37) using (4.39) yields

$$(\mathcal{I} + h_n \Delta t \mathcal{L}) \boldsymbol{e}_*^{n,[p-1]} = (\mathcal{I} + h_n \Delta t \mathcal{L}) (\boldsymbol{e}_{\mathbf{u}}^{n,[p-1]} + \boldsymbol{e}_{\mathbf{c}}^{n,[p-1]}),$$
 (4.40)

from which the result $e_*^{n,[p-1]} = e_u^{n,[p-1]} + e_c^{n,[p-1]}$ follows by inversion of the streaming operator. The fact that $e_*^{n,[p-1]} = e^{n,[p-1]}$ follows by simply combining the collided and uncollided terms of the scattering source in (4.37) and observing that the resulting equation for $e_*^{n,[p-1]}$ is the same as that of the equation (4.36) for $e^{n,[p-1]}$.

4.2.3 Two-grid LS-IDC methods

As was the case for the non-hybrid and hybrid-I methods discussed in section 3.2, the hybrid-II LS-IDC methods are obtained through simple modifications of the residual terms in the correction equations for the hybrid-II IDC methods. Modifying the residual terms in (4.33) yields the following equations from which the components of the error approximation for the hybrid-IIb and hybrid-IIc LS-IDC methods are computed:

$$e_{\mathbf{u}}^{n,[p-1]} = -h_{n} \Delta t \widetilde{\mathcal{L}}_{\mathbf{u}} e_{\mathbf{u}}^{n,[p-1]} - \Delta t \sum_{\ell=1}^{n-1} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}} \right) \boldsymbol{\psi}_{*}^{\ell,[p]} - \boldsymbol{q}^{\ell} \right]$$

$$+ \left(\boldsymbol{\psi}_{*}^{n-1,[p]} - \boldsymbol{\psi}_{*}^{n,[p-1]} \right) - \Delta t \sum_{\ell=n}^{N} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}} \right) \boldsymbol{\psi}_{*}^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right],$$

$$(4.41a)$$

$$\boldsymbol{e}_{\mathrm{c}}^{n,[p-1]} = -h_{n}\Delta t (\widetilde{\mathcal{L}}_{\mathrm{c}} - \mathcal{S}_{\mathrm{c}} \mathcal{P}_{\mathrm{c}}) \boldsymbol{e}_{\mathrm{c}}^{n,[p-1]} + h_{n}\Delta t \mathcal{S}_{\mathrm{c}} \mathcal{P}_{\mathrm{u}} \boldsymbol{e}_{\mathrm{u}}^{n,[p-1]}. \tag{4.41b}$$

Algorithm 4.7 Hybrid-II IDC correction.

```
1: for p = 1, ..., P do
                                                                                                                                                                                                                             ➤ Correction iterations.
  2:
                for n = 1, ..., N do
                        Compute e_{11}^{n,[p-1]} by solving (4.33a):
  3:
                    \left(\widetilde{\mathcal{L}}_{\mathbf{u}} + \frac{1}{h_n \Delta t}\right) e_{\mathbf{u}}^{n,[p-1]} = \frac{1}{h_n \Delta t} \left(\boldsymbol{\psi}_*^{n-1,[p]} - \boldsymbol{\psi}_*^{n,[p-1]}\right) - \sum_{\ell=1}^{N} \frac{\gamma_{n,\ell}}{h_n} \left[ \left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}}\right) \boldsymbol{\psi}_*^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right].
                        Compute e_c^{n,[p]} by solving (4.33a):
  4:
                                                                                  \left(\tilde{\mathcal{L}}_{C} + \frac{1}{h_{u} \wedge t} - \mathcal{S}_{C} \mathcal{P}_{C}\right) e_{C}^{n,[p-1]} = \mathcal{S}_{C} \mathcal{P}_{U} e_{U}^{n,[p-1]}
  5:
                        switch hybrid method do
                                case hybrid-IIb:
  6:
                                        \psi_*^{n,[p]} \leftarrow \psi_*^{n,[p-1]} + e_{11}^{n,[p-1]} + \mathcal{R}e_{C}^{n,[p-1]}
  7:
                                case hybrid-IIc:
  9:
                                        Compute \boldsymbol{\psi}_{*}^{n,[p]} by solving (4.38):
10:
                                \left(\mathcal{L}_{\mathrm{U}} + \frac{1}{h_{\mathrm{tr}} \wedge t}\right) \psi_{*}^{n,[p]} = \mathcal{S}_{\mathrm{U}} \left(\mathcal{P}_{\mathrm{U}} e_{\mathrm{U}}^{n,[p-1]} + \mathcal{P}_{\mathrm{C}} e_{\mathrm{C}}^{n,[p-1]}\right) + \frac{1}{h_{\mathrm{tr}} \wedge t} \psi_{*}^{n-1,[p]} + \mathcal{L}_{\mathrm{U}} \psi_{*}^{n,[p-1]}
                                                                                           -\sum_{\ell=1}^{N} \frac{\gamma_{n,\ell}}{h_n} \left[ (\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}}) \boldsymbol{\psi}_*^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right].
                                end case
11:
                         end switch
12:
13:
                end for
14: end for
        return: \boldsymbol{\psi}_{*}^{N,[P]}
```

The reconstructed error approximations for the hybrid-IIb LS-IDC method are defined by (4.34). For the hybrid-IIc LS-IDC method, reconstructed error approximations for each correction are computed by solving a modified form of (4.37):

$$e_{*}^{n,[p-1]} = -h_{n} \Delta t \widetilde{\mathcal{L}}_{\mathbf{u}} e_{*}^{n,[p-1]} + h_{n} \Delta t \left(\mathcal{P}_{\mathbf{u}} e_{\mathbf{u}}^{n,[p-1]} + \mathcal{P}_{\mathbf{c}} e_{\mathbf{c}}^{n,[p-1]} \right) + \left(\boldsymbol{\psi}_{*}^{n-1,[p]} - \boldsymbol{\psi}_{*}^{n,[p-1]} \right)$$

$$-\Delta t \sum_{\ell=1}^{n-1} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}} \right) \boldsymbol{\psi}_{*}^{\ell,[p]} - \boldsymbol{q}^{\ell} \right] - \Delta t \sum_{\ell=n}^{N} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}} \right) \boldsymbol{\psi}_{*}^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right]$$

$$(4.42)$$

Alternatively, one may compute the corrected solution directly using

$$\boldsymbol{\psi}_{*}^{n,[p]} = -h_{n} \Delta t \mathcal{L}_{\mathbf{u}} \left(\boldsymbol{\psi}_{*}^{n,[p]} - \boldsymbol{\psi}_{*}^{n-1,[p]} \right) + h_{n} \Delta t \mathcal{S}_{\mathbf{u}} \left(\mathcal{P}_{\mathbf{u}} \boldsymbol{e}_{\mathbf{u}}^{n,[p-1]} + \mathcal{P}_{\mathbf{c}} \boldsymbol{e}_{\mathbf{c}}^{n,[p-1]} \right) + \boldsymbol{\psi}_{*}^{n-1,[p]}$$

$$-\Delta t \sum_{\ell=1}^{n-1} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}} \right) \boldsymbol{\psi}_{*}^{\ell,[p]} - \boldsymbol{q}^{\ell} \right] - \Delta t \sum_{\ell=n}^{N} \gamma_{n,\ell} \left[\left(\mathcal{L}_{\mathbf{u}} - \mathcal{S}_{\mathbf{u}} \mathcal{P}_{\mathbf{u}} \right) \boldsymbol{\psi}_{*}^{\ell,[p-1]} - \boldsymbol{q}^{\ell} \right],$$

$$(4.43)$$

Algorithm 4.8 Hybrid-II LS-IDC correction.

```
1: for p = 1, ..., P do
                                                                                                                                                                                                                                                                             ➤ Correction iterations.
  2:
                    for n = 1, ..., N do
                             Compute e_{11}^{n,[p-1]} by solving (4.41a):
  3:
                                \left(\widetilde{\mathcal{L}}_{\mathrm{u}} + \frac{1}{h_{n}\Delta t}\right) \boldsymbol{e}_{\mathrm{u}}^{n,[p-1]} = \frac{1}{h_{n}\Delta t} \left(\boldsymbol{\psi}_{*}^{n-1,[p]} - \boldsymbol{\psi}_{*}^{n,[p-1]}\right) + \sum_{\ell=1}^{N} \frac{\gamma_{n,\ell}}{h_{n}} \boldsymbol{q}^{\ell}
                                                                                                             -\sum_{\ell=1}^{n-1} \frac{\gamma_{n,\ell}}{h_n} (\mathcal{L}_{\mathrm{u}} - \mathcal{S}_{\mathrm{u}} \mathcal{P}_{\mathrm{u}}) \boldsymbol{\psi}_*^{\ell,[p]} - \sum_{\ell=n}^{N} \frac{\gamma_{n,\ell}}{h_n} (\mathcal{L}_{\mathrm{u}} - \mathcal{S}_{\mathrm{u}} \mathcal{P}_{\mathrm{u}}) \boldsymbol{\psi}_*^{\ell,[p-1]}.
                             Compute e_c^{n,[p]} by solving (4.33a):
  4:
                                                                                                   \left(\widetilde{\mathcal{L}}_{\mathsf{C}} + \frac{1}{h_n \Delta t} - \mathcal{S}_{\mathsf{C}} \mathcal{P}_{\mathsf{C}}\right) e_{\mathsf{C}}^{n,[p-1]} = \mathcal{S}_{\mathsf{C}} \mathcal{P}_{\mathsf{U}} e_{\mathsf{U}}^{n,[p-1]}.
                             switch hybrid method do
  5:
                                      case hybrid-IIb:  \boldsymbol{\psi}_*^{n,[p]} \leftarrow \boldsymbol{\psi}_*^{n,[p-1]} + \boldsymbol{e}_{\mathrm{u}}^{n,[p-1]} + \mathcal{R}\boldsymbol{e}_{\mathrm{c}}^{n,[p-1]}. 
  6:
  9:
                                               Compute \psi_*^{n,[p]} by solving (4.43):
10:
                 \left(\mathcal{L}_{\mathbf{u}} + \frac{1}{h_{n}\Delta t}\right)\boldsymbol{\psi}_{*}^{n,[p]} = \mathcal{S}_{\mathbf{u}}\left(\mathcal{P}_{\mathbf{u}}\boldsymbol{e}_{\mathbf{u}}^{n,[p-1]} + \mathcal{P}_{\mathbf{c}}\boldsymbol{e}_{\mathbf{c}}^{n,[p-1]}\right) + \frac{1}{h_{n}\Delta t}\boldsymbol{\psi}_{*}^{n-1,[p]} + \mathcal{L}_{\mathbf{u}}\boldsymbol{\psi}_{*}^{n,[p-1]} + \sum_{\ell=1}^{N} \frac{\gamma_{n,\ell}}{h_{n}}\boldsymbol{q}^{\ell}
                                                                                         -\sum_{\ell=1}^{n-1}\frac{\gamma_{n,\ell}}{h_n}(\mathcal{L}_{\mathrm{u}}-\mathcal{S}_{\mathrm{u}}\mathcal{P}_{\mathrm{u}})\boldsymbol{\psi}_*^{\ell,[p]}-\sum_{\ell=n}^N\frac{\gamma_{n,\ell}}{h_n}(\mathcal{L}_{\mathrm{u}}-\mathcal{S}_{\mathrm{u}}\mathcal{P}_{\mathrm{u}})\boldsymbol{\psi}_*^{\ell,[p-1]}.
11:
                                       end case
12:
                              end switch
                    end for
13:
14: end for
          return: \psi_*^{N,[P]}
```

which is analogous to (4.38). A result analogous to that of proposition 5 for the hybrid-IIc LS-IDC method is stated in proposition 10. The algorithm for the prediction stage of the LS-IDC hybrid-II methods is the same as that of algorithm 4.6, and the algorithm for the correction stage of the LS-IDC hybrid-II methods is outlined in algorithm 4.8.

Proposition 10. Suppose that the same quadrature is used to compute the initial provisional solutions $\psi_{\rm u}^{n,[0]}$ and $\psi_{\rm c}^{n,[0]}$ for the hybrid-IILS-IDC method and that the reconstructed provisional solutions $\psi_{*}^{n,[0]}$ are computed using the Nyström reconstruction formula (4.16). Suppose further that this quadrature is used to compute the collided and uncollided components of the error approximations of the hybrid-II

LS-IDC method (4.41) and that the reconstructed error approximations $\mathbf{e}_{*}^{n,[p-1]}$ are computed through (4.42). Then the reconstructed error approximations $\mathbf{e}_{*}^{n,[p-1]}$ satisfy $\mathbf{e}_{*}^{n,[p-1]} = \mathbf{e}_{\mathrm{u}}^{n,[p-1]} + \mathbf{e}_{\mathrm{c}}^{n,[p-1]}$ and $\mathbf{e}_{*}^{n,[p-1]} = \mathbf{e}^{n,[p-1]}$ for each substep $n = 1,\ldots,N$, where $\mathbf{e}^{n,[p-1]}$ is the error approximation for the standard IDC discrete ordinates method (4.36).

Proof. The proof is essentially the same as that of proposition 9.

4.3 Numerical tests for Nyström methods

The same set of tests used in section 3.4 are repeated for the hybrid methods using Nyström reconstruction procedures. Results for two-grid methods (i.e., hybrid-II IDC and LS-IDC methods) are given in section 4.4. Descriptions of the problems and test procedures are not repeated here – see section 3.4 for detailed descriptions of these tests.

4.3.1 One spatial dimension

We first examine the convergence behavior of hybrid methods using Nyström reconstruction procedures. The results of the convergence tests for DIRK integrators are shown in figures 4.2 to 4.5 and for STDG integrators in figure 4.1.

The hybrid methods introduced in chapter 3 displayed two phenomena in these convergence tests: convergence saturation and order reduction. The saturation behavior of the hybrid methods remains unchanged when Nyström reconstruction procedures are used, but significant differences are observed with respect to the order reduction behavior of the methods. Whereas the $S_{32,2}$ hybrid methods using reconstruction procedures based on polynomial interpolation were observed to converge with second-order accuracy after the saturation point, we observe that $S_{32,2}$ hybrid methods using Nyström reconstruction procedures converge with fifth-order accuracy after the saturation point. This suggests a convergence rate of $2N_{\rm C}+1$ for hybrid methods using Nyström reconstruction procedures and an angular quadrature of order $N_{\rm C}$ in the collided component. Preliminary investigations into the order-reduction phenomenon observed in the hybrid methods support this observation²: Specifically, investigations into

²Since the results of these investigations were not fully available when the present work was published, this claim is, unfortunately, stated without justification. We hope to include a thorough justification for the observed order reduction phenomenon in future publications.

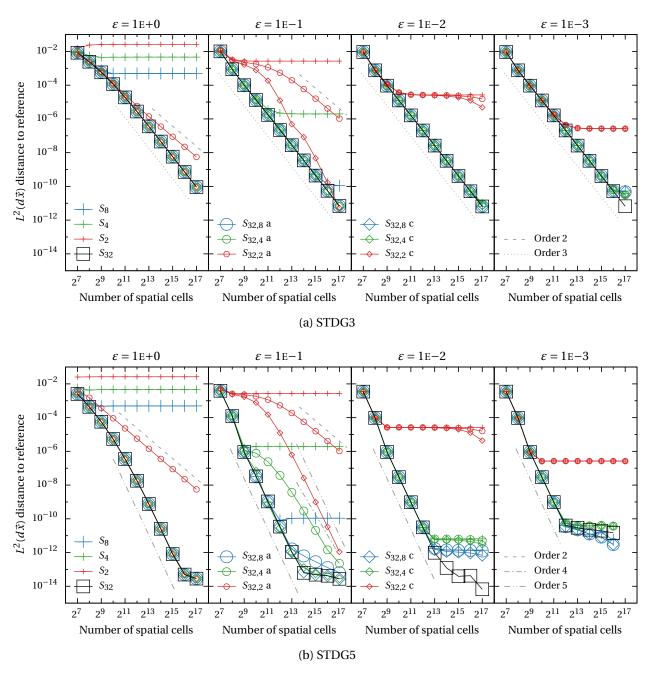


Figure 4.1: L^2 convergence studies of hybrid STDG methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

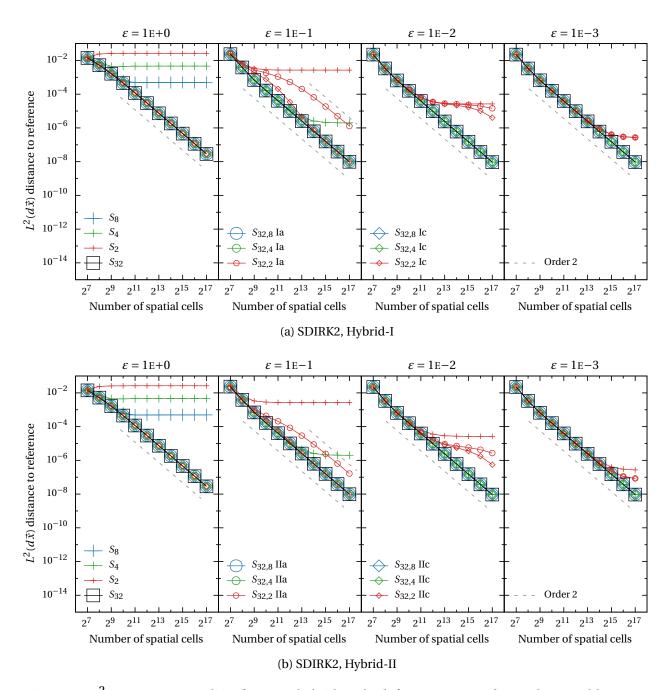


Figure 4.2: L^2 convergence studies of SDIRK2 hybrid methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

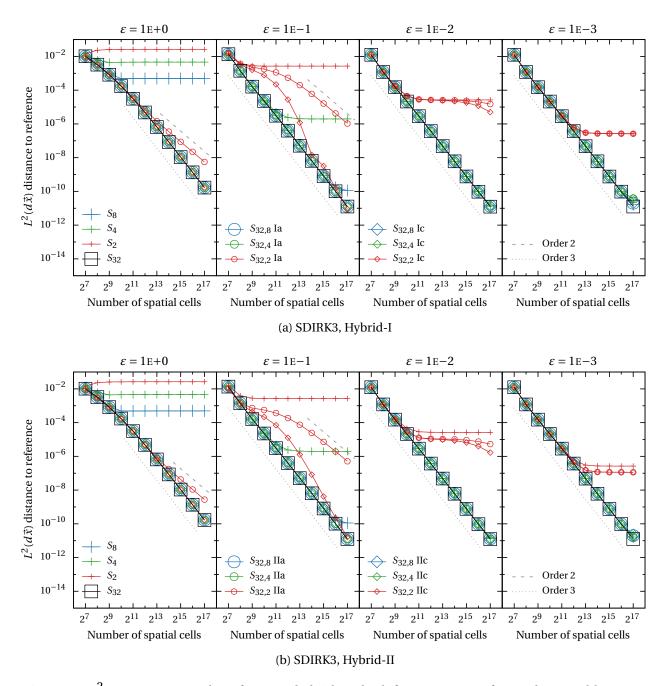


Figure 4.3: L^2 convergence studies of SDIRK3 hybrid methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

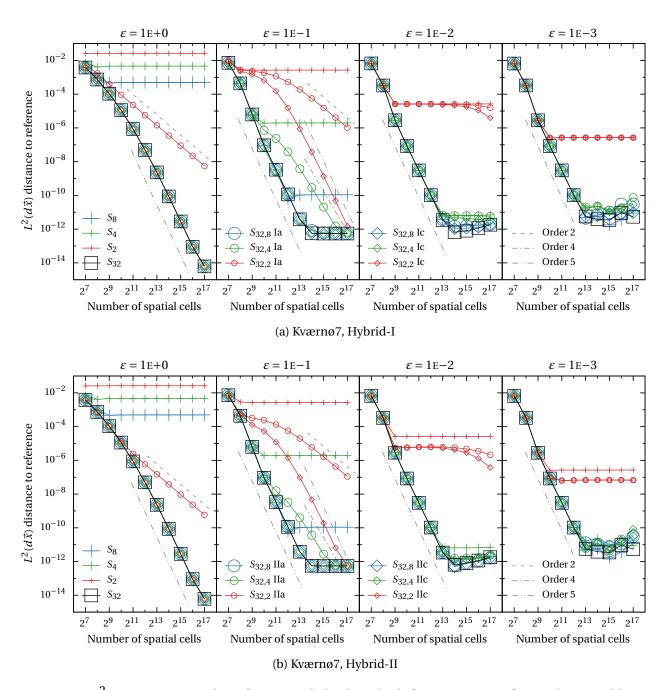


Figure 4.4: L^2 convergence studies of Kværnø7 hybrid methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

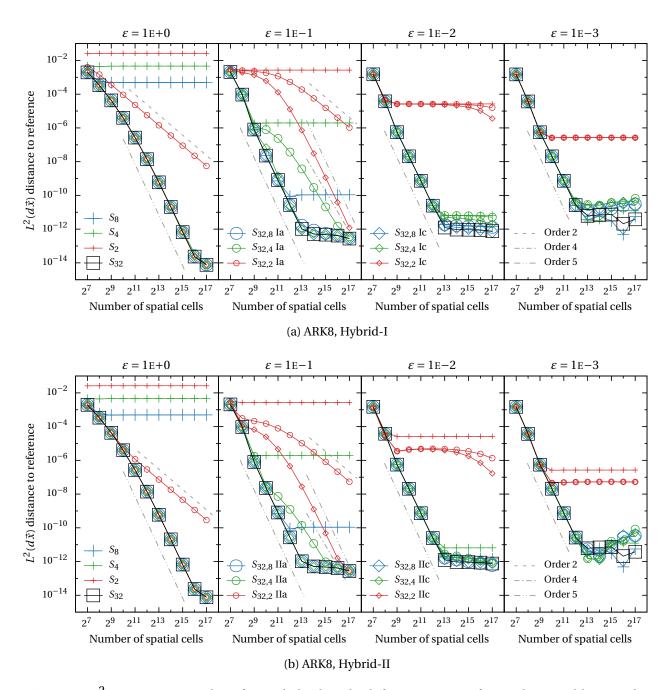


Figure 4.5: L^2 convergence studies of ARK8 hybrid methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

the order-reduction phenomenon indicate the existence of superconvergence behavior in the angular moments of the hybrid solution with respect to the length of time that elapses between applications of the reconstruction operation. The results suggest that the truncation error in the zeroth angular moment corresponds to a global convergence rate of $2N_{\rm C}+1$. Since the values obtained through reconstruction procedures based on polynomial interpolation contain contributions from all of the angular moments of the collided component, the superconvergence behavior of the zeroth moment is obscured by the lower convergence rates of the higher-degree angular moments. In contrast, the Nyström reconstruction procedures use only the zeroth angular moment of the collided component, and hence the superconvergence behavior in the zeroth moment is preserved during reconstruction. The fact that the truncation error of the hybrid methods using Nyström reconstruction procedures is observed to be significantly higher than that of previous methods alleviates concerns over order reduction for most practical purposes. However, the issue of convergence saturation remains.

Overall, the use of Nyström reconstruction procedures improves the accuracy of the hybrid methods by a significant margin. Whereas the $S_{32,4}$ hybrid methods using reconstruction procedures based on polynomial interpolation displayed increases in error over the reference S_{32} approximations in a small number of tests (specifically, for fifth-order methods and $\varepsilon = 1\text{E}-1$), the $S_{32,4}$ hybrid methods using Nyström reconstruction procedures do not display any discernible increase in error over the reference S_{32} approximation. While saturation is still a concern for $S_{32,2}$ hybrid methods, accuracy in regimes with low to moderate collision rates is significantly improved due to the fact that the Nyström reconstruction procedures yield truncation errors with a much higher order of accuracy than reconstruction procedures based on polynomial interpolation.

4.3.2 Two spatial dimensions

Hybrid methods using Nyström reconstruction procedures are now applied to the hohlraum test problem in order to evaluate their computational efficiency. The problem setup remains the same as before; i.e., the final time, CFL, and spatial meshes used are the same as those outlined in section 3.4.2. Results are presented in three forms as before, with minor modifications to each presentation.

First, we show plots comparing the scalar flux density and the error distribution of standard discrete ordinates methods and hybrid discrete ordinates methods using Nyström reconstruction proce-

dures and a selection of angular quadrature resolutions. For CL quadratures, standard discrete ordinates methods using an angular quadrature of order 64 are compared with hybrid discrete ordinates methods using an angular quadrature of order 64 in the uncollided component and angular quadratures of various resolutions in the collided component. For T_N quadratures, standard discrete ordinates methods using an angular quadrature of order 32 are compared with hybrid discrete ordinates methods using an angular quadrature of order 32 in the uncollided component and angular quadratures of various resolutions in the collided component. Results for DIRK methods are shown in figures 4.6 to 4.13 and results for STDG methods are shown in figures 4.14 and 4.15. The solve time (in seconds) and an approximation of the L^{∞} error are included for each solution.

As before, we observe that hybrid methods require significantly less solve time than the standard discrete ordinates methods that they are compared to, and the reduced solve times are accompanied by increases in solution error when very coarse angular approximations are applied to the collided component of the hybrid methods. However, the structure of the additional error introduced by the hybrid splitting is substantially different when Nyström reconstruction procedures are used than was observed in the hybrid methods in chapter 3. Whereas the use of lower-resolution angular quadratures in the collided component of the hybrid methods in chapter 3 introduced prominent ray effects in the scalar flux distributions, it is notable that no visibly apparent ray effects due to the lower-resolution angular quadrature can be observed in the scalar flux distributions obtained from hybrid methods using Nyström reconstruction procedures. Indeed, the plots of the error distributions clearly show that the Nyström reconstruction procedures yield errors that are more smoothly distributed over the spatial domain.

The disparities that are observed in the error distributions of the various hybrid methods due to the choice of reconstruction procedure stem from inherent differences between the reconstruction approaches. The reconstruction procedures used in chapter 3 are applied locally at each spatial point and map the values of the collided component to the higher-resolution angular quadrature without applying any scaling to the original or resulting values. Thus any ray effects present in the collided component are transferred into the reconstructed solution. The Nyström reconstruction procedures are composed of two essential components that deviate significantly from this paradigm. First, an approximation of the scattering source is computed using the collided and uncollided components: In doing so, the values of the collided component are scaled by the material cross section in each region. For regions with very low collision probabilities (where the low-resolution angular quadrature would be expected to exhibit

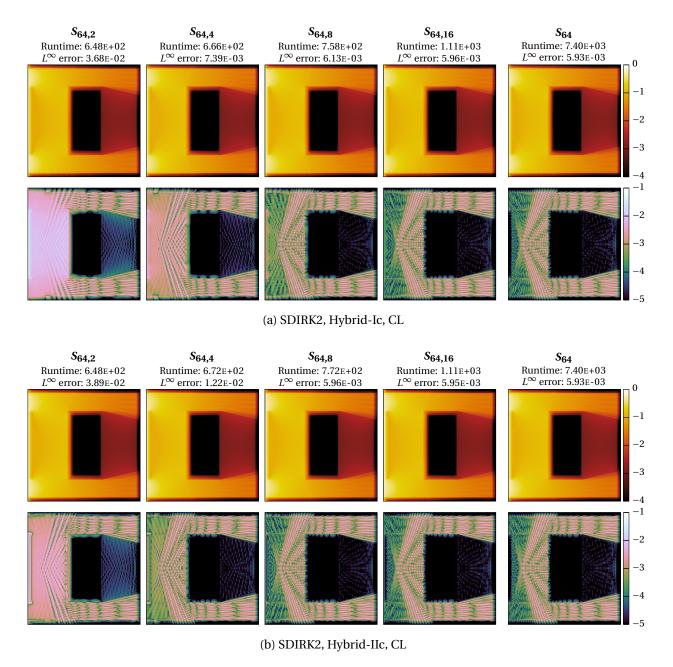


Figure 4.6: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK2 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

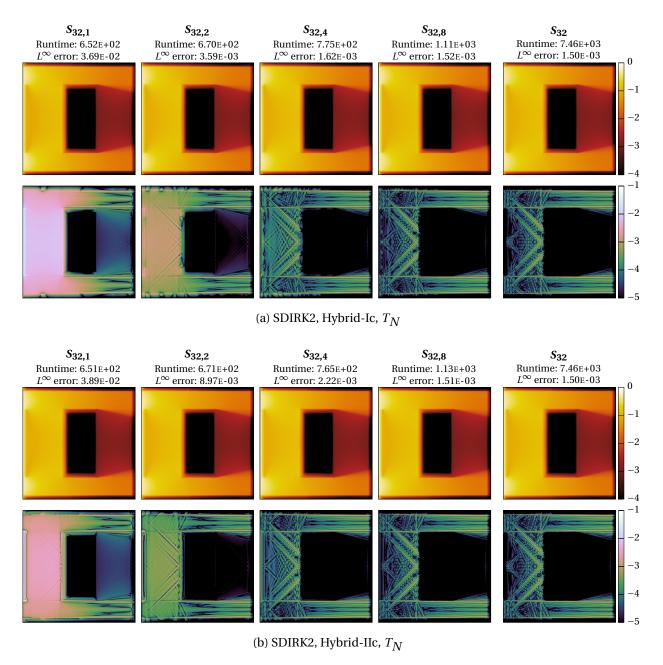


Figure 4.7: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK2 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution.

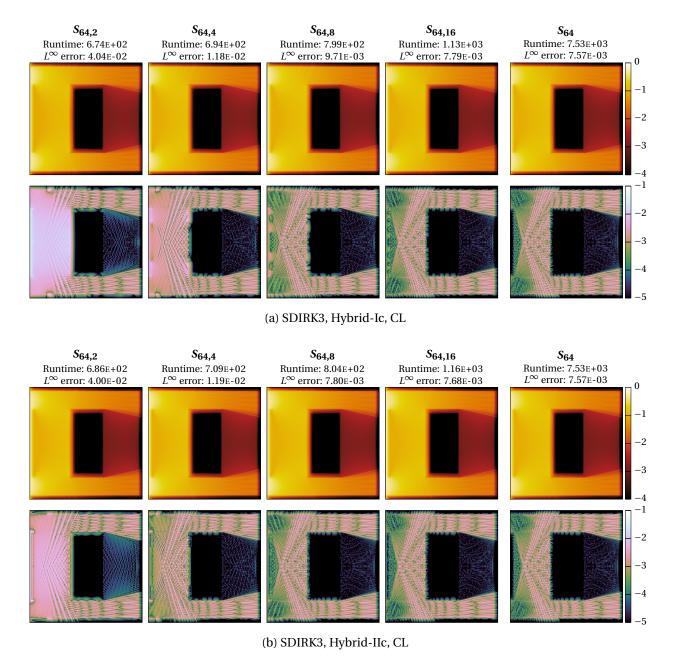


Figure 4.8: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK3 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

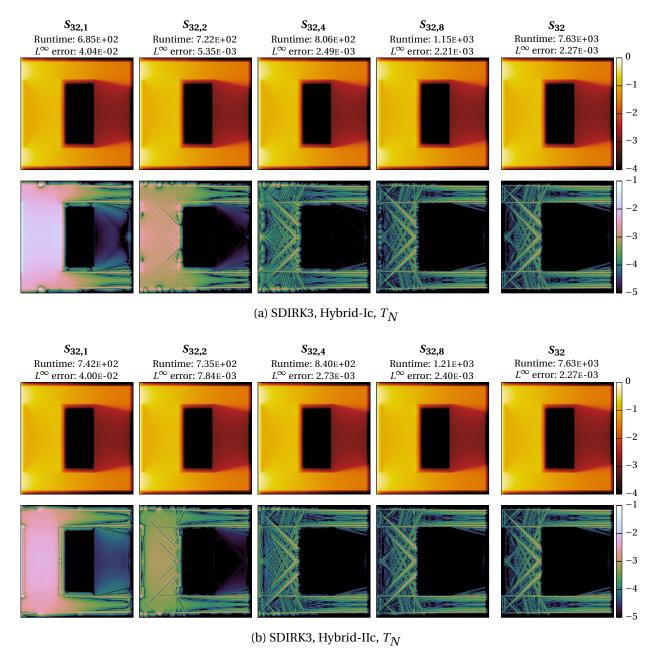
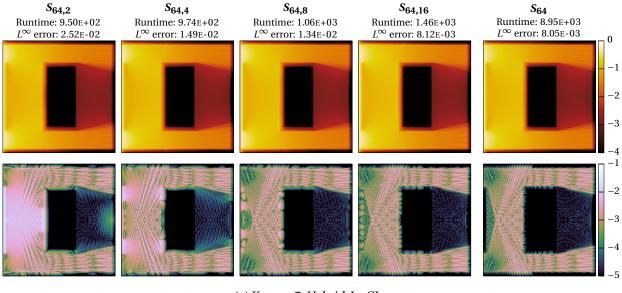
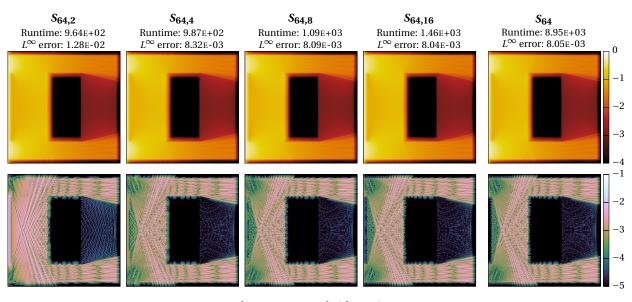


Figure 4.9: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid SDIRK3 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

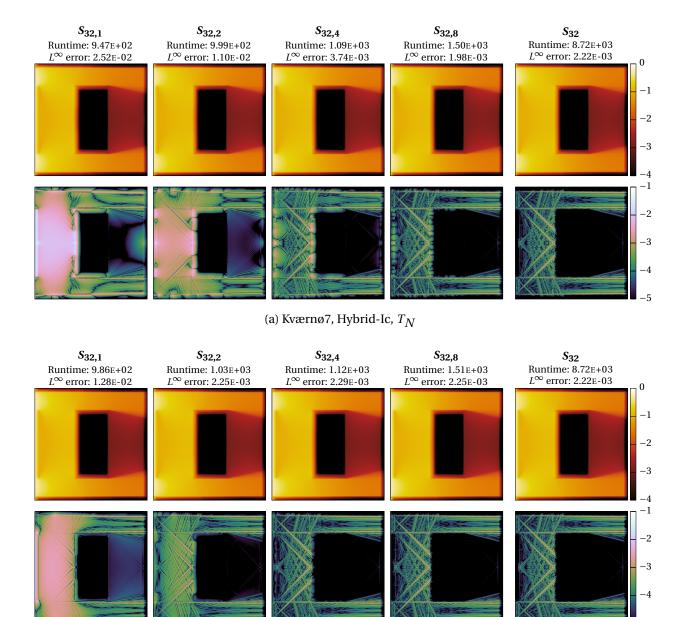


(a) Kværnø7, Hybrid-Ic, CL



(b) Kværnø7, Hybrid-IIc, CL

Figure 4.10: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid Kværnø7 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.



(b) Kværnø7, Hybrid-IIc, T_N

Figure 4.11: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid Kværnø7 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution.

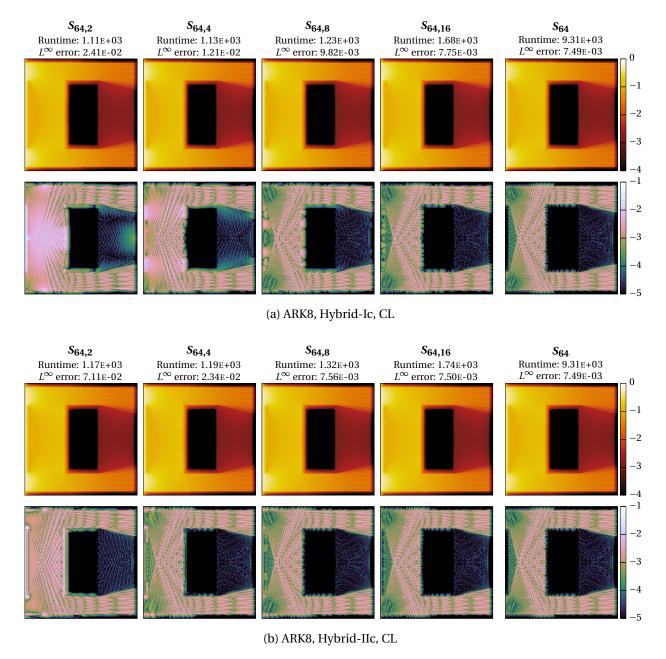


Figure 4.12: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid ARK8 methods using CL quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

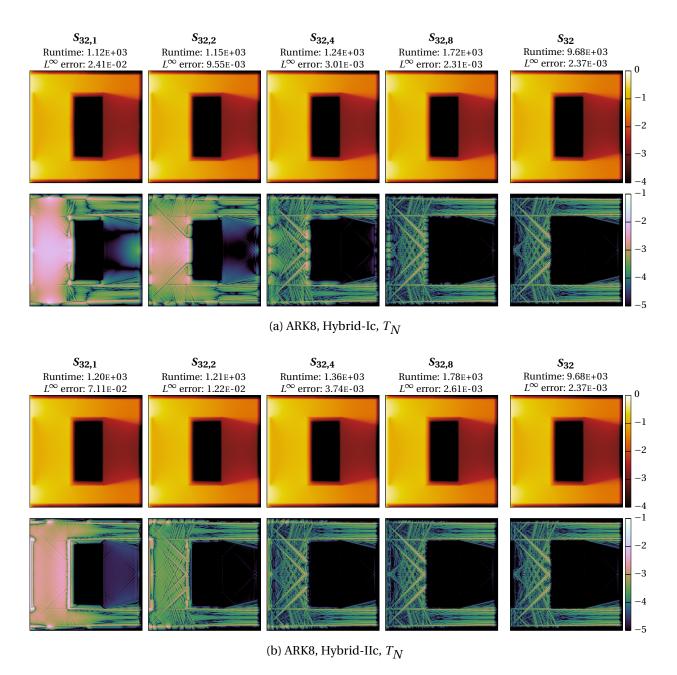


Figure 4.13: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid ARK8 methods using T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

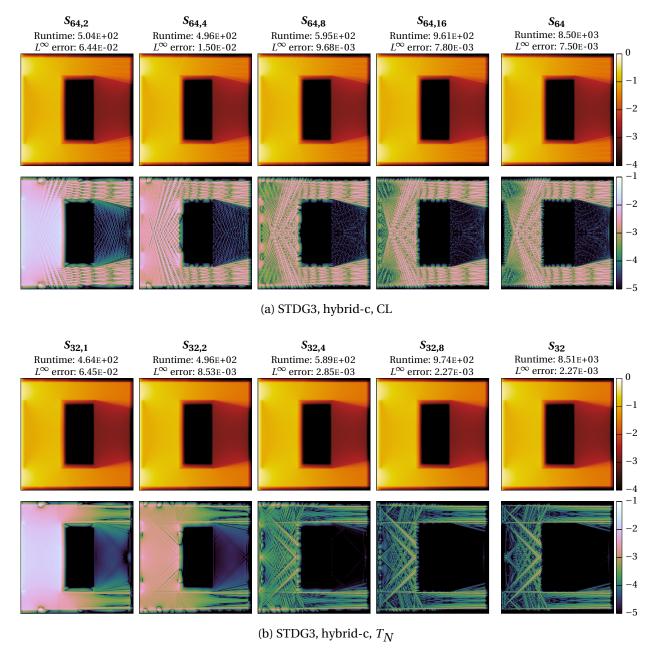


Figure 4.14: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid STDG3 methods using CL and T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

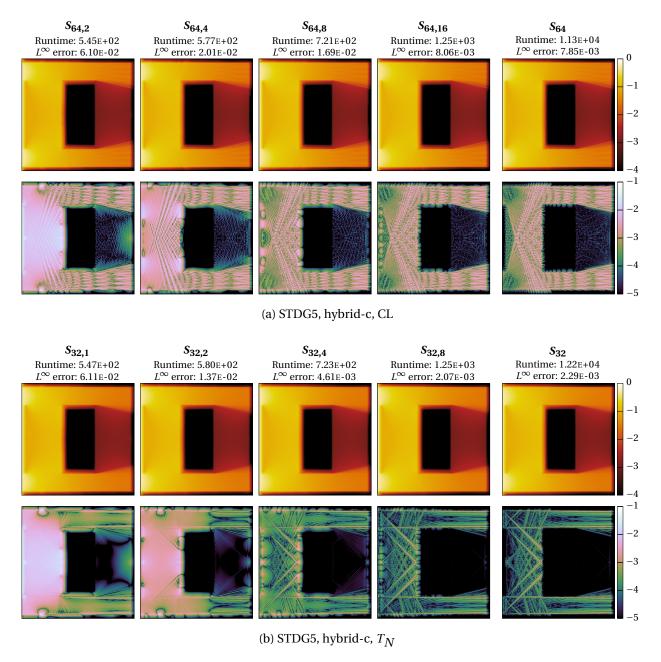


Figure 4.15: Comparison of scalar flux density (top rows) and deviation of scalar flux density from reference solution (bottom rows) of hybrid and non-hybrid STDG5 methods using CL and T_N quadratures applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

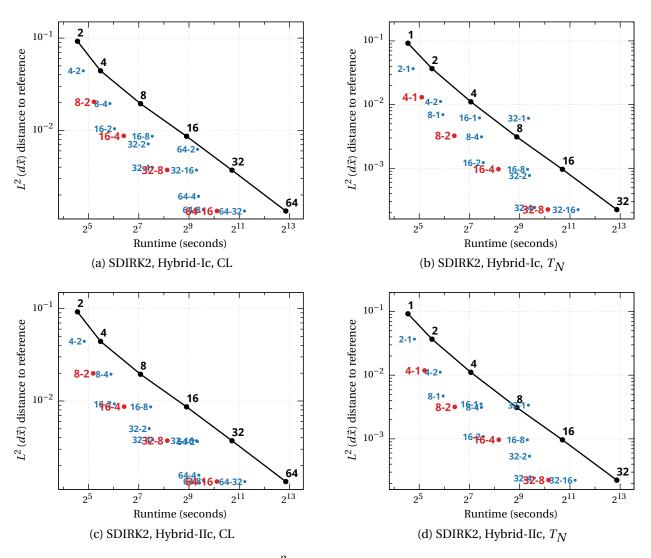


Figure 4.16: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the SDIRK2 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

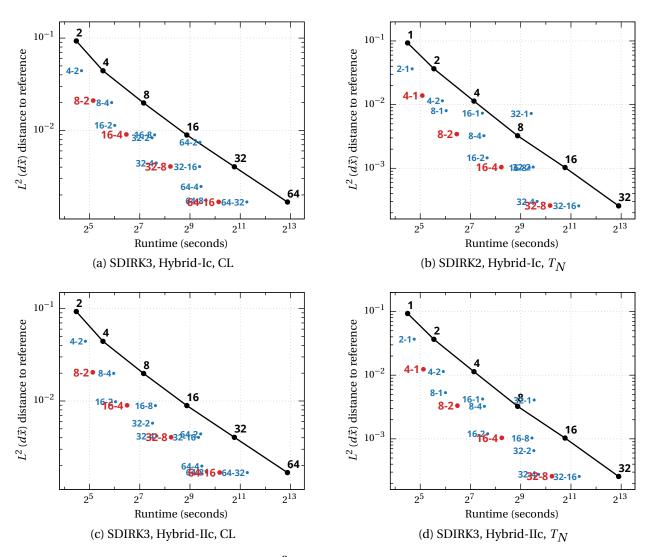


Figure 4.17: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the SDIRK3 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

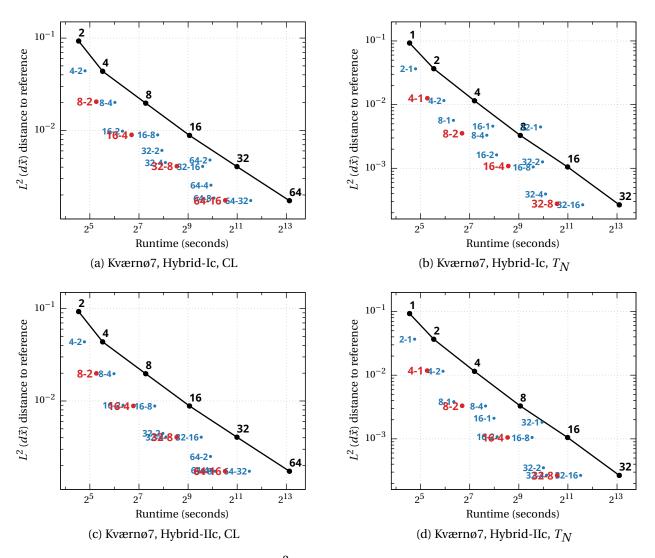


Figure 4.18: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the Kværnø7 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

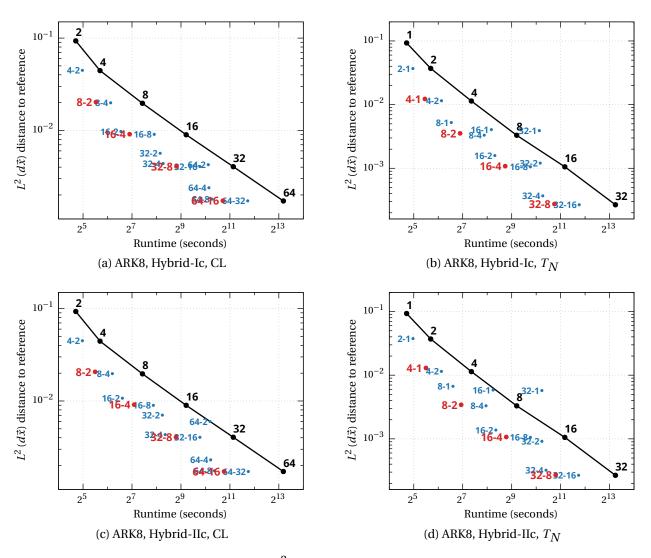


Figure 4.19: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on the ARK8 integrator applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

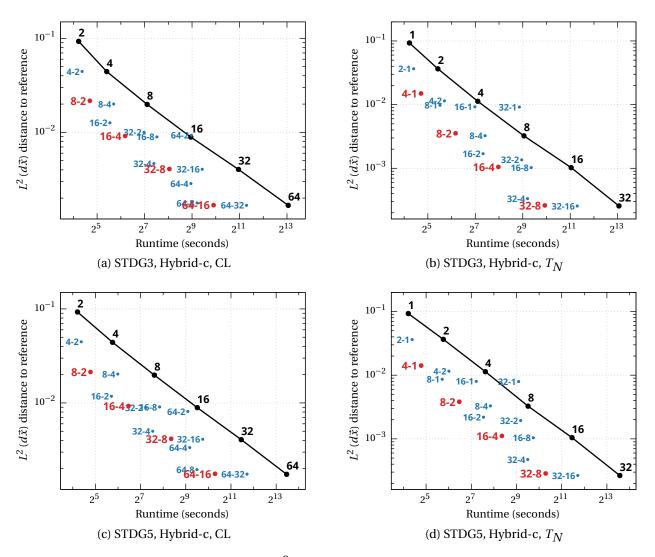


Figure 4.20: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference for standard (black) and hybrid (red and blue) discrete ordinates methods based on STDG3 and STDG5 integrators applied to the holraum test problem. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue. Note the differences in scales.

ray effects) this scaling has a damping effect. The scattering source is therefore concentrated in regions with higher collision probabilities where the low-resolution angular quadrature is likely to be sufficient for acceptable accuracy. Hence, the amount of ray effects present in the scattering source is expected to be limited. The computed approximation of the scattering source can nevertheless be inaccurate due to ray effects in neighboring regions even though ray effects may not be present in the scattering source itself. The second component of the Nyström reconstruction procedures is the inversion of the high-resolution streaming operator via a transport sweep. This operation is non-local and distributes particles from the approximate scattering source – and any errors contained therein – throughout the spatial domain. Because this operation is computed using the higher-resolution angular quadrature, the amount of ray effects introduced in this step is limited, resulting in the smoother error distributions that are observed.

In addition to the differences observed in the overall structure of the error distributions, the Nyström reconstruction procedures also yield solutions that are overall more accurate than the reconstruction procedures considered in chapter 3. For example, the $S_{32,2}$ hybrid-Ic DIRK methods using T_N quadratures yield solutions with L^{∞} errors that are smaller than those of the corresponding hybrid-Ia DIRK methods by factors of between 7 and 13. On the other hand, the Nyström reconstruction procedures are generally more expensive than the procedures used in chapter 3. When Nyström reconstruction procedures are used, two transport sweeps must be performed for each stage using the higher-resolution angular quadrature: one to compute the uncollided component and another to computed the reconstructed solution. This means that hybrid methods using Nyström reconstruction procedures require at most twice the solve time of the hybrid methods described in chapter 3.

Comparisons of the computational efficiency of hybrid methods using Nyström reconstruction procedures for a broader range of angular resolutions are shown for DIRK methods in figures 4.16 to 4.19 and for STDG methods in figure 4.20. As before, we observe that hybrid methods using angular quadratures that differ in order by a factor of 4 tend to consistently lead to good computational efficiency relative to standard discrete ordinates methods and other hybrid methods. However, since using a Nyström reconstruction procedure significantly increases the overall accuracy of the hybrid solutions, the efficiency comparisons that can be made against standard discrete ordinates methods are similarly improved. In order to attain a level of accuracy comparable to or better than that of a standard discrete ordinates method using an angular quadrature of order N, the hybrid methods in chapter 3 required

that an angular quadrature of order 2N be used in the uncollided component. This is no longer the case when Nyström reconstruction procedures are used: We observe that hybrid methods using Nyström reconstruction procedures and angular quadratures of order N and N/4 in the uncollided and collided components, respectively, yield solutions with comparable errors in less time than a standard discrete ordinates method using an angular quadrature of order N. These hybrid methods reduce solve times by factors of between 4 and 8 for the hohlraum test problem, with larger reductions observed when higher-resolution angular quadratures are used.

We note that using angular quadratures that differ in order by a factor of 8 may also be an acceptable choice in certain cases. However, the variability observed in the efficacy of these methods prevents us from issuing a general recommendation for their use. When the angular quadratures applied to the components of the hybrid methods differ in order by a factor of 8 or more, the vast majority of the overall solve time of the methods is spent on the uncollided component. Hence further coarsening the resolution of the angular quadrature that is applied to the collided component yields diminishing returns in computational cost savings. Weighing this fact against the potential uncertainty in the accuracy of these methods, we find that using angular quadratures that differ in order by a factor of 4 is likely a better choice in most cases.

In chapter 3 it was observed that hybrid-IIa DIRK methods produced significantly more accurate solutions than hybrid-Ia DIRK methods. The difference in accuracy between hybrid-IIc and hybrid-Ic DIRK methods does not appear to be quite as significant. When angular quadratures that differ in order by a factor of 4 or less are used, hybrid-IIc and hybrid-Ic DIRK methods yield solutions with comparable levels of accuracy. When angular quadratures that differ in order by a factor of more than 8 are used, hybrid-IIc DIRK methods yield marginally smaller errors than hybrid-Ic DIRK methods. However, the difference in accuracy between hybrid-IIc and hybrid-Ic DIRK methods is much less than the difference that was observed between hybrid-IIa and hybrid-Ia DIRK methods.

Comparisons between standard discrete ordinates methods using CL and T_N quadratures and hybrid discrete ordinates methods using CL and T_N quadratures that differ by a factor of 4 are shown in figures 4.21 to 4.23. Hybrid methods using Nyström reconstruction procedures and the reconstruction procedures described in chapter 3 are included. In each case, using a Nyström reconstruction approach improves the accuracy of the hybrid methods. While the difference is comparatively small for hybrid-II DIRK methods, the improvements observed for STDG hybrid methods and hybrid-I DIRK methods are

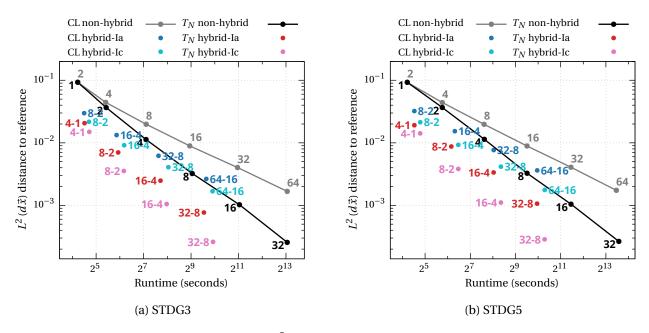


Figure 4.21: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray), hybrid-a (red and blue), and hybrid-c (pink and teal) discrete ordinates methods based on STDG integrators and using various angular quadrature resolutions applied to the hohlraum test problem. CL (gray, blue, and teal) and T_N (black, red, and pink) angular quadratures are considered.

quite significant. It is clear that the computational cost of applying the Nyström reconstruction procedures is appreciably larger than that of the reconstruction procedures used in chapter 3. Hybrid methods using Nyström reconstruction procedures require between 15 to 55 percent more solve time than previous methods when angular quadratures that differ in order by a factor of four are used. However, the increase in computational cost is more than outweighed by the increase in accuracy observed in most cases.

Nyström reconstruction procedures improve the accuracy and computational efficiency of the hybrid methods by a clear margin. As a result, the effective memory requirements of the hybrid methods are significantly reduced. The recommendation given in chapter 3 that a standard discrete ordinates method using an angular quadrature of order N be replaced with a hybrid discrete ordinates method using angular quadratures of order 2N and N/2 for the uncollided and collided components, respectively, yields an increase in memory consumption by a factor of approximately 4.25. The revised recommendations given for hybrid methods using Nyström reconstructions – that angular quadratures of order N and N/4 be used for the uncollided and collided components, respectively – requires only a minor increase in memory footprint over standard discrete ordinates methods. Since the resolution of the

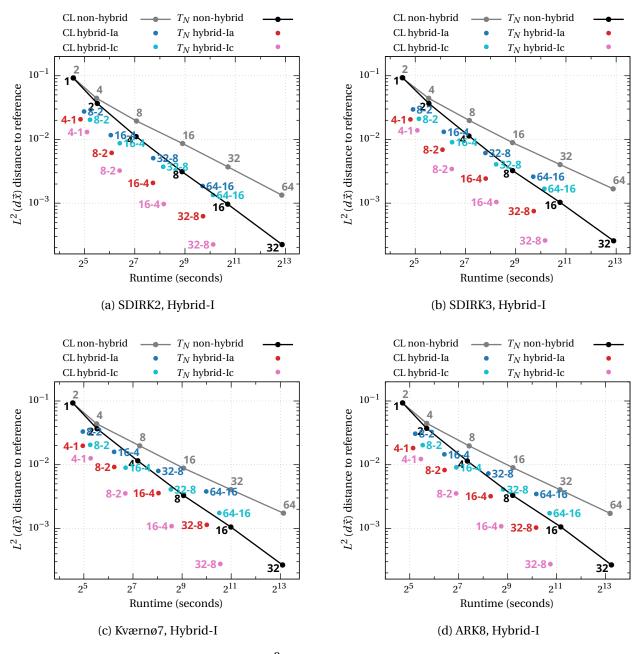


Figure 4.22: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray), hybrid-Ia (red and blue), and hybrid-Ic (pink and teal) discrete ordinates methods based on DIRK integrators and using various angular quadrature resolutions applied to the hohlraum test problem. CL (gray, blue, and teal) and T_N (black, red, and pink) angular quadratures are considered.

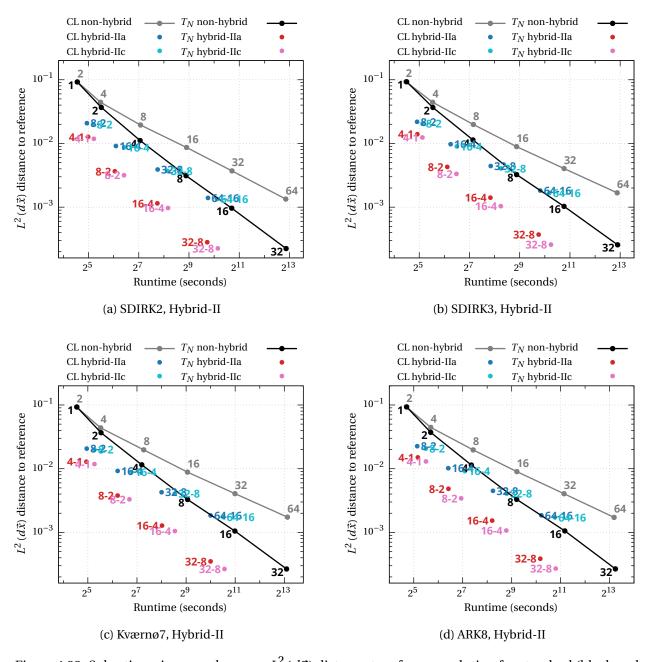


Figure 4.23: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black and gray), hybrid-IIa (red and blue), and hybrid-IIc (pink and teal) discrete ordinates methods based on DIRK integrators and using various angular quadrature resolutions applied to the hohlraum test problem. CL (gray, blue, and teal) and T_N (black, red, and pink) angular quadratures are considered.

angular quadrature used for the uncollided component is the same as that of the corresponding standard discrete ordinates method, storing the uncollided components of the stages requires no increase in memory footprint. The only increase comes from storing the collided components of the stages, which requires one-sixteenth the memory of the corresponding non-hybrid method. The result is that, for each time integrator, replacing a standard discrete ordinates method with a hybrid discrete ordinates method using Nyström reconstruction increases memory consumption by approximately 6 percent.

4.4 Numerical tests for two-grid methods

The same set of tests used in sections 3.4 and 4.3 are repeated for two-grid methods based on IDC and LS-IDC integrators. Descriptions of the problems and test procedures are not repeated here – see section 3.4 for detailed descriptions of these tests.

4.4.1 One spatial dimension

We first examine the convergence behavior of the two-grid methods. Convergence tests for two-grid methods are presented in a format that is different than that of previous convergence tests. In previous convergence tests, results for one or more hybrid methods based on a single time-integration scheme using angular quadratures of order 2, 4, and 8 in the collided component were presented in each figure. Here each figure shows results for a single time integrator, hybrid method, and fixed angular resolution with varying numbers of correction iterations. While all of the hybrid methods shown in these figures are computed using IDC5 or LS-IDC5 integrators, the standard discrete ordinates methods that are included are computed using the collocation integrator associated with each defect correction method, which corresponds to the STDG5 scheme in both cases.

As additional correction iterations are performed, the solutions obtained from the IDC methods would be expected to converge to the solutions obtained from the associated collocation scheme. However, the order of accuracy of the IDC methods is expected to be less than that of the collocation scheme until sufficiently many correction iterations have been performed. In this case, one prediction and four correction iterations are required in order to observe fifth-order accuracy from the IDC methods. As discussed in section 2.3.3, the linear stability of IDC methods is generally somewhat less robust than that of the associated collocation scheme, and additional correction iterations can improve the stability of the methods and diminish the extent of any order reduction that may occur. This is particularly true for

LS-IDC methods, which tend to converge more slowly towards the collocation solution than traditional IDC methods.

The colors and styles of point marks used for hybrid methods in these convergence plots are selected based on the above observations. Hybrid methods using four and five correction iterations are the focus of these convergence plots, and hence a large starburst symbol has been used to emphasize these methods: Point marks for methods using other numbers of correction iterations are denoted using circles of different sizes. The color that is used for each point mark emphasizes which solutions would ideally have accuracy comparable to that of the associated collocation solution. Since the IDC5 and LS-IDC5 methods attain their maximum theoretical order of accuracy after four correction iterations, methods using more than four correction iterations should have accuracy similar to that of the collocation scheme unless their accuracy has been inhibited in some manner. Therefore, a red color is used for point marks of hybrid methods using more than four correction iterations, and a blue color is used for point marks of hybrid methods using four or fewer correction iterations. The result is that, unless the accuracy of the methods has been inhibited in some way, point marks for hybrid methods should satisfy the following conditions: (i) red-colored marks should lie mostly or entirely within the black squares, and (ii) starburst marks of any color should lie close to the black squares, though it would not necessarily be expected that they lie within the black squares.

We also note that a different norm is used to compute solution errors than was used in previous convergence tests. Errors in previous convergence tests were measured using the scalar flux of the solution with respect to the norm (2.52) at the final simulation time. Here errors are measured using the angular flux of the solution with respect to the semi-discrete $L^2(d\vec{x}d\vec{\Omega})$ norm

$$e(t) = \left(\sum_{k=1}^{K_{\mathbf{u}}} \omega_{\mathbf{u},k} \int_{X} \left[\psi\left(\vec{x}, \vec{\Omega}_{\mathbf{u},k}, t\right) - \psi^{\text{ref}}\left(\vec{x}, \vec{\Omega}_{\mathbf{u},k}, t\right) \right]^{2} d\vec{x} \right)^{1/2}, \tag{4.44}$$

where $\{\vec{\Omega}_{\mathrm{u},k},\omega_{\mathrm{u},k}\}_{k=1}^{K_{\mathrm{u}}}$ is the high-resolution angular quadrature used for the reference solution and the uncollided components of the hybrid solutions (S_{32} in this case). The motivation for this is based on the fact that the metric (4.44) tends to be somewhat more sensitive than (2.52). Hence behavior that may otherwise be obscured when considering only the scalar flux should be more readily apparent in these convergence tests.

The results of the convergence tests for methods based on the IDC5 integrator are shown in figures 4.24 and 4.26 and results for methods based on the LS-IDC5 integrator are shown in figures 4.25

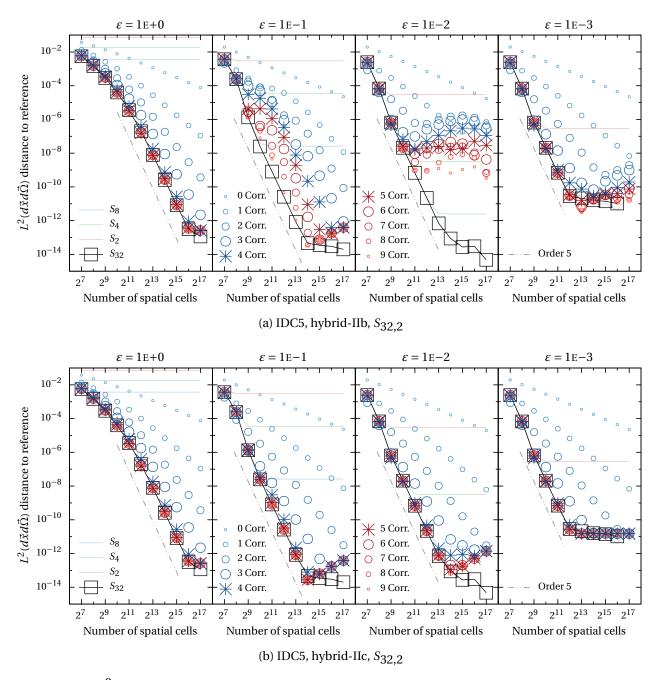


Figure 4.24: L^2 convergence studies of IDC5 hybrid-IIb and hybrid-IIc $S_{32,2}$ methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

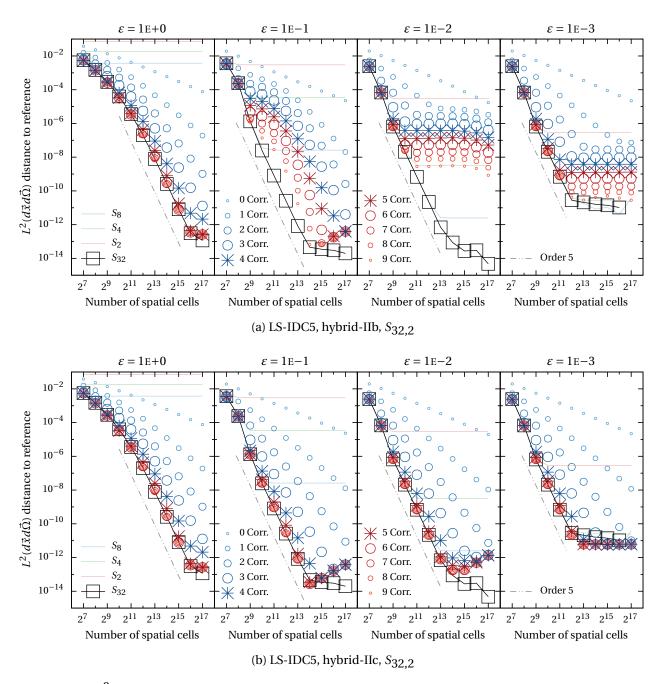


Figure 4.25: L^2 convergence studies of LS-IDC5 hybrid-IIb and hybrid-IIc $S_{32,2}$ methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

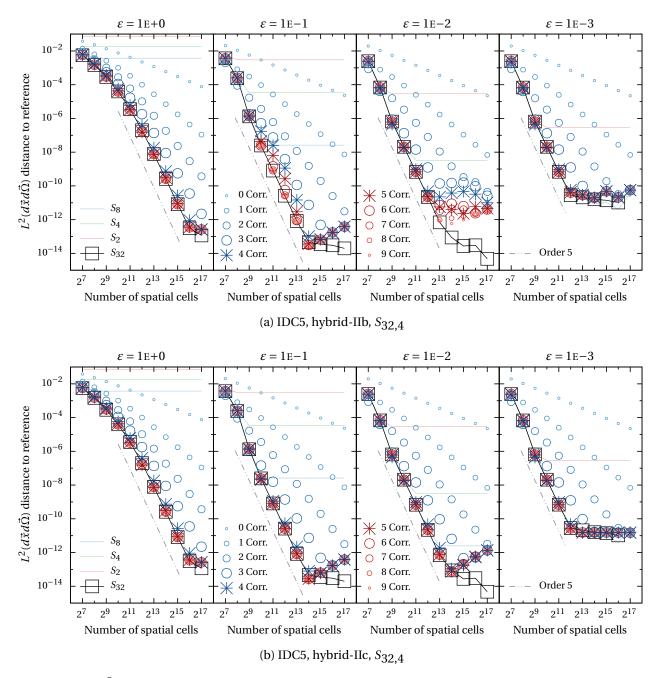


Figure 4.26: L^2 convergence studies of IDC5 hybrid-IIb and hybrid-IIc $S_{32,4}$ methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

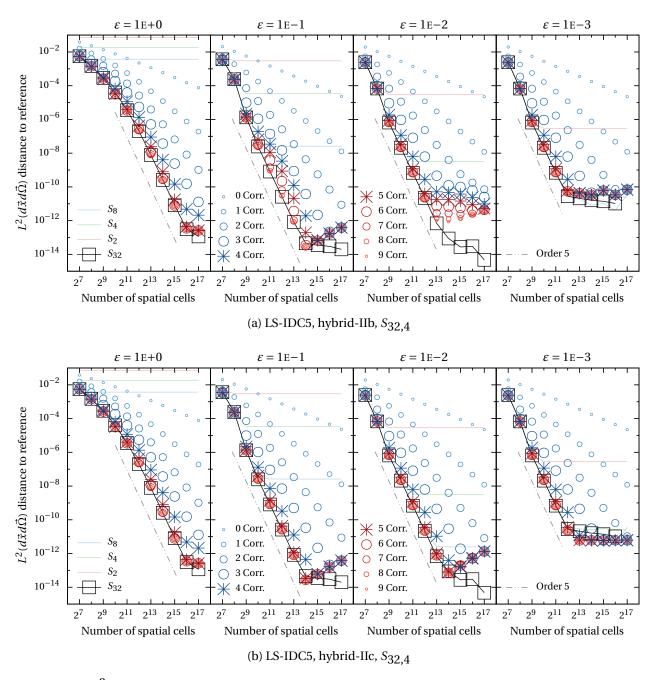


Figure 4.27: L^2 convergence studies of LS-IDC5 hybrid-IIb and hybrid-IIc $S_{32,4}$ methods for a sequence of smooth 1D problems under a diffusive scaling with $\sigma_t = \sigma_s = 1$, spatial domain $z \in [-4,4]$, final time 0.5, and a CFL of 8 for various values of ε approaching 0. The error in the scalar flux of each solution is approximated by comparing to a reference solution using 2^{18} spatial cells and an S_{32} discrete ordinates approximation in angle. Point marks for hybrid solutions whose convergence is unhindered lie within the black boxes.

and 4.27. Overall, it is clear that the two-grid defect correction procedures significantly increase the accuracy of the hybrid methods. The saturation and order reduction that was observed with previous hybrid methods are both significantly reduced by the two-grid approach.

The degree of improvement that is observed strongly depends on the method. For the hybrid-IIb methods, the order reduction phenomenon that was observed in previous hybrid methods is effectively eliminated. Whereas reduction to second-order accuracy was observed with previous $S_{32,2}$ hybrid methods using reconstruction procedures based on polynomial interpolation, the hybrid-IIb IDC and LS-IDC methods are able to recover the expected fifth-order convergence rate. Convergence saturation is still observed in the hybrid-IIb methods, though the error level at which this phenomenon occurs in the hybrid-IIb IDC and LS-IDC methods is significantly smaller than was observed for previous hybrid methods. Reductions in the saturation error of the $S_{32,2}$ hybrid-IIb methods as additional correction iterations are performed can be clearly observed in figures 4.24a and 4.25a. However, this saturation persists very strongly in many tests even when more than four correction iterations are used, and the amount by which the error in the saturation regime is reduced for each additional correction iteration is rather small. It is interesting to note that the convergence behavior of the hybrid-IIb methods in the saturation regime is distinctly different for IDC and LS-IDC integrators. In particular, the convergence of the hybrid-IIb IDC5 method is non-monotonic and somewhat erratic, while the behavior of the hybrid-IIb LS-IDC5 method is much more consistent. At this time it is not clear what properties of the IDC and LS-IDC methods may affect this behavior, though we note that the erratic behavior of the IDC5 method does not seem unreasonable given the fact that it is known that the two-grid method on which the hybrid-II methods are based does not converge monotonically (c.f. section 4.1.2 and [17]).

We observe that the hybrid-IIc methods are overall much more accurate than the hybrid-IIb methods. A significant degree of convergence saturation is observed for $S_{32,2}$ hybrid-IIb methods, and a moderate amount of saturation is observed for $S_{32,4}$ hybrid-IIb methods below very small tolerances. In contrast, the $S_{32,2}$ hybrid-IIc methods show errors that could indicate a very small amount of saturation behavior for only two data points (for $\varepsilon = 1\text{E}-1$ and mesh sizes of 2^{11} and 2^{12}), and the $S_{32,4}$ hybrid-IIc methods do not display any discernible measure of saturation behavior in any of these tests. It is interesting to note that the advantage of using a Nyström reconstruction procedure for hybrid methods based on DIRK and STDG integrators was rather different than what is observed here. For hybrid methods based on DIRK and STDG integrators, the advantage of using a Nyström reconstruction pro-

cedure was isolated to the order-reduction behavior of the methods – the saturation behavior remained nearly identical. For hybrid-II IDC and LS-IDC methods, the two-grid approach eliminates the impact of the order reduction phenomenon rather effectively irrespective of the reconstruction procedure that is used, and the Nyström reconstruction procedure seems to primarily serve to reduce the impact of the saturation phenomenon.

Comparing the IDC and LS-IDC methods, we observe differences in overall accuracy similar to what was observed in section 2.4.1. Specifically, the LS-IDC methods generally produce larger errors for a given number of correction iterations than the equivalent IDC method. We further observe that: (i) ignoring the impact of the saturation phenomenon, the difference in accuracy between the IDC and LS-IDC methods is typically rather small, and (ii) the difference in accuracy seems to vary depending on the value of ε , with larger differences observed for larger values of ε . In every case, an LS-IDC method using p+1 correction iterations yields an error comparable to or better than that of the corresponding IDC method using p correction iterations. The increased computational cost of performing an additional correction iteration is likely to be a worthwhile penalty in certain applications in exchange for the significantly decreased memory requirements of the LS- IDC methods. This penalty is also expected to be much less pronounced in cases where other considerations, such as the use of a combined space-time defect correction approach, may motivate the use of additional correction iterations to reduce the error contributions from components other than the time discretization.

4.4.2 Two spatial dimensions

Two-grid methods are now applied to the hohlraum test problem in order to evaluate their computational efficiency. The problem setup remains the same as before; i.e., the final time, CFL, and spatial meshes used are the same as those outlined in section 3.4.2. As before, results are presented in three forms, but modifications are made to each collection of figures to better illuminate the characteristics of the two-grid methods.

First, we show plots comparing the error distribution of standard discrete ordinates methods and two-grid discrete ordinates methods using a selection of angular quadrature resolutions. For CL quadratures, standard discrete ordinates methods using an angular quadrature of order 64 are compared with two-grid discrete ordinates methods using an angular quadrature of order 64 in the uncollided compo-

nent and angular quadratures of various resolutions in the collided component. For T_N quadratures, standard discrete ordinates methods using an angular quadrature of order 32 are compared with two-grid discrete ordinates methods using an angular quadrature of order 32 in the uncollided component and angular quadratures of various resolutions in the collided component. As in the convergence tests in the preceding section, we present results for IDC and LS-IDC methods using a range of correction iterations.

Results are shown for hybrid methods based on the IDC3 scheme in figures 4.28 to 4.31, for hybrid methods based on the LS-IDC3 scheme in figures 4.32 to 4.35, for hybrid methods based on the IDC5 scheme in figures 4.36 to 4.39, and for hybrid methods based on the LS-IDC5 scheme in figures 4.40 to 4.43. In these figures, each of the four leftmost columns shows the error distribution for one of the two-grid methods and the rightmost column shows the error distribution for the corresponding standard discrete ordinates method. The solutions in each row are computed using the specified number of correction iterations.

Many of the same characteristics observed for the hybrid methods examined previously can be observed here as well. Specifically, we observe that the hybrid methods require significantly less solve time than the standard discrete ordinates methods using the same number of correction iterations, and the reduced solve times are accompanied by increases in solution error when very coarse angular approximations are applied to the collided component of the hybrid methods and only a small number of correction iterations are used. We also observe that using a Nyström reconstruction procedure yields an error profile that is more smoothly distributed over the spatial domain and does not contain the ray effects from the low-resolution angular quadrature that are found in the solutions of methods using other reconstruction approaches.

In each case, the error of the hybrid methods is visibly reduced as additional correction iterations are performed, and many methods are able to recover most of the accuracy of the standard discrete ordinates approximation that they are compared to within several correction iterations. However, this reduction in error is not accompanied by a change in the overall character of the error distribution. For example, ray effects are still present in the hybrid-IIb solutions due to the lower-resolution angular quadrature that is applied to the collided component, though they are reduced in magnitude after each iteration. Similarly, the additional error in the hybrid-IIc solutions remains rather smoothly distributed over the spatial domain. We also observe that the amount by which each correction iteration reduces each solution's error

is different for IDC and LS-IDC schemes. In particular, the error for IDC schemes is often marginally smaller than for LS-IDC schemes. However, as in the convergence tests in section 4.4.1, the convergence of the IDC schemes is not consistent: For example, in figure 4.38 the L^{∞} error increases between one and two correction iterations for both the $S_{32,2}$ and $S_{32,4}$ methods. The LS-IDC schemes to not appear to exhibit this phenomenon.

Comparisons of the computational efficiency of hybrid-II methods for a broader range of angular quadrature resolutions are shown for hybrid methods based on the IDC3 scheme in figures 4.44 to 4.47, for hybrid methods based on the IDC5 scheme in figures 4.48 to 4.51, for hybrid methods based on the LS-IDC3 scheme in figures 4.52 to 4.55, and for hybrid methods based on the LS-IDC5 scheme in figures 4.56 to 4.59.

Because the prediction steps are computed using the implicit Euler method, the efficiency results obtained for the prediction steps are similar to what was observed with previous methods. However, the accuracy of the hybrid methods improves significantly as additional correction iterations are performed. These results suggest that using even very low resolution angular quadratures in the collided component of the two-grid methods is reasonable provided that sufficiently many correction iterations are performed. In theory, these methods could be iterated until some prescribed convergence tolerance has been satisfied: this is how traditional two-grid methods for integral equations are used. However, we are interested in asking whether it is possible to avoid applying an arbitrary number of correction iterations. More specifically, we ask in what cases is it reasonable to apply only the minimum number of correction iterations required to attain the maximum order of accuracy possible from the IDC scheme, and no more, without sacrificing accuracy because of the hybrid splitting? We remind the reader that the minimum number of correction iterations required for the IDC3 and LS-IDC3 schemes is two, and the minimum number required for the IDC5 and LS-IDC5 schemes is four.

In previous sections of this work it was observed that hybrid methods using angular quadratures that differ in order by a factor of four consistently produced good computational efficiency relative to standard discrete ordinates methods and other hybrid methods. We observe that this is also the case for the two-grid methods discussed here. In particular, we observe that two-grid methods using the minimum number of correction iterations prescribed by the IDC scheme and angular quadratures of order N and N/4 in the uncollided and collided components respectively, yield solutions with comparable errors in less time than a standard discrete ordinates method using an angular quadrature of order N. This obser-

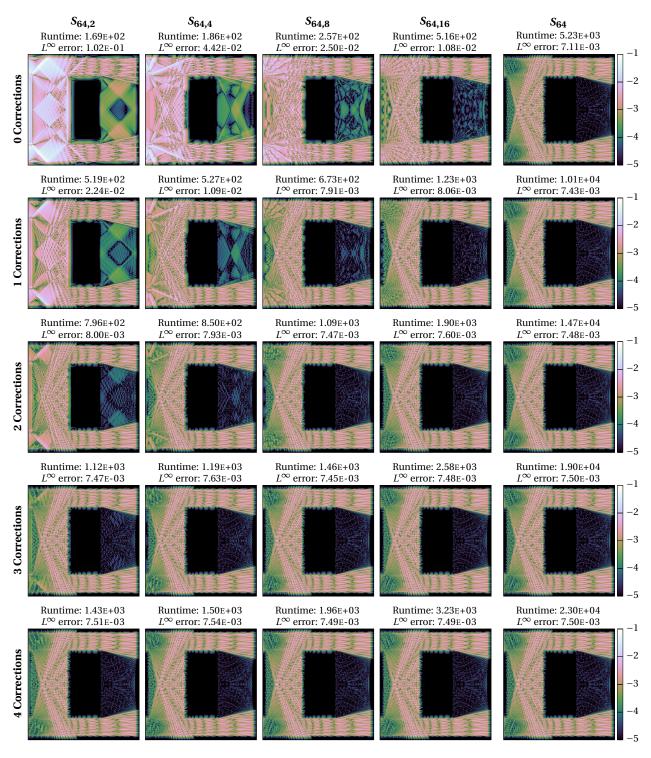


Figure 4.28: Deviation of scalar flux density from reference solution for hybrid-IIb IDC3 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

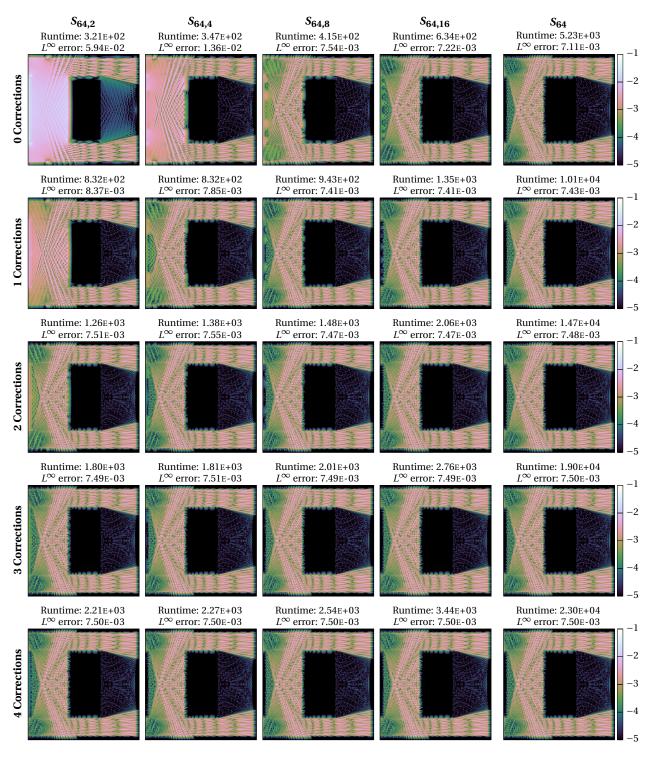


Figure 4.29: Deviation of scalar flux density from reference solution for hybrid-IIc IDC3 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

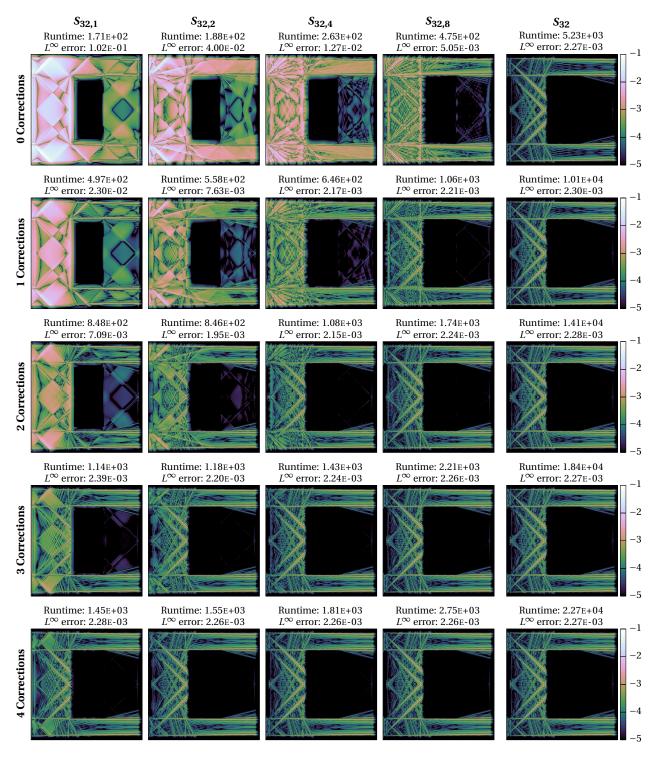


Figure 4.30: Deviation of scalar flux density from reference solution for hybrid-IIb IDC3 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

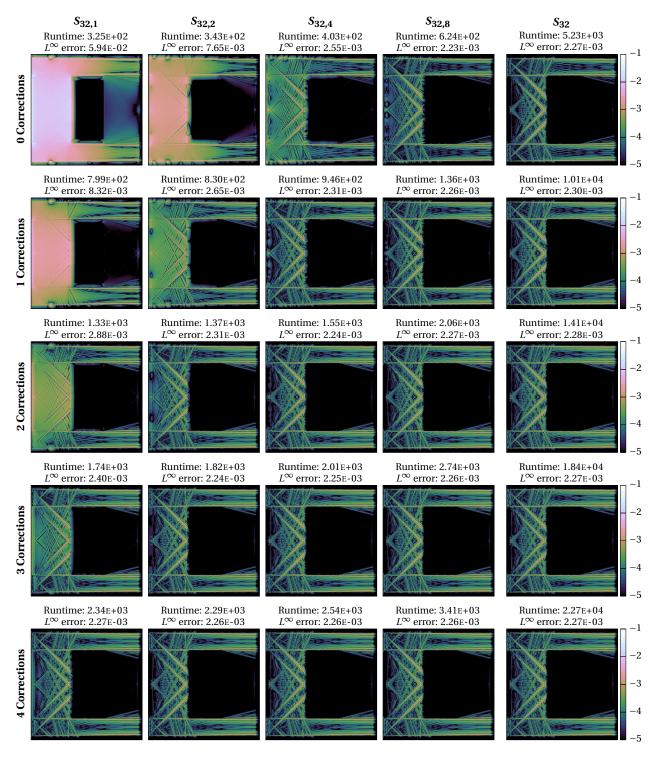


Figure 4.31: Deviation of scalar flux density from reference solution for hybrid-IIc IDC3 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

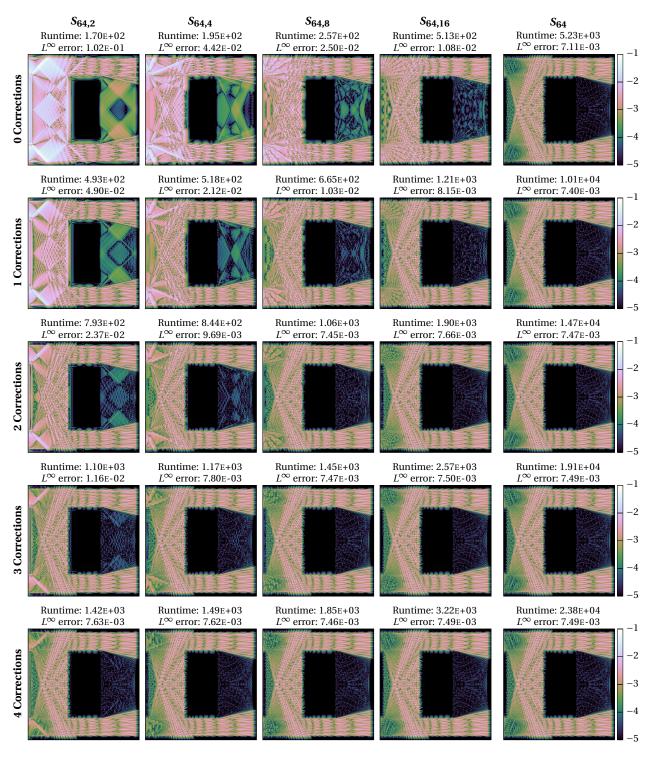


Figure 4.32: Deviation of scalar flux density from reference solution for hybrid-IIb LS-IDC3 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

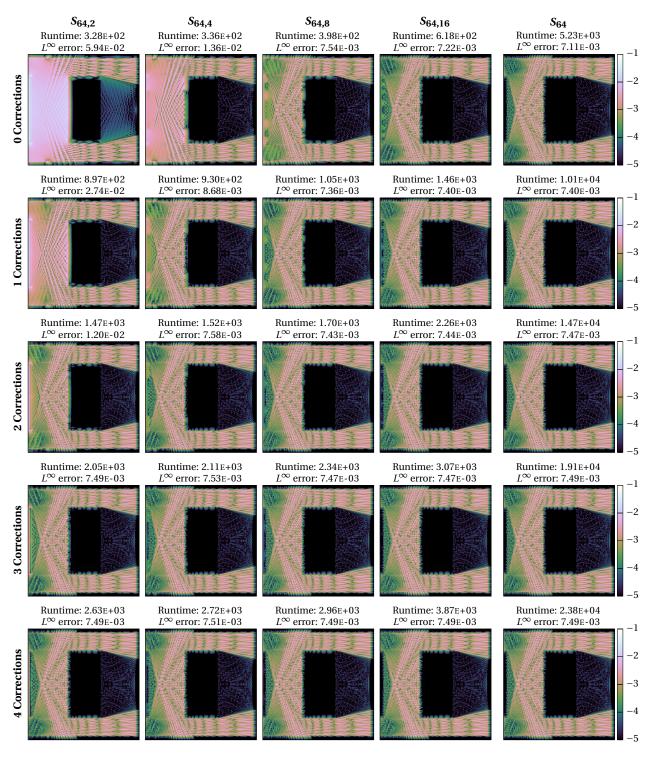


Figure 4.33: Deviation of scalar flux density from reference solution for hybrid-IIc LS-IDC3 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

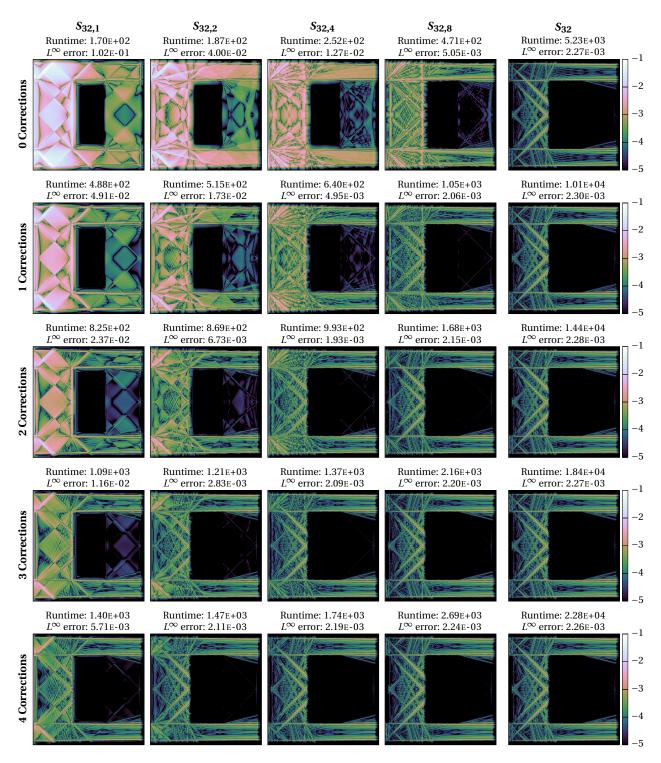


Figure 4.34: Deviation of scalar flux density from reference solution for hybrid-IIb LS-IDC3 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^∞ error are included for each solution.

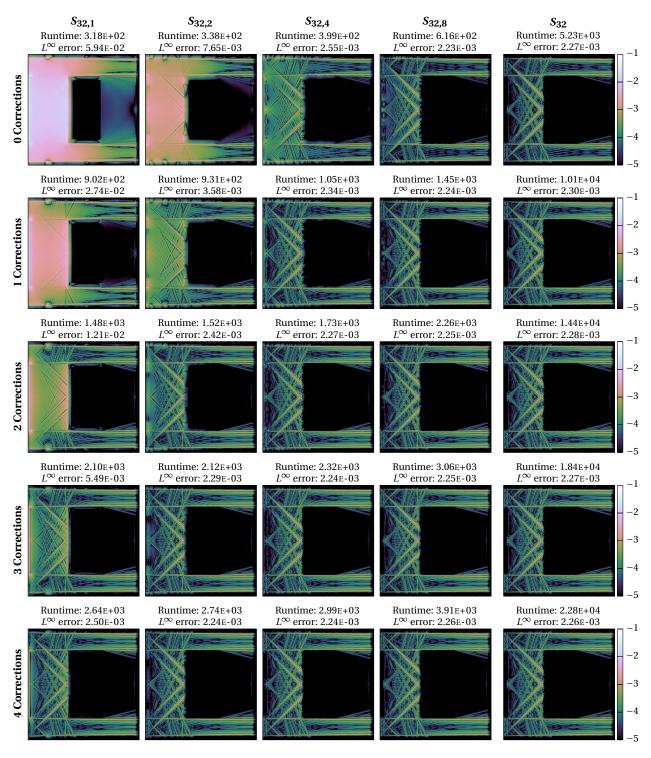


Figure 4.35: Deviation of scalar flux density from reference solution for hybrid-IIc LS-IDC3 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

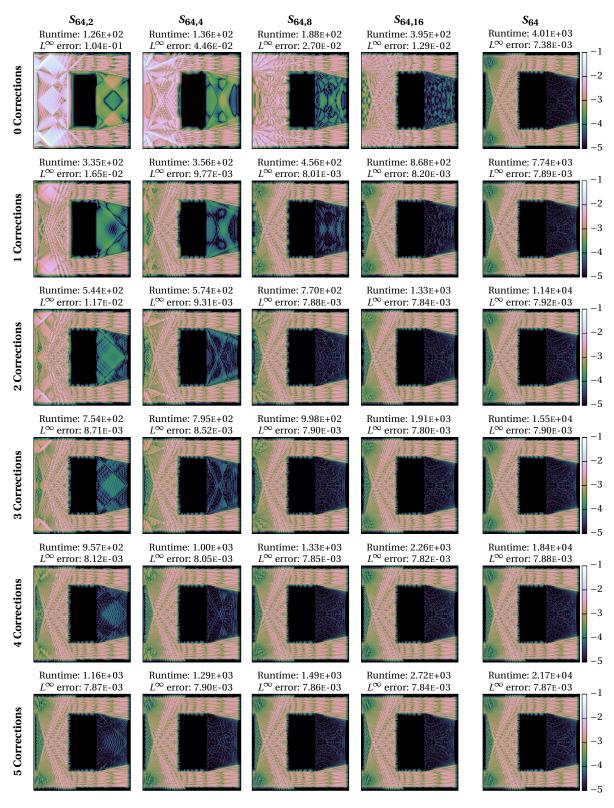


Figure 4.36: Deviation of scalar flux density from reference solution for hybrid-IIb IDC5 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.



Figure 4.37: Deviation of scalar flux density from reference solution for hybrid-IIc IDC5 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

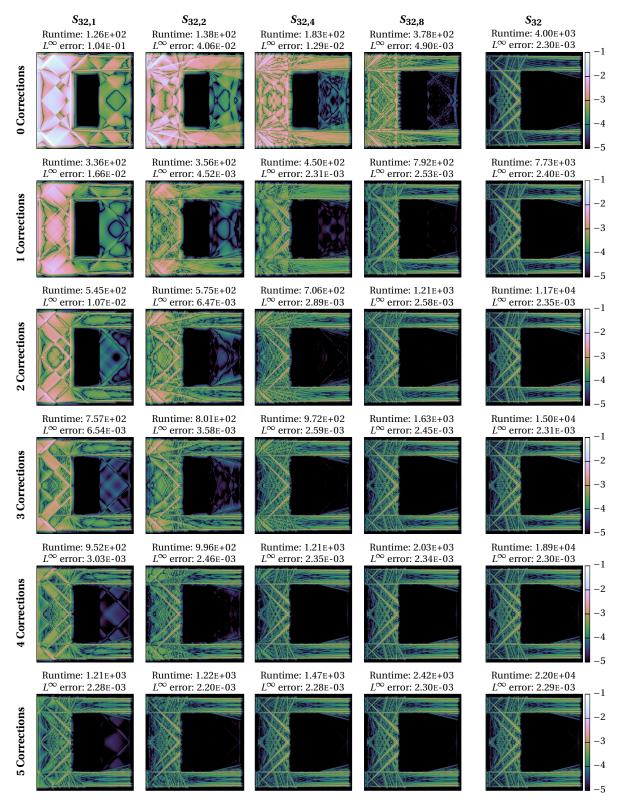


Figure 4.38: Deviation of scalar flux density from reference solution for hybrid-IIb IDC5 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

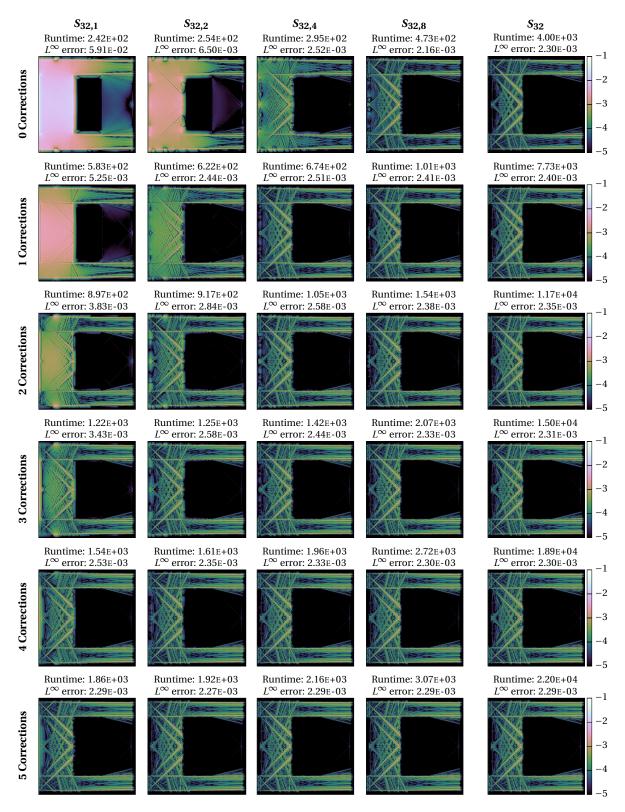


Figure 4.39: Deviation of scalar flux density from reference solution for hybrid-IIc IDC5 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

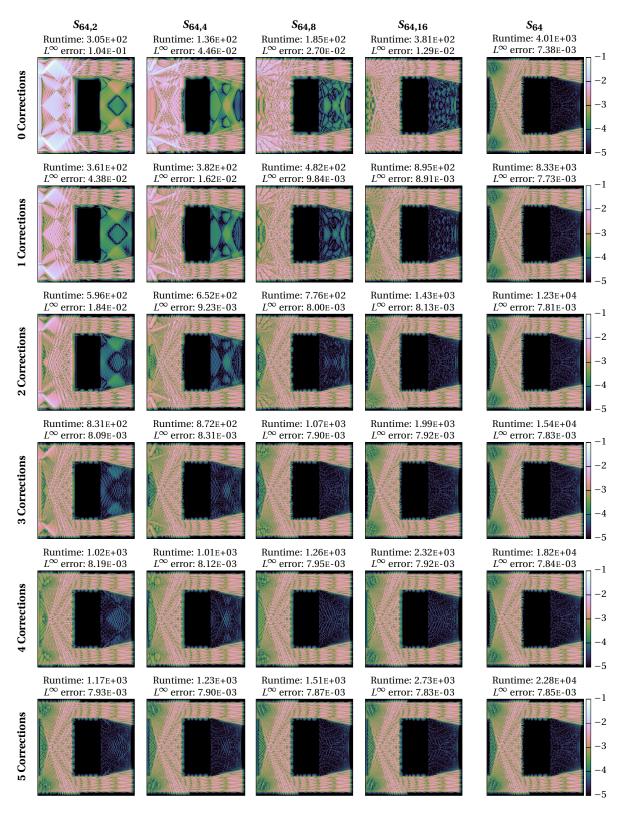


Figure 4.40: Deviation of scalar flux density from reference solution for hybrid-IIb LS-IDC5 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.



Figure 4.41: Deviation of scalar flux density from reference solution for hybrid-IIc LS-IDC5 methods using CL quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

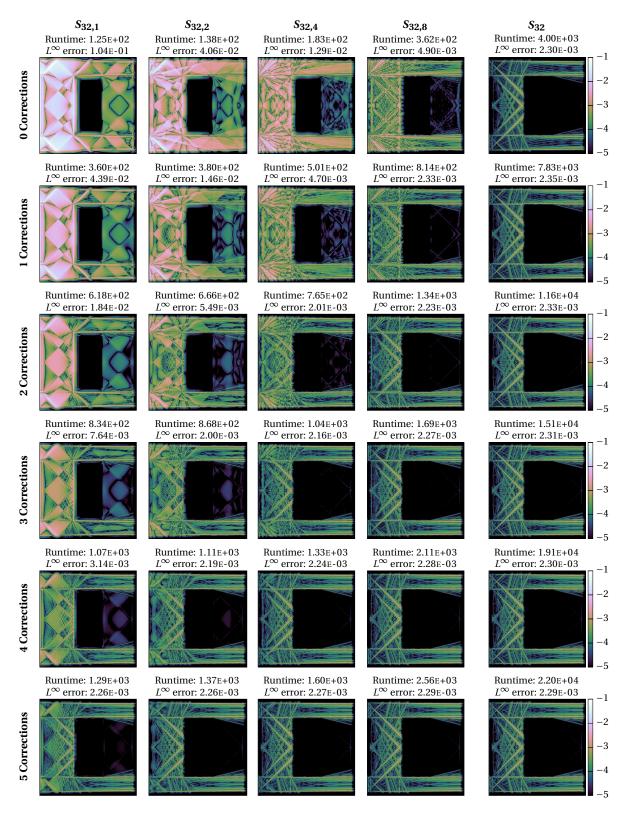


Figure 4.42: Deviation of scalar flux density from reference solution for hybrid-IIb LS-IDC5 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

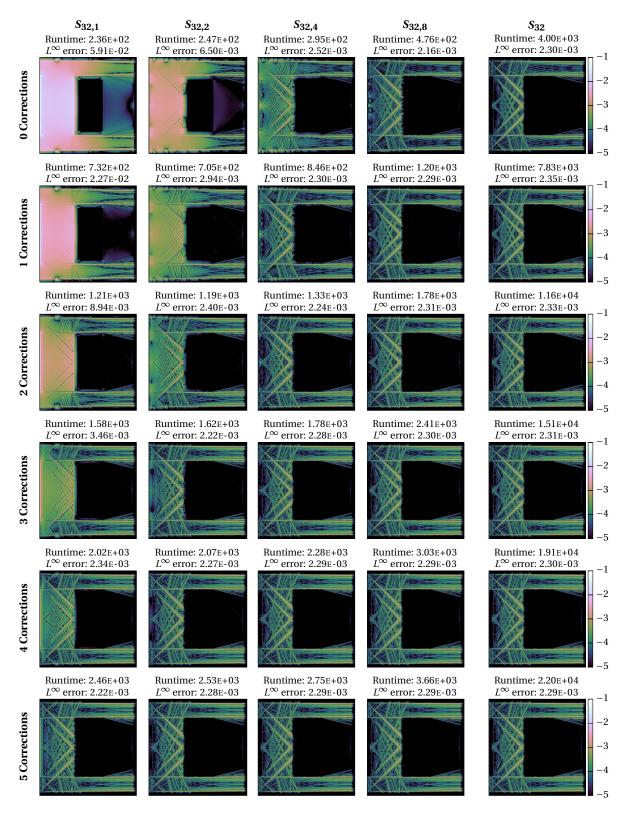


Figure 4.43: Deviation of scalar flux density from reference solution for hybrid-IIc LS-IDC5 methods using T_N quadratures and different numbers of correction iterations applied to the hohlraum test problem. The solve time (in seconds) and the L^{∞} error are included for each solution.

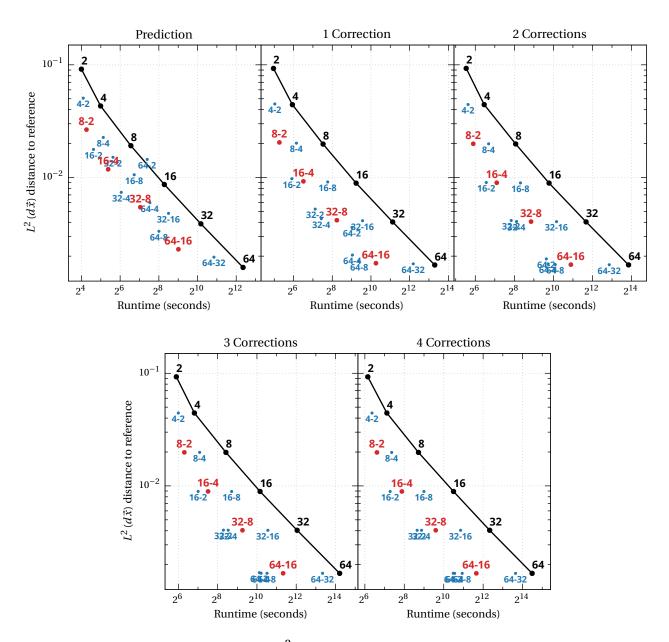


Figure 4.44: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

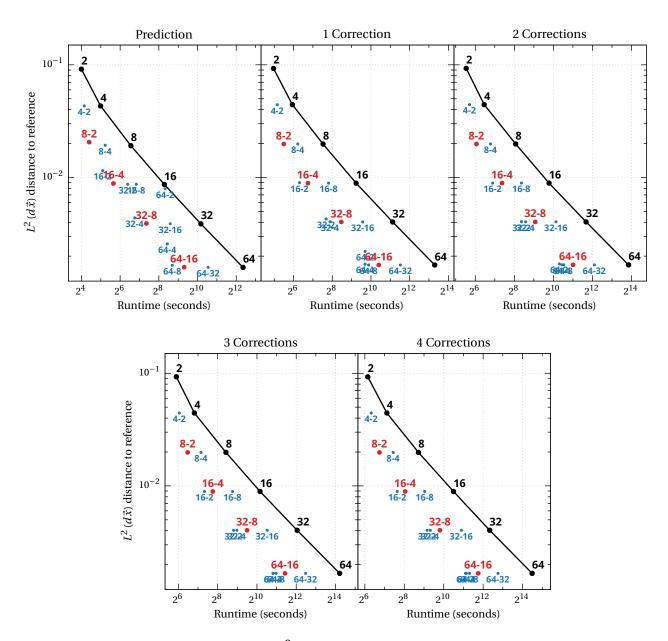


Figure 4.45: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

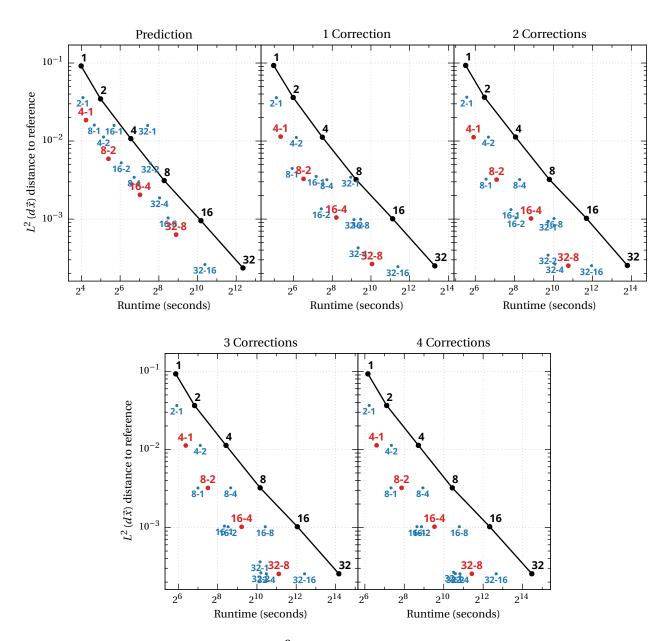


Figure 4.46: Solve times in seconds versus L^2 ($d\vec{x}$) distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

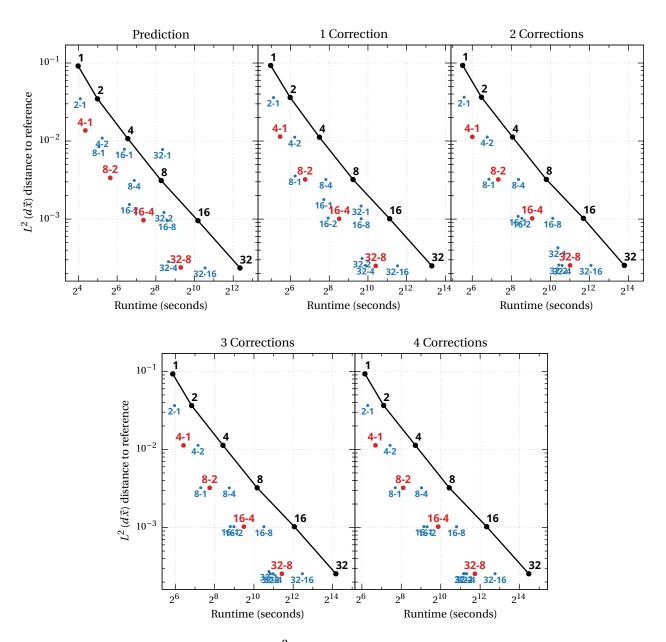


Figure 4.47: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

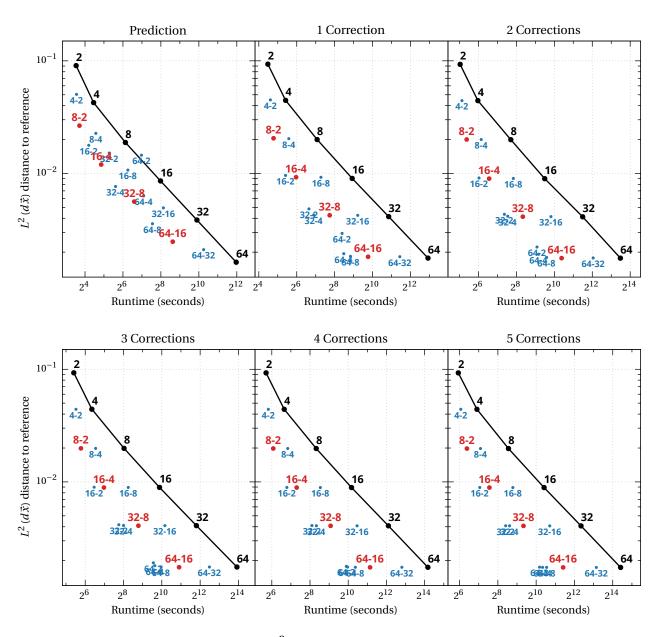


Figure 4.48: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

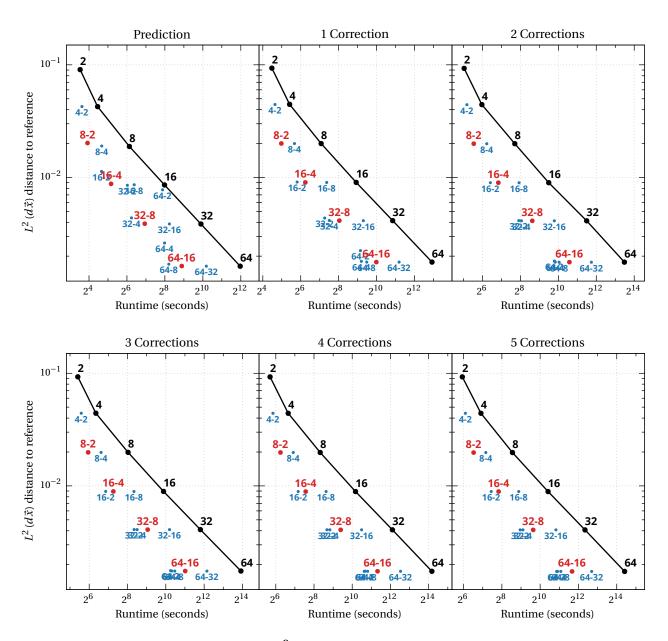


Figure 4.49: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

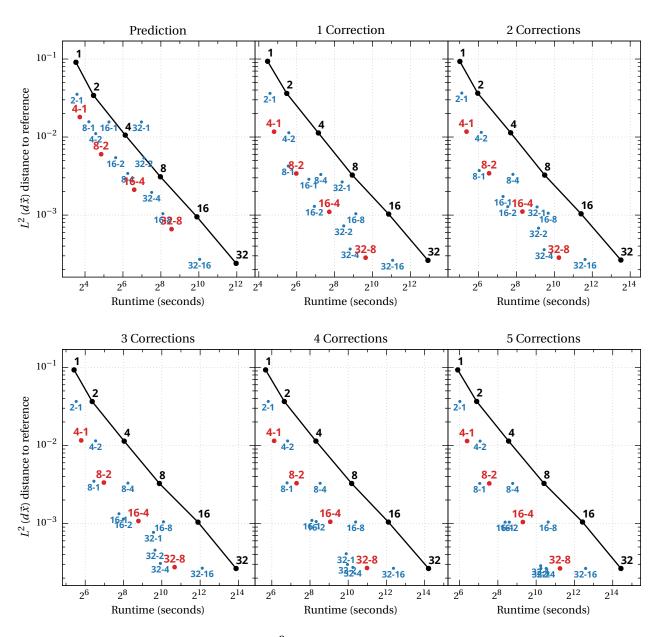


Figure 4.50: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

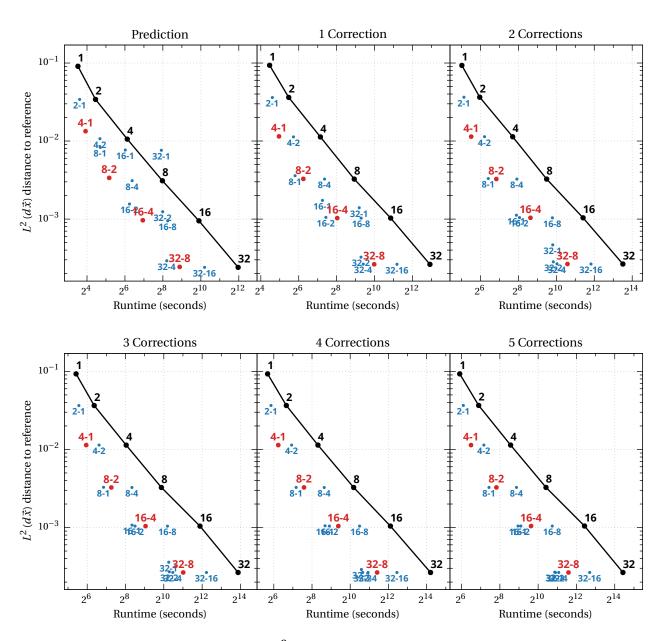


Figure 4.51: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

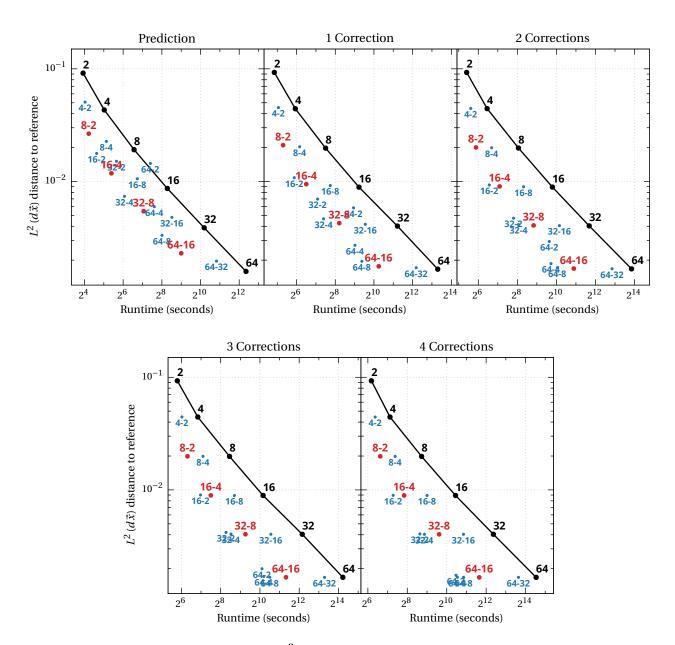


Figure 4.52: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the LS-IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

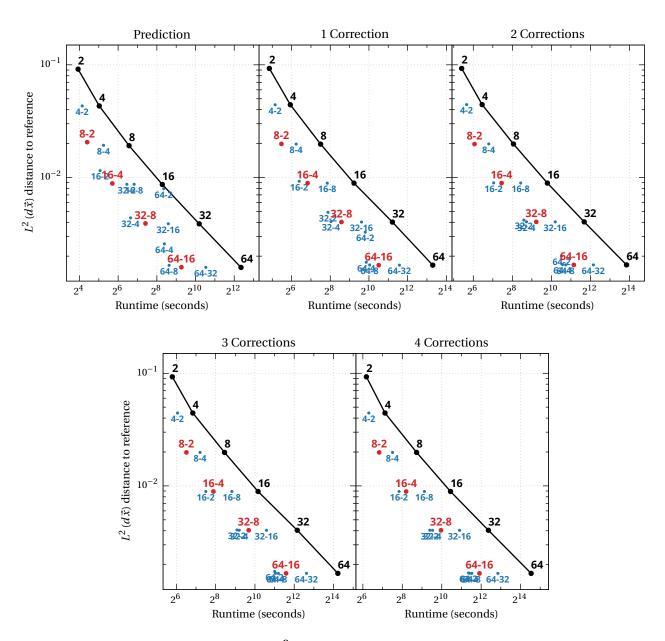


Figure 4.53: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the LS-IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

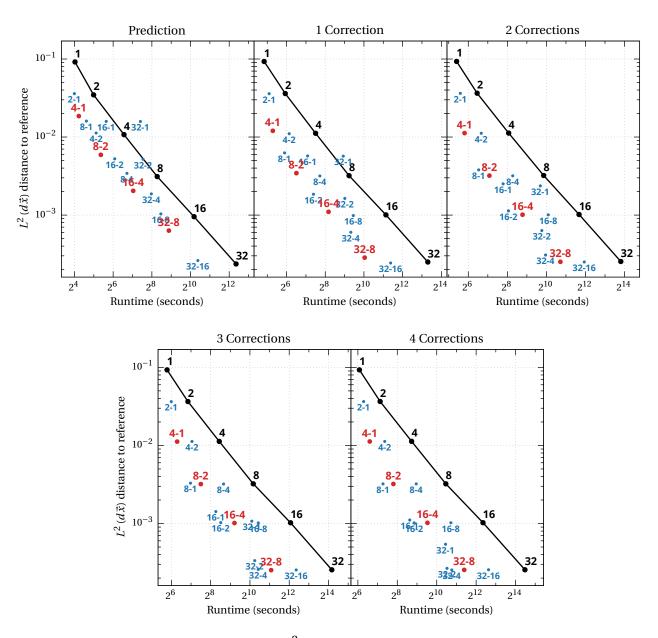


Figure 4.54: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the LS-IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

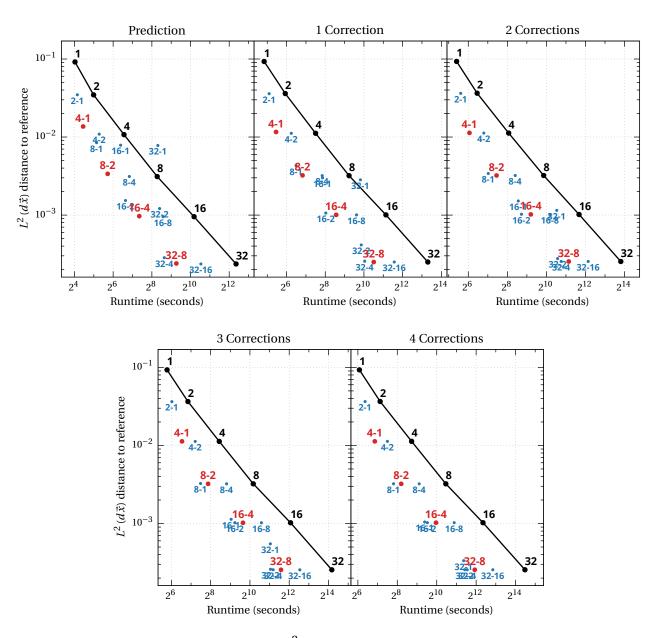


Figure 4.55: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the LS-IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

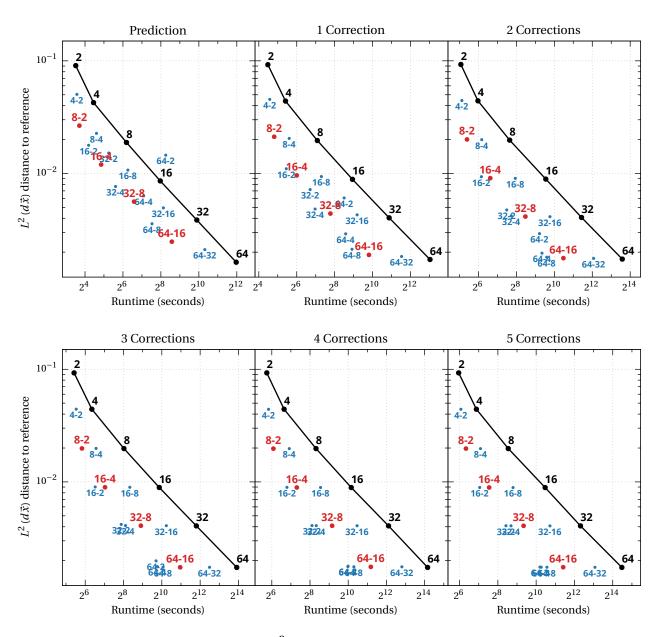


Figure 4.56: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the LS-IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

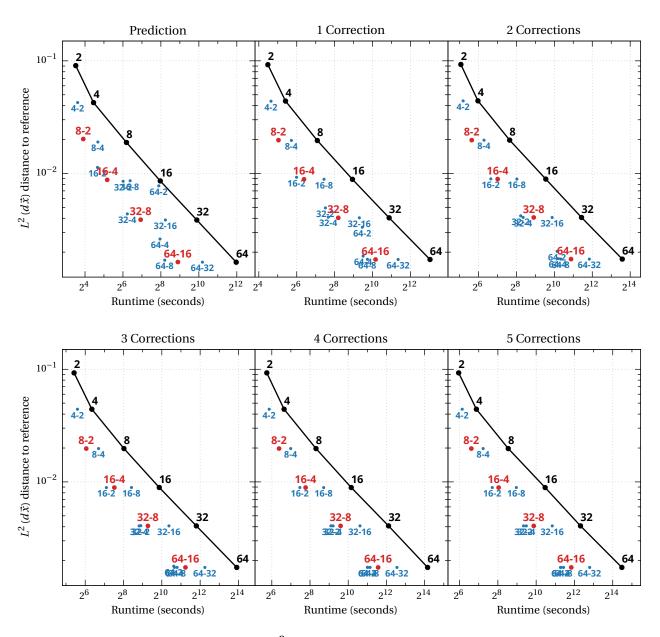


Figure 4.57: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the LS-IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

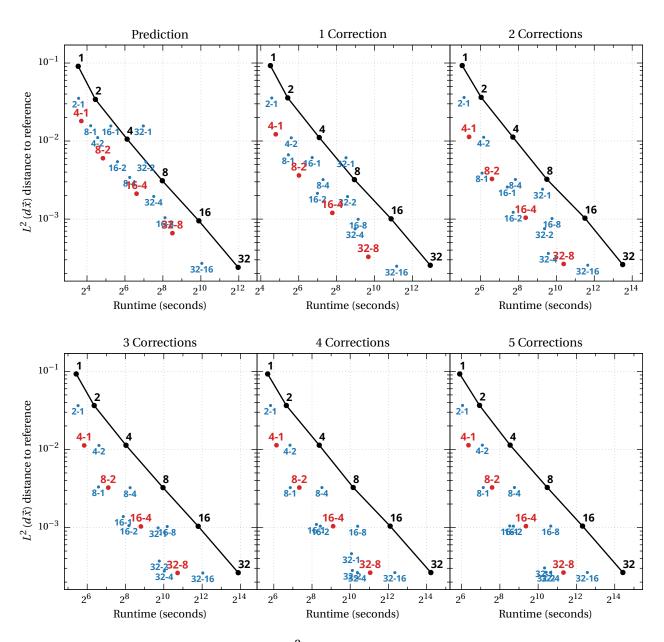


Figure 4.58: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIb (red and blue) discrete ordinates methods based on the LS-IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

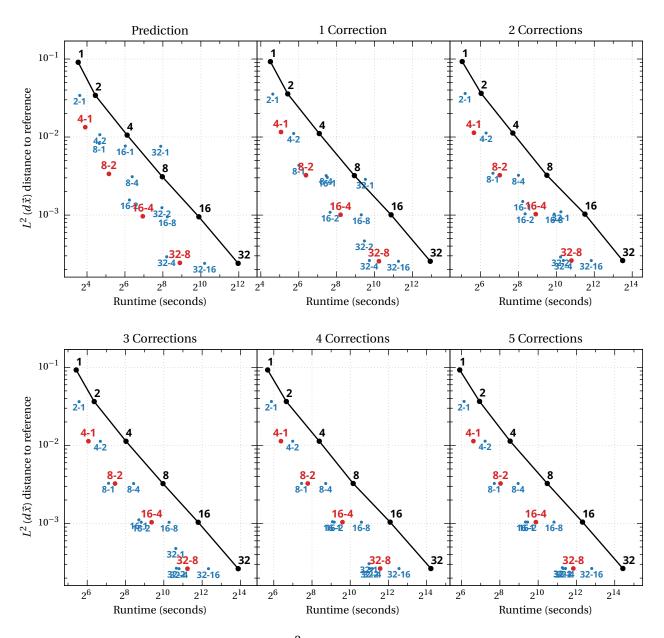


Figure 4.59: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid-IIc (red and blue) discrete ordinates methods based on the LS-IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations. Hybrid methods using angular quadratures that differ in order by a factor of 4 are shown in red and other hybrid methods are shown in blue.

vation holds for all two-grid methods regardless of the reconstruction procedure that is used. However, two-grid methods using angular quadratures that differ in order by a factor of eight appear to be a superior choice in many cases: two-grid methods using angular quadratures that differ in order by a factor of eight generally require between 20 and 40 percent less solve time that the equivalent method using angular quadratures that differ in order by a factor of four. In a few cases – for the hybrid-IIb IDC3 and LS-IDC3 methods using T_N quadratures – the two-grid methods using angular quadratures that differ in order by a factor of eight yield a small increase in the solution error, but if one additional correction iteration is performed (i.e., three correction iterations in total) then this deviation in accuracy is not longer present. This is not an issue for the hybrid-IIc methods using Nyström reconstruction procedures, and it seems that the hybrid-IIc methods may in fact support further coarsening the resolution of the quadrature that is applied to the collided component. However, doing so reduces overall solve time by only a very small amount (less than 10 percent), because the uncollided portion begins to dominate the solve time. Since other problems may be more sensitive to further reducing the resolution of the collided quadrature, we conclude that reductions beyond a factor of eight difference in quadrature order are unlikely to be worthwhile.

Each figure in the remaining set is designed to compare four types of methods based on IDC or LS-IDC integrators: (i) standard discrete ordinates methods, (ii) hybrid-Ib discrete ordinates methods using angular quadratures that differ in order by a factor of four, (iii) hybrid-IIb discrete ordinates methods using angular quadratures that differ in order by a factor of eight, and (iii) hybrid-IIc discrete ordinates methods using angular quadratures that differ in order by a factor of eight. All of the methods shown in each figure are constructed using the same time-integration scheme and angular quadrature types. Results are again presented for IDC and LS-IDC integrators using various numbers of correction iterations. Results are shown for methods based on the IDC3 scheme in figures 4.60 and 4.61, for methods based on the IDC5 scheme in figures 4.62 and 4.63, for methods based on the LS-IDC3 scheme in figures 4.64 and 4.65, and for methods based on the LS-IDC5 scheme in figures 4.66 and 4.67.

In each case, the two-grid methods (hybrid-IIb and hybrid-IIc) offer significant improvements in computational efficiency relative to both the standard discrete ordinates methods and the hybrid-Ib methods. It is also clear that the hybrid-IIc methods are more expensive than an equivalent hybrid-IIb method. For example, the $S_{32,4}$ hybrid-IIb methods using T_N quadratures require between one-thirteenth and one-sixteenth the solve time of the corresponding standard discrete ordinates method

using a T_N quadrature of order 32, while the $S_{32,4}$ hybrid-IIc methods require between one-eighth and one-tenth the solve time of the S_{32} method. However, the time savings offered by the hybrid-IIb methods do not come without sacrifice. First, whereas the hybrid-IIc methods offer excellent computational efficiency regardless of the number of correction iterations that are performed, the efficacy of the hybrid-IIb methods depends much more strongly on the iteration count. Because of this, the hybrid-IIb methods may be much less effective when applied to other types of problems. In addition, the Nyström reconstruction procedures used by the hybrid-IIc methods are much more flexible than the quadrature-specific reconstruction procedures used by the hybrid-IIb methods. In particular, the implementation of the hybrid-IIc methods is much simpler since it does not depend in any way on the choice of angular quadrature.

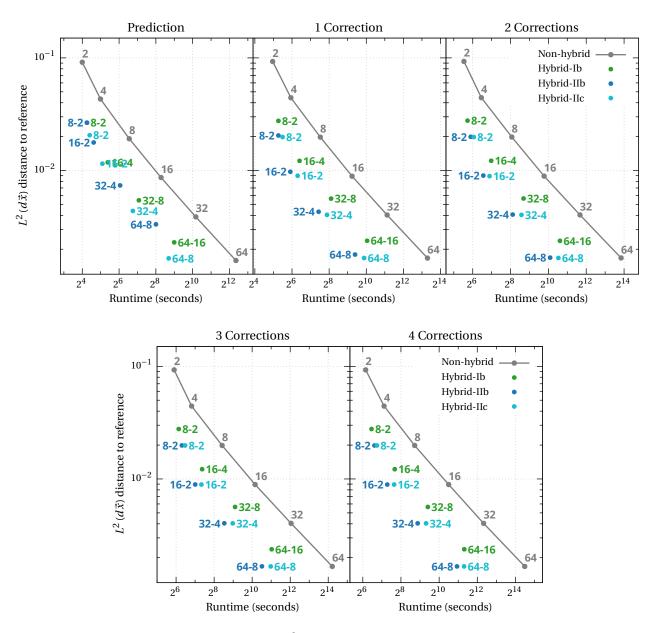


Figure 4.60: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (blue, teal, and green) discrete ordinates methods based on the IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations.

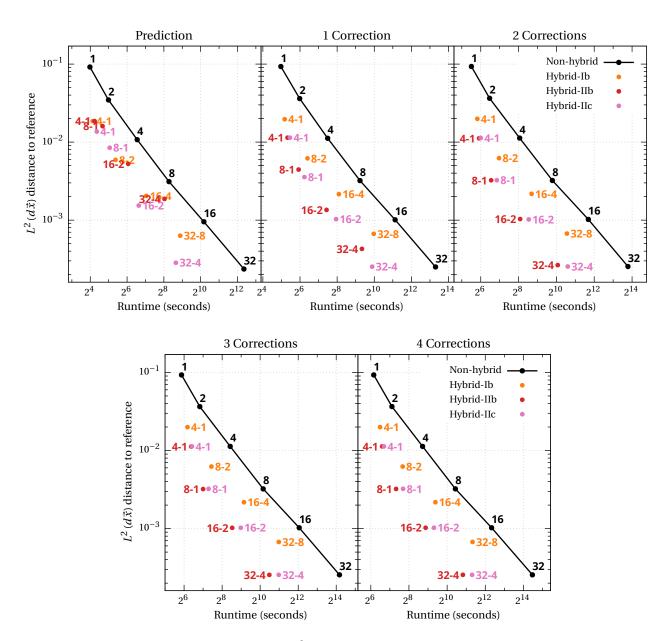


Figure 4.61: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (red, pink, and orange) discrete ordinates methods based on the IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations.

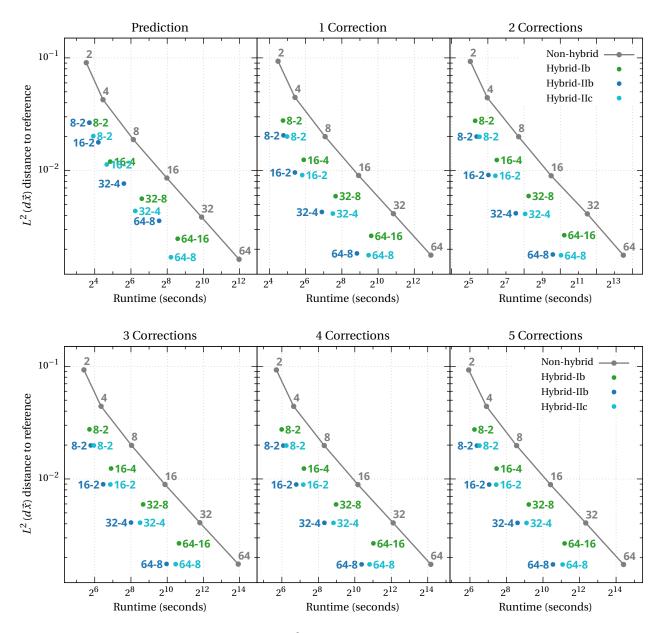


Figure 4.62: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (blue, teal, and green) discrete ordinates methods based on the IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations.

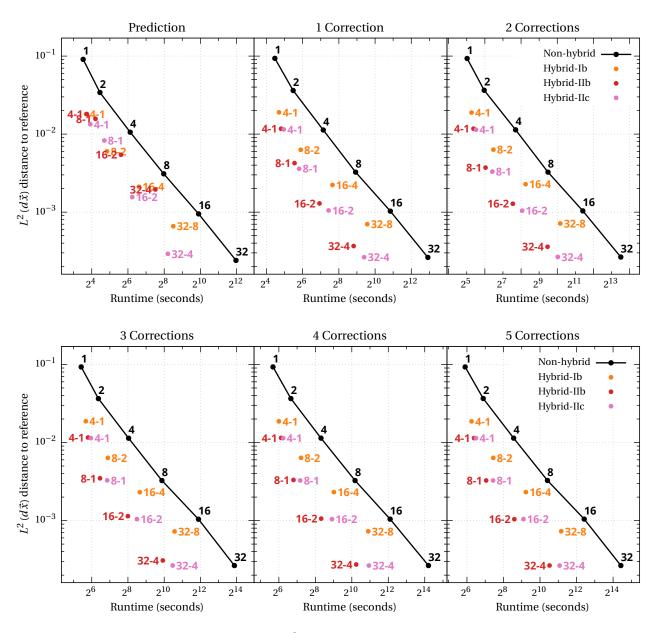


Figure 4.63: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (red, pink, and orange) discrete ordinates methods based on the IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations.

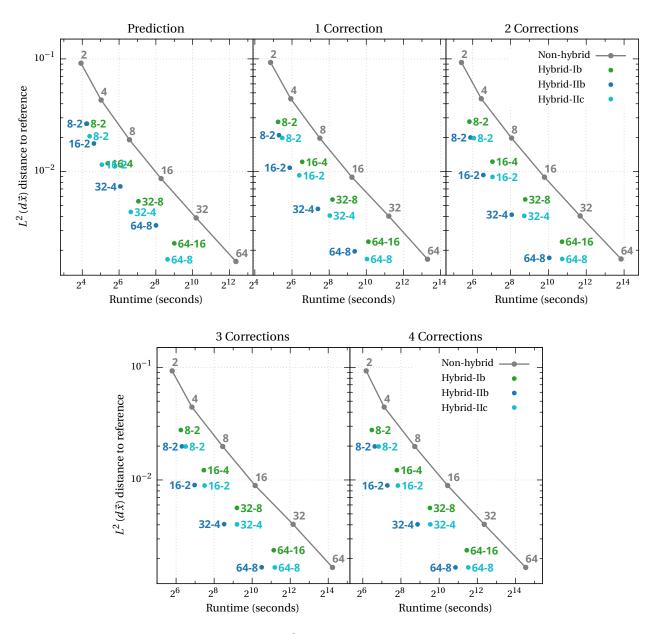


Figure 4.64: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (blue, teal, and green) discrete ordinates methods based on the LS-IDC3 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations.

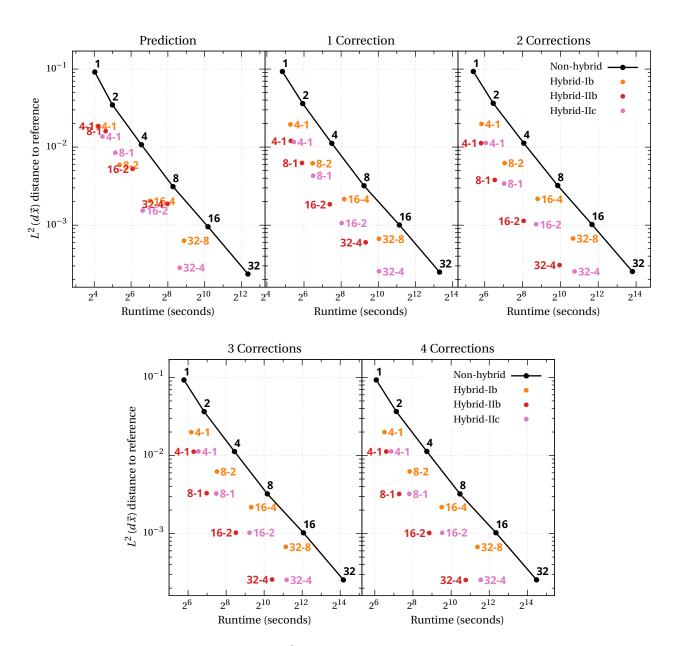


Figure 4.65: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (red, pink, and orange) discrete ordinates methods based on the LS-IDC3 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations.

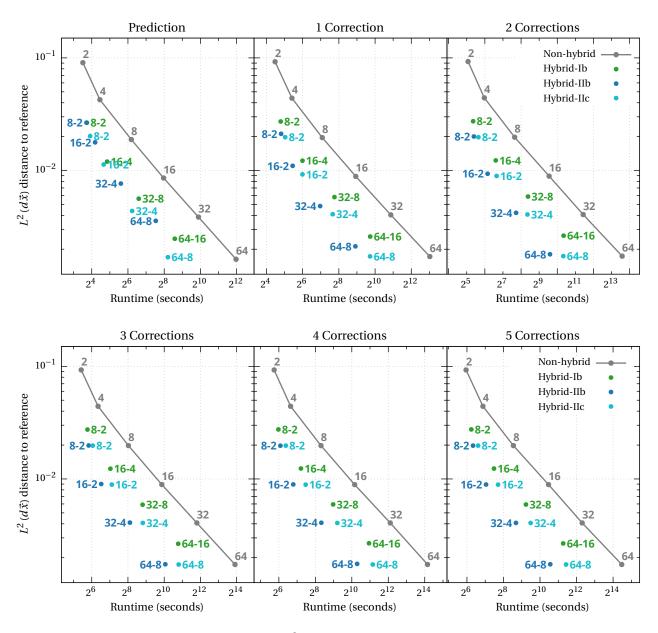


Figure 4.66: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (blue, teal, and green) discrete ordinates methods based on the LS-IDC5 integrator using CL quadratures applied to the hohlraum test problem with various numbers of correction iterations.

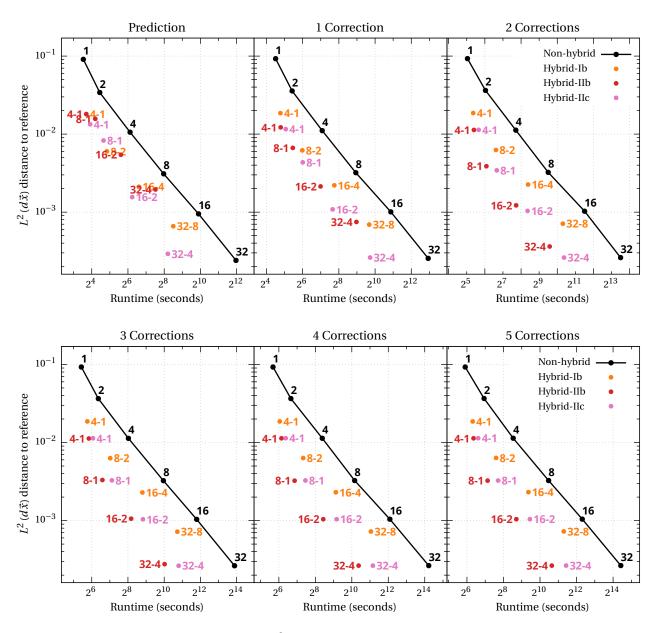


Figure 4.67: Solve times in seconds versus $L^2(d\vec{x})$ distance to reference solution for standard (black) and hybrid (red, pink, and orange) discrete ordinates methods based on the LS-IDC5 integrator using T_N quadratures applied to the hohlraum test problem with various numbers of correction iterations.

CHAPTER 5

CONCLUSIONS AND FUTURE WORK

5.1 Low-storage integral deferred correction methods for scientific computing

In chapter 2, a simple modification of traditional integral deferred correction (IDC) time-integration schemes was proposed that significantly reduced the storage requirements of the methods. These methods, which we refer to as low-storage IDC or LS-IDC methods, require storing only one copy of each stage vector throughout the iteration process. This reduction was achieved by splitting the spectral integration matrix into upper-triangular and lower-triangular parts, and applying the lower-triangular portion to the most recent set of stages in the residual calculation instead of the previous set. This results in lower storage requirements than other implicit methods of the same order of accuracy.

An analytic proof was given, showing that the LS-IDC methods converge with the same order of accuracy as traditional IDC methods for linear, constant-coefficient problems, and numerical investigations were used to examine the linear stability properties of some LS-IDC methods. It was found that the low-storage modification negatively impacts the stability properties of many LS-IDC methods, particularly for methods using right-biased quadratures that include the right endpoint of the timestep interval but not the left. This suggests that the stability of LS-IDC methods is insufficient for very stiff problems. Many LS-IDC methods using Gauss-Lobatto quadratures of order up to 10 were found to have better contraction properties in the strongly-stiff limit than IDC methods using the same quadratures, but were found to have inferior stability behavior for values near the imaginary axis. These methods may be an attractive option for certain non-stiff systems in which stability for values with large negative real part is more important than for values near the imaginary axis.

Two test problems were used to demonstrate the efficacy of LS-IDC methods. For a simple nonlinear ODE problem, LS-IDC methods were found to yield moderately larger errors than equivalent IDC methods using the same number of correction iterations, but similar accuracy when one additional correction iteration was applied. When applied to a linear transport equation solved by the method of lines, LS-IDC methods were found to yield similar or better accuracy than other methods without applying additional correction iterations. These results suggest that even though the low-storage modification negatively

impacts the stability properties of many LS-IDC methods, these methods are still an attractive option for solving high-dimensional partial differential equations that lack a significant degree of numerical stiffness.

In general, the reduction in memory footprint offered by LS-IDC methods depends strongly on the implicit solver that is used. The memory savings that were obtained when solving the transport equation (1.2) relied on the fact that the solvers used for discrete ordinates methods – either the source iteration method (3.13) or the GMRES formulation (3.14) and (3.15) – require storing only one angular flux vector throughout the iterative process. Similar implementations that rely on simple fixed-point methods or solving systems of reduced dimension should derive comparable benefits from the low-storage modification. However, there is little apparent advantage to LS-IDC methods in cases where, for example, Krylov methods are applied to the full system or Jacobian matrices must be constructed.

At this time it is unclear if the stability properties of LS-IDC methods can be improved. The modification of the traditional IDC approach used to produce the LS-IDC methods is similar to the work of [184], in which linear algebraic arguments were used to tune the contractive properties of spectral deferred correction methods to particular scenarios. An analogous technique could be applied to LS-IDC methods to improve their stability properties by selecting a low-order DIRK scheme to replace the implicit Euler steps. Since the low-order scheme only needs to provide a stable, consistent discretization, the additional degrees of freedom that are present when multiple Euler steps are combined could be augmented to improve the stability properties of the LS-IDC iteration. However, it would likely be necessary to hand-tune these modifications to each desired quadrature, and it is not obvious what sort of stability improvements might be attainable or how such modifications would impact other properties of the methods.

5.2 High-order hybrid methods for radiation transport

In chapter 3, extensions of a collision-based hybrid method for time-dependent radiation transport simulations were discussed. The hybrid methods were constructed by splitting the radiation flux into collided and uncollided components to which low- and high-resolution discrete ordinates approximations were applied, respectively. A reconstruction procedure was defined to map the collided component to the high-resolution angular quadrature between timesteps. The use of arbitrarily high-order numerical approximations was emphasized, with particular attention paid to high-order time-integration

schemes. A range of fully-implicit time integrators were considered, including DIRK, STDG, IDC, and LS-IDC methods of up to fifth-order accuracy, and multiple approaches for constructing hybrid methods were compared.

A sequence of convergence tests under a diffusive scaling was conducted in one-dimensional slab geometry. While the hybrid methods exhibited convergence stagnation and order reduction in certain scenarios, the overall accuracy of the hybrid approximations was found to be comparable to standard discrete ordinates methods in many cases. The stagnation phenomenon that was observed is readily explained: Since particles are moved to the collided distribution when they undergo collisions, the overall accuracy of the hybrid methods degrades to the accuracy of the low-resolution approximation used for the collided component when problem conditions are such that the majority of particles experience collisions over each timestep. However, a rigorous explanation for the observed order reduction phenomenon is less straightforward. We believe that this phenomenon is the result of a truncation error in time that is induced by the relabeling or reconstruction procedure that is applied. A rigorous study of the mechanism by which this phenomenon develops is currently underway.

With two goals in mind, the hybrid methods were applied to a test problem in two-dimensional xy-geometry consisting of a mock-up of a standard hohlraum configuration. The first goal was to investigate how the resolutions of the angular quadratures applied to the components of the hybrid methods could be chosen to adequately balance solve time and solution accuracy. Hybrid methods using angular quadratures with orders differing by a factor of four were found to consistently attain good efficiency relative to standard discrete ordinates methods and other hybrid methods. The second goal was to compare the computational efficiency of the hybrid methods with standard discrete ordinates methods. It was discovered that replacing a standard discrete ordinates method using an angular quadrature of order N with a hybrid discrete ordinates method using angular quadratures of order N and N in the uncollided and collided components, respectively, reduced solve time by a factor of 2 or more and in many cases also increased solution accuracy by a factor of up to 2, depending on the type of angular quadrature used. The suggested replacement was noted to have the disadvantage of increasing memory consumption by a factor of approximately 4.25, which may make these hybrid methods undesirable in applications where available system memory is limited.

Two types of angular quadratures were considered for multidimensional geometries. The first type of quadratures, the Chebyshev-Legendre (CL) quadratures, are constructed using a tensor-product ap-

proach and are designed to satisfy large numbers of moment conditions 1 . The second type of quadratures, the T_{N} quadratures, are constructed using a geometric approach that more evenly distributes quadrature nodes over the surface of the sphere. It was observed that the advantage in computational efficiency derived from the hybrid approach was dependent on the angular quadrature that was used, and that standard discrete ordinates methods using angular quadratures tailored towards the characteristics of the problem under consideration could yield similar errors in the same amount of solve time as hybrid methods using a less efficient choice of angular quadratures. Just as the choice of angular quadrature can significantly impact the accuracy of standard discrete ordinates methods, judiciously selecting the angular quadrature for the collided and uncollided components of the hybrid methods is a crucial component of their overall efficiency.

The angular quadratures used in this work were selected based on the fact that (i) they permit an arbitrary level of refinement with strictly positive quadrature weights, and (ii) their construction is relatively straightforward. Having demonstrated the feasibility of the hybrid approach, we believe that further investigation into the choice of angular quadratures is warranted. Future explorations should include more advanced quadratures, e.g. [6, 7, 107, 131, 132, 133], and should examine whether any benefit may be derived from selecting angular quadratures for the components of the hybrid methods that are of different types. For example, using angular quadratures for the collided component that satisfy a certain number of moment conditions may improve accuracy in highly collisional regimes, while it would be more effective to emphasize other properties in the uncollided quadrature. Adaptive refinement strategies could also be investigated, though we would hypothesize that an adaptive strategy would be far more beneficial for the high-resolution quadrature than the low-resolution quadrature.

In chapter 4, two mechanisms for increasing the effectiveness of the hybrid methods were presented. First, a novel reconstruction procedure for mapping between arbitrary discrete ordinates quadratures within the context of these hybrid methods was described. This approach, called Nyström reconstruction, was shown to have a number of significant advantages over previous reconstruction methods. Re-

¹More specifically, these quadratures were chosen because an arbitrary number of moment conditions may be satisfied by simply increasing the order of the quadrature. It should be noted that these quadratures are rather inefficient in the sense that the number of moment conditions satisfied by these quadratures is relatively small compared to the number of nodes in each quadrature. More efficient quadratures exist (e.g., [6, 131, 132, 133]) but they are generally much more difficult to construct, particularly for very high orders.

peating the convergence tests conducted previously, it was found that hybrid methods using Nyström reconstruction procedures suffered from the same convergence stagnation phenomenon observed previously, but concerns over the order reduction phenomenon were significantly reduced. Hybrid methods using a Gauss-Legendre quadrature of order 2 in the collided component had previously been observed to converge with second-order accuracy when reconstruction procedures based on polynomial interpolation were used, but fifth-order accuracy was observed when Nyström reconstruction procedures were used. This suggests a superconvergence rate of order 2N + 1 for hybrid methods based on Nyström reconstructions when a Gauss-Legendre quadrature of order N is used in the collided component.

When applied to the hohlraum test problem, the Nyström reconstruction procedures were found to significantly increase the accuracy of the hybrid methods. Whereas pronounced ray effects were observed in many of the solutions obtained from hybrid methods using other reconstruction procedures, no visibly apparent ray effects due to the lower-resolution angular quadrature were observed in the scalar flux distributions obtained from hybrid methods using Nyström reconstructions. For the previously recommended quadrature resolutions, the Nyström reconstruction procedures were found to increase the solve time of the hybrid methods by up to 55 percent. However, the increased solve time was generally outweighed by increased solution accuracy. The accuracy of hybrid methods using Nyström reconstruction procedures and angular quadratures of order N and N/4 for the uncollided and collided components, respectively, was consistently found to be comparable to that of a standard discrete ordinates method using an angular quadrature of order N. In each case, the hybrid methods required solve times that were between 4 and 8 times smaller than the standard discrete ordinates methods, with larger reductions observed when higher-resolution angular quadratures were used. By eliminating the need to increase the resolution of the angular quadrature that is applied to the uncollided component of the hybrid methods, the Nyström reconstruction procedures remedy the effective increase in memory consumption required by other hybrid methods. The revised recommendations given for hybrid methods using Nyström reconstruction – that angular quadratures of order N and N/4 be used for the uncollided and collided components, respectively – increases overall memory consumption by approximately 6 percent compared to a standard discrete ordinates method using an angular quadrature of order N.

Aside from increasing solution accuracy and reducing effective memory requirements to acceptable levels, the Nyström reconstruction procedures have a number of additional benefits. First, the Nyström reconstruction procedures are extremely flexible and eliminate the need to construct quadrature-

specific mapping procedures. The same reconstruction algorithm can be used in combination with any pair of angular quadratures. This approach could also be used to combine angular discretizations other than discrete ordinates methods. For example, a diffusion approximation, moment method, or finite element method could all be applied to the collided component instead of a low-resolution discrete ordinates approximation. In general, the Nyström reconstruction procedures only require that an approximation of the scalar flux distribution be constructed from the collided and uncollided components in a manner that is consistent with the spatial discretization applied to the uncollided flux. This flexibility should greatly simplify future studies that examine the effectiveness of combining different angular approximations within the context of these hybrid methods.

In certain applications (e.g., thermal radiative transfer) maintaining positivity in the scalar flux distribution is important not only because negative values are non-physical, but severe instabilities may occur when the scalar flux contains negative values. While we did not specifically examine mechanisms for maintaining positivity in this work, Nyström reconstruction procedures should be more well-behaved in this regard than other reconstruction approaches. Some methods, including the reconstruction procedures discussed in section 3.3.2 that were based on polynomial interpolation or hyperinterpolation steps, may introduce negative values into the reconstructed solution even if the initial values of the collided component are strictly non-negative. So long as the spatial and angular discretizations applied to the components of the hybrid methods yield positive representations of these components, the Nyström reconstruction procedures should preserve this property. That is, the Nyström reconstruction procedures should not introduce negative values into the solution where they do not already exists.

In the final portion of this work, it was demonstrated that hybrid methods based on IDC integrators could be modified to iteratively improve the accuracy of the hybrid splitting in addition to the temporal accuracy of the time discretization. These methods, which were referred to as two-grid IDC methods, demonstrate the effectiveness of a combined defect correction approach based on IDC integrators. When the convergence tests conducted previously were repeated for the two-grid methods, it was found that the saturation and order reduction that was observed with previous hybrid methods were both significantly reduced by the two-grid approach, and almost completely eliminated when Nyström reconstruction procedures were used.

When applied to the hohlraum test problem, it was found that the defect correction approach allowed for the resolution of the angular quadrature applied to the collided component to be further re-

duced without sacrificing solution accuracy. Specifically, two-grid methods using angular quadratures that differ in order by a factor of eight were found to consistently attain good efficiency relative to standard discrete ordinates methods and other hybrid methods. Additionally, the accuracy of two-grid methods using angular quadratures of order N and N/8 for the uncollided and collided components, respectively, was consistently found to be comparable to that of a standard discrete ordinates method using an angular quadrature of order N. In each case, the two-grid methods required solve times that were between 6 and 16 times smaller than the standard discrete ordinates methods, depending on the reconstruction procedure and angular quadratures that were applied.

These results demonstrate that the two-grid approach is able to iteratively decrease the error of the hybrid methods as intended, but several questions remain. In particular, it is not fully clear when such methods would have an efficiency advantage over hybrid methods based on other time-integration schemes. LS-IDC methods were demonstrated to have smaller memory requirements than other time integrators of the same order, but the total number of stages that must be computed is higher than for other methods. For example, the SDIRK3 and LS-IDC3 schemes require 3 and 6 total stages, respectively, and the ARK8 and LS-IDC5 schemes require 7 and 15 implicit stages, respectively. It should be noted that the computational cost associated with the stages of implicit IDC methods is not necessarily the same as that of implicit DIRK methods. Since the previous set of provisional solutions provides a rather good initial guess for the iterative solves required during subsequent correction iterations, the number of solver iterations required for each stage can decrease significantly as additional correction iterations are performed. Hence a one-to-one comparison of stage counts may not reliably quantify the differences in computational cost between time-integration schemes. A thorough comparison of the computational efficiency of hybrid methods constructed using different time integration schemes should be established, particularly for calculations in multidimensional geometries.

When Nyström reconstruction procedures were used, the performance of different hybrid methods was rather consistent relative to standard discrete ordinates methods when applied to the hohlraum test problem. However, the properties of the various hybrid methods provides some direction for the types of problems in which the performance of the two-grid methods may be superior to that of hybrid methods based on other time integrators. In general, the first-collision source technique works well for problems that transition from optically thin to optically thick, but is ineffective for problems that transition from optically thin. The motivation for introducing a relabeling or reconstruction procedure

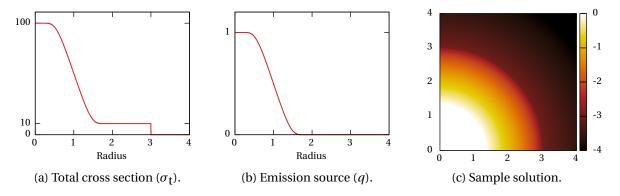


Figure 5.1: Layout and sample solution for the inhomogeneous sphere problem. Color scale of solution is log_{10} .

is to allow first-collision source methods to more accurately capture thick to thin transitions. The Nyström reconstruction procedures introduced in this work significantly increase the effectiveness of hybrid methods based on the first-collision source technique in the thick to thin transition, particularly when this transition is sharp, as is the case for the hohlraum test problem. A common problem used to test radiation transport codes is the homogeneous sphere problem [1, 43, 163, 165, 172], which emphasizes a thick to thin transition at a sharp material interface. Since hybrid methods using Nyström reconstruction seem to perform very well when this transition is sharp, ongoing work to evaluate the hybrid methods discussed in this work has focused on a modification of the homogeneous sphere problem in which the thick to thin transition occurs more slowly. Some initial results for this modified problem, which we refer to as the inhomogeneous sphere problem, are now briefly presented.

The inhomogeneous sphere problem consists of a spherically symmetric body of varying density radiating into a vacuum. We consider a two-dimensional version of the problem on a spatial domain of $[-4,4] \times [-4,4]$ with vacuum (zero outflow) boundary conditions on all sides. The parameters of the problem and one quadrant of a reference solution are shown in figure 5.1. It is assumed that particle collisions are composed of pure scattering; i.e., $\sigma_S = \sigma_t$. The convergence behavior of hybrid-Ic SDIRK3 methods and hybrid-IIc LS-IDC3 methods are compared in figure 5.2. All solutions are computed using T_N quadratures. As with previous convergence tests, we observe that (i) the saturation phenomenon is an issue for the hybrid-Ic DIRK methods, and (ii) the two-grid defect correction methods reduce the impact of this saturation phenomenon. Because the geometry of this problem is much more complex than that of the one-dimensional convergence tests conducted in previous portions of this work, the accuracy of the hybrid-IIc LS-IDC3 methods suffers when very coarse angular approximations are applied to the

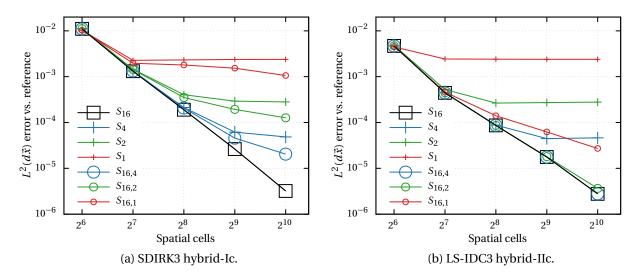


Figure 5.2: L^2 -convergence studies of hybrid-Ic SDIRK3 methods and hybrid-IIc LS-IDC3 methods applied to the inhomogeneous sphere problem with a final time of 8 and a CFL of 128. All solutions are computed using T_N quadratures.

collided component. However, we note that the $S_{16,1}$ approximations use quadratures that differ more significantly in their resolution than might be recommended: whereas the order 16 quadrature contains 256 nodes per octant, the order 1 quadrature contains only 1 quadrature node per octant.

There are many directions in which the two-grid hybrid methods may be extended in future work. First, the convergence of the two-grid methods using IDC integrators must be formally established. Additionally, there are many practical issues that must be resolved before these methods can be considered for widespread use. One issue of particular importance is preserving positivity in the solution. At present, we are not aware of any work that examines enforcing positivity constraints when IDC time integrators are used. IDC methods attain high-order accuracy through the use of residual terms in the correction equations that act as fictitious sources. Since these terms are purely mathematical in nature, they may be either positive or negative, depending on the properties of the solution and the high-order quadrature that is used to evaluate the residual. It is not clear how positivity may be enforced while preserving high-order accuracy when negative source terms are present.

It was noted in section 4.1 that the two-grid iteration on which the two-grid hybrid methods are based does not converge in a satisfactory manner. Indeed, the numerical results in section 4.4 showed that the convergence of the two-grid iterations is generally not monotonic. However, other approaches to two-grid methods exist that do converge monotonically, and their convergence is also generally much faster [17, Section 6.2.2]. Future work should investigate whether additional benefits may be derived

from incorporating these methods into the hybrid approach.

BIBLIOGRAPHY

BIBLIOGRAPHY

- [1] E. ABDIKAMALOV, A. BURROWS, C. D. OTT, F. LÖFFLER, E. O'CONNOR, J. C. DOLENCE, AND E. SCHNETTER, *A new Monte Carlo method for time-dependent neutrino radiation transport*, The Astrophysical Journal, 755 (2012), p. 111, doi:10.1088/0004-637x/755/2/111.
- [2] I. K. ABU-SHUMAYS, *Angular quadratures for improved transport computations*, Transport Theory and Statistical Physics, 30 (2001), pp. 169–204, doi:10.1081/TT-100105367.
- [3] M. L. Adams, *Discontinuous finite element transport solutions in thick diffusive problems*, Nuclear Science and Engineering, 137 (2001), pp. 298–333, doi:10.13182/nse00-41.
- [4] M. L. Adams and E. W. Larsen, *Fast iterative methods for discrete-ordinates particle transport calculations*, Progress in Nuclear Energy, 40 (2002), pp. 3–159, doi:10.1016/S0149-1970(01)00023-3.
- [5] M. P. Adams, M. L. Adams, W. D. Hawkins, T. Smith, L. Rauchwerger, N. M. Amato, T. S. Bailey, and R. D. Falgout, *Provably optimal parallel transport sweeps on regular grids*, in International Conference on Mathematics and Computational Methods Applied to Nuclear Science & Engineering (M&C 2013), Sun Valley, Idaho, USA, May 5–9, 2013, American Nuclear Society.
- [6] C. Ahrens and G. Beylkin, *Rotationally invariant quadratures for the sphere*, Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences, 465 (2009), pp. 3103–3125, doi:10.1098/rspa.2009.0104.
- [7] C. D. Ahrens, *Highly efficient, exact quadratures for three-dimensional discrete ordinates transport calculations*, Nuclear Science and Engineering, 170 (2012), pp. 98–101, doi:10.13182/nse10-69tn.
- [8] R. E. Alcouffe, *A first collision source method for coupling Monte Carlo and discrete ordinates for localized source problems*, in Monte-Carlo Methods and Applications in Neutronics, Photonics and Statistical Physics, R. Alcouffe, R. Dautray, A. Forster, G. Ledanois, and B. Mercier, eds., vol. 240 of Lecture Notes in Physics, Springer, Berlin, Heidelberg, 1985, pp. 352–366, doi:10.1007/BFb0049062.
- [9] R. E. ALCOUFFE, R. D. O'DELL, AND F. W. BRINKEY, JR., A first-collision source method that satisfies discrete S_N transport balance, Nuclear Science and Engineering, 105 (1990), pp. 198–203, doi:10.13182/NSE90-A23749.
- [10] R. ALEXANDER, *Diagonally implicit Runge-Kutta methods for stiff O.D.E.'s*, SIAM Journal on Numerical Analysis, 14 (1977), pp. 1006–1021, doi:10.1137/0714068.
- [11] P. ALLU AND S. MAZUMDER, *Hybrid ballistic-diffusive solution to the frequency-dependent phonon Boltzmann transport equation*, International Journal of Heat and Mass Transfer, 100 (2016), pp. 165–177, doi:10.1016/j.ijheatmasstransfer.2016.04.049.
- [12] D. N. Arnold, *An interior penalty finite element method with discontinuous elements*, SIAM Journal on Numerical Analysis, 19 (1982), pp. 742–760, doi:10.1137/0719052.
- [13] D. N. ARNOLD, F. BREZZI, B. COCKBURN, AND D. MARINI, *Unified analysis of discontinuous Galerkin methods for elliptic problems*, SIAM Journal on Numerical Analysis, 39 (2002), pp. 1749–1779, doi:10.1137/S0036142901384162.

- [14] T. ASAOKA, Y. NAKAHARA, AND K. SAITO, *Multiple collision method for neutron transport problems*, Journal of Nuclear Energy. Parts A/B. Reactor Science and Technology, 18 (1964), pp. 665–701, doi:10.1016/0368-3230(64)90122-3.
- [15] U. M. ASHER AND L. R. PETZOLD, Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations, SIAM, 1998.
- [16] K. Atkinson and W. Han, *Spherical Harmonics and Approximations on the Unit Sphere: An Introduction*, vol. 2044 of Lecture Notes in Mathematics, Springer, Berlin, Heidelberg, 2012, doi:10.1007/978-3-642-25983-8.
- [17] K. E. ATKINSON, *The Numerical Solution of Integral Equations of the Second Kind*, vol. 4 of Cambridge Monographs on Applied and Computational Mathematics, Cambridge University Press, Cambridge, United Kingdom, 1997.
- [18] W. AUZINGER, H. HOFSTÄTTER, W. KREUZER, AND E. WEINMÜLLER, *Modified defect correction algorithms for ODEs part I: General theory*, Numerical Algorithms, 36 (2004), pp. 135–155, doi:10.1023/B:NUMA.0000033129.73715.7f.
- [19] W. AUZINGER, H. HOFSTÄTTER, W. KREUZER, AND E. WEINMÜLLER, *Modified defect correction algorithms for ODEs part II: Stiff initial value problems*, Numerical Algorithms, 40 (2005), pp. 285–303, doi:10.1007/s11075-005-5327-4.
- [20] M. BADRI, P. JOLIVET, B. ROUSSEAU, AND Y. FAVENNEC, *High performance computation of radiative transfer equation using the finite element method*, Journal of Computational Physics, 360 (2018), pp. 74–92, doi:10.1016/j.jcp.2018.01.027.
- [21] T. S. BAILEY AND R. D. FALGOUT, *Analysis of massively parallel discrete-ordinates transport sweep algorithms with collisions*, in International Conference on Mathematics, Computational Methods, and Reactor Physics (M&C 2009), Saratoga Springs, New York, USA, May 3–7, 2009, American Nuclear Society.
- [22] F. BASSI AND S. REBAY, A high-order accurate discontinuous finite element method for the numerical solution of the compressible Navier-Stokes equations, Journal of Computational Physics, 131 (1997), pp. 267–279, doi:10.1006/jcph.1996.5572.
- [23] H. Berninger, E. Frénod, M. Gander, M. Liebendörfer, and J. Michaud, *Derivation of the isotropic diffusion source approximation (IDSA) for supernova neutrino transport by asymptotic expansions*, SIAM Journal on Mathematical Analysis, 45 (2013), pp. 3229–3265, doi:10.1137/12089243x.
- [24] H. BERNINGER, E. FRÉNOD, M. J. GANDER, M. LIEBENDÖRFER, J. MICHAUD, AND N. VASSET, *A mathematical description of the IDSA for supernova neutrino transport, its discretization and a comparison with a finite volume scheme for Boltzmann's equation*, ESAIM: Proceedings, 38 (2012), pp. 163–182, doi:10.1051/proc/201238009.
- [25] K. BÖHMER, *Discrete Newton methods and iterated defect corrections*, Numerische Mathematik, 37 (1981), pp. 167–192, doi:10.1007/bf01398251.
- [26] K. BÖHMER, P. W. HEMKER, AND H. J. STETTER, *The defect correction approach*, in Defect Correction Methods: Theory and Applications, H. Böhmer and H. Stetter, eds., vol. 5 of Computing Supplementa, Springer, Vienna, Vienna, 1984, pp. 1–32, doi:10.1007/978-3-7091-7023-6_1.

- [27] C. BÖRGERS, E. W. LARSEN, AND M. L. ADAMS, *The asymptotic diffusion limit of a linear discontinuous discretization of a two-dimensional linear transport equation*, Journal of Computational Physics, 98 (1992), pp. 285–300, doi:10.1016/0021-9991(92)90143-m.
- [28] S. BOSCARINO, L. PARESCHI, AND G. RUSSO, *Implicit-explicit Runge-Kutta schemes for hyperbolic systems and kinetic equations in the diffusion limit*, SIAM Journal on Scientific Computing, 35 (2013), pp. A22–A51, doi:10.1137/110842855.
- [29] J. P. BOYD, Chebyshev and Fourier Spectral Methods, Dover Publications, Inc., second ed., 2001.
- [30] F. Brezzi, G. Manzini, D. Marini, P. Pietra, and A. Russo, *Discontinuous Galerkin approximations for elliptic problems*, Numerical Methods for Partial Differential Equations, 16 (2000), pp. 365–378, doi:10.1002/1098-2426(200007)16:4<365::AID-NUM2>3.0.CO;2-Y.
- [31] T. A. Brunner, *Forms of approximate radiation transport*, Tech. Report SAND2002-1778, Sandia National Laboratories, Albuquerque, New Mexico, 2002.
- [32] B. G. Carlson, *Tables of equal weight quadrature* EQ_n *over the unit sphere*, Tech. Report LA-4734, Los Alamos Scientific Laboratory, Los Alamos, New Mexico, 1971.
- [33] B. G. Carlson and C. E. Lee, *Mechanical quadrature and the transport equation*, Tech. Report LAMS-2573, Los Alamos Scientific Laboratory, Los Alamos, New Mexico, 1961.
- [34] E. CARSON AND N. J. HIGHAM, Accelerating the solution of linear systems by iterative refinement in three precisions, SIAM Journal on Scientific Computing, 40 (2018), pp. A817–A847, doi:10.1137/17m1140819.
- [35] J. CASH, A variable order deferred correction algorithm for the numerical solution of nonlinear two point boundary value problems, Computers & Mathematics with Applications, 9 (1983), pp. 257–265, doi:10.1016/0898-1221(83)90129-3.
- [36] P. CASTILLO, B. COCKBURN, I. PERUGIA, AND D. SCHÖTZAU, An a priori error analysis of the local discontinuous Galerkin method for elliptic problems, SIAM Journal on Numerical Analysis, 38 (2000), pp. 1676–1706, doi:10.1137/S0036142900371003.
- [37] R. O. CASTRO AND J. P. TRELLES, *Spatial and angular finite element method for radiative transfer in participating media*, Journal of Quantitative Spectroscopy and Radiative Transfer, 157 (2015), pp. 81–105, doi:10.1016/j.jqsrt.2015.02.008.
- [38] C. CERCIGNANI AND D. H. SATTINGER, *Scaling Limits and Models in Physical Processes*, Birkhäuser Basel, 1998, doi:10.1007/978-3-0348-8810-3.
- [39] G. CHEN, *Ballistic-diffusive heat-conduction equations*, Physical Review Letters, 86 (2001), pp. 2297–2300, doi:10.1103/physrevlett.86.2297.
- [40] G. CHEN, *Ballistic-diffusive equations for transient heat conduction from nano to macroscales*, Journal of Heat Transfer, 124 (2002), p. 320, doi:10.1115/1.1447938.
- [41] G.-Q. CHEN, C. D. LEVERMORE, AND T.-P. LIU, *Hyperbolic conservation laws with stiff relaxation terms and entropy*, Communications on Pure and Applied Mathematics, 47 (1994), pp. 787–830, doi:10.1002/cpa.3160470602.

- [42] Y. CHENG AND C.-W. SHU, A discontinuous Galerkin finite element method for time dependent partial differential equations with higher order derivatives, Mathematics of Computation, 77 (2008), pp. 699–730, doi:10.1090/S0025-5718-07-02045-5.
- [43] P. CHIDYAGWAI, M. FRANK, F. SCHNEIDER, AND B. SEIBOLD, A comparative study of limiting strategies in discontinuous Galerkin schemes for the M_1 model of radiation transport, Journal of Computational and Applied Mathematics, 342 (2018), pp. 399–418, doi:10.1016/j.cam.2018.04.017.
- [44] A. CHRISTLIEB, W. Guo, M. Morton, and J.-M. Qiu, *A high order time splitting method based on integral deferred correction for semi-Lagrangian Vlasov simulations*, Journal of Computational Physics, 267 (2014), pp. 7–27, doi:10.1016/j.jcp.2014.02.012.
- [45] A. CHRISTLIEB, M. MORTON, B. ONG, AND J. QIU, *Semi-implicit integral deferred correction constructed with additive Runge-Kutta methods*, Communications in Mathematical Sciences, 9 (2011), pp. 879–902, doi:10.4310/CMS.2011.v9.n3.a10.
- [46] A. CHRISTLIEB AND B. ONG, *Implicit parallel time integrators*, Journal of Scientific Computing, 49 (2011), pp. 167–179, doi:10.1007/s10915-010-9452-4.
- [47] A. CHRISTLIEB, B. ONG, AND J.-M. QIU, *Integral deferred correction methods constructed with high order Runge-Kutta integrators*, Mathematics of Computation, 79 (2010), pp. 761–783, doi:10.1090/S0025-5718-09-02276-5.
- [48] A. J. CHRISTLIEB, R. D. HAYNES, AND B. W. ONG, *A parallel space-time algorithm*, SIAM Journal on Scientific Computing, 34 (2012), pp. C233–C248, doi:10.1137/110843484.
- [49] A. J. CHRISTLIEB, Y. LIU, AND Z. XU, *High order operator splitting methods based on an inte-gral deferred correction framework*, Journal of Computational Physics, 294 (2015), pp. 224–242, doi:10.1016/j.jcp.2015.03.032.
- [50] A. J. CHRISTLIEB, C. B. MACDONALD, AND B. W. ONG, *Parallel high-order integrators*, SIAM Journal on Scientific Computing, 32 (2010), pp. 818–835, doi:10.1137/09075740X.
- [51] A. J. CHRISTLIEB, B. W. ONG, AND J.-M. QIU, *Comments on high-order integrators embedded within integral deferred correction methods*, Communications in Applied Mathematics and Computational Science, 4 (2009), pp. 27–56, doi:10.2140/camcos.2009.4.27.
- [52] M. CLEVELAND AND A. WOLLABER, *Corrected implicit Monte Carlo*, Journal of Computational Physics, 359 (2018), pp. 20–44, doi:10.1016/j.jcp.2017.12.038.
- [53] B. COCKBURN, S. HOU, AND C.-W. SHU, *The Runge-Kutta local projection discontinuous Galerkin finite element method for conservation laws IV: The multidimensional case*, Mathematics of Computation, 54 (1990), pp. 545–581, doi:10.1090/S0025-5718-1990-1010597-0.
- [54] B. COCKBURN, G. KANSCHAT, I. PERUGIA, AND D. SCHÖTZAU, Superconvergence of the local discontinuous Galerkin method for elliptic problems on Cartesian grids, SIAM Journal on Numerical Analysis, 39 (2001), pp. 264–285, doi:10.1137/S0036142900371544.
- [55] B. COCKBURN, S.-Y. LIN, AND C.-W. SHU, *TVB Runge-Kutta local projection discontinuous Galerkin finite element method for conservation laws III: One-dimensional systems*, Journal of Computational Physics, 84 (1989), pp. 90–113, doi:10.1016/0021-9991(89)90183-6.

- [56] B. COCKBURN AND C. SHU, *The local discontinuous Galerkin method for time-dependent convection-diffusion systems*, SIAM Journal on Numerical Analysis, 35 (1998), pp. 2440–2463, doi:10.1137/S0036142997316712.
- [57] B. COCKBURN AND C.-W. SHU, *TVB Runge-Kutta local projection discontinuous Galerkin finite ele*ment method for conservation laws II: General framework, Mathematics of Computation, 52 (1989), pp. 411–435, doi:10.1090/S0025-5718-1989-0983311-4.
- [58] B. COCKBURN AND C.-W. SHU, *The Runge-Kutta local projection P*¹-discontinuous-Galerkin finite element method for scalar conservation laws, ESIAM: Mathematical Modeling and Numerical Analysis Modèlisation Mathématique et Analyse Numérique, 25 (1991), pp. 337–361, doi:10.1051/m2an/1991250303371.
- [59] B. COCKBURN AND C.-W. SHU, *The Runge-Kutta discontinuous Galerkin method for conservation laws V: Multidimensional systems*, Journal of Computational Physics, 141 (1998), pp. 199–224, doi:10.1006/jcph.1998.5892.
- [60] J. C. COMPTON AND C. J. CLOUSE, *Tiling models for spatial decomposition in AMTRAN*, in Proceedings of Joint Russian-American Five-Laboratory Conference on Computational Mathematics/Physics, Vienna, Austria, June 19–23, 2005.
- [61] M. M. CROCKATT, A. J. CHRISTLIEB, C. K. GARRETT, AND C. D. HAUCK, *An arbitrary-order, fully implicit, hybrid kinetic solver for linear radiative transport using integral deferred correction*, Journal of Computational Physics, 346 (2017), pp. 212–241, doi:10.1016/j.jcp.2017.06.017.
- [62] B. DAVISON, Neutron Transport Theory, Clarendon Press, Oxford, England, 1957.
- [63] T. DEAKIN, S. McIntosh-Smith, M. Martineau, and W. Gaudin, *An improved parallelism scheme for deterministic discrete ordinates transport*, The International Journal of High Performance Computing Applications, (2016), pp. 1–14, doi:10.1177/1094342016668978.
- [64] M. Delfour, W. Hager, and F. Trochu, *Discontinuous Galerkin methods for ordinary differential equations*, Mathematics of Computation, 36 (1981), pp. 455–455, doi:10.1090/s0025-5718-1981-0606506-0.
- [65] J. D. DENSMORE, *Interface methods for hybrid Monte Carlo-diffusion radiation-transport simula-tions*, Annals of Nuclear Energy, 33 (2006), pp. 343–353, doi:10.1016/j.anucene.2005.11.007.
- [66] J. D. DENSMORE, T. J. URBATSCH, T. M. EVANS, AND M. W. BUKSAS, *A hybrid transport-diffusion method for Monte Carlo radiative-transfer simulations*, Journal of Computational Physics, 222 (2007), pp. 485–503, doi:10.1016/j.jcp.2006.07.031.
- [67] P. DEUFLHARD AND A. HOHMANN, *Numerical Analysis in Modern Scientific Computing*, Springer New York, 2003, doi:10.1007/978-0-387-21584-6.
- [68] J. DOUGLAS, JR. AND T. DUPONT, *Interior penalty procedures for elliptic and parabolic Galerkin methods*, in Computing Methods in Applied Sciences, vol. 58 of Lecture Nodes in Physics, Springer, Berlin, Heidelberg, 1976, pp. 207–216, doi:10.1007/BFb0120591.
- [69] J. J. DUDERSTADT AND W. R. MARTIN, Transport Theory, John Wiley & Sons, Inc., 1979.

- [70] A. DUTT, L. GREENGARD, AND V. ROKHLIN, Spectral deferred correction methods for ordinary differential equations, Bit Numerical Mathematics, 40 (2000), pp. 241–266, doi:10.1023/a:1022338906936.
- [71] T. Endo and A. Yamamoto, *Development of new solid angle quadrature sets to satisfy even- and odd-moment conditions*, Journal of Nuclear Science and Technology, 44 (2007), pp. 1249–1258, doi:10.1080/18811248.2007.9711368.
- [72] A. FILIMON, M. DUMBSER, AND C.-D. MUNZ, Discontinuous Galerkin schemes with defect corrections based on reconstruction, Computers & Fluids, 62 (2012), pp. 71–85, doi:10.1016/j.compfluid.2012.03.005.
- [73] A. FILIMON AND C.-D. MUNZ, *The application of iterated defect corrections based on WENO reconstruction*, in Management and Minimization of Uncertainties and Errors in Numerical Aerodynamics, B. Eisfeld, H. Barnewitz, W. Fritz, and F. Thiele, eds., vol. 122 of Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Springer, Berlin, Heidelberg, 2013, pp. 129–154, doi:10.1007/978-3-642-36185-2_6.
- [74] W. A. FIVELAND, *Discrete ordinate methods for radiative heat transfer in isotropically and anisotropically scattering media*, Journal of Heat Transfer, 109 (1987), p. 809, doi:10.1115/1.3248167.
- [75] J. FLECK AND J. CUMMINGS, *An implicit Monte Carlo scheme for calculating time and frequency dependent nonlinear radiation transport*, Journal of Computational Physics, 8 (1971), pp. 313–342, doi:10.1016/0021-9991(71)90015-5.
- [76] R. Frank and C. W. Ueberhuber, *Iterated defect correction for the efficient solution of stiff systems of ordinary differential equations*, BIT, 17 (1977), pp. 146–159, doi:10.1007/bf01932286.
- [77] R. Frank and C. W. Ueberhuber, *Collocation and iterated defect correction*, in Numerical Treatment of Differential Equations, R. Bulirsch, R. Grigorieff, and J. Schröder, eds., vol. 631 of Lecture Notes in Mathematics, Springer, Berlin, Heidelberg, 1978, pp. 19–34, doi:10.1007/bfb0067461.
- [78] B. GANAPOL, *Multiple collision solution of the linearized Vlasov equation*, Progress in Nuclear Energy, 8 (1981), pp. 151–161, doi:10.1016/0149-1970(81)90009-3.
- [79] B. D. Ganapol, *The multiple collision solution of Kramer's equation for Brownian motion*, Transport Theory and Statistical Physics, 11 (1982), pp. 217–231, doi:10.1080/00411458208245742.
- [80] B. D. GANAPOL, Exact time-dependent moments for the monoenergetic neutron transport equation with isotropic scattering, Nuclear Science and Engineering, 89 (1985), pp. 256–260, doi:10.13182/nse85-a17546.
- [81] B. D. GANAPOL, R. S. BAKER, J. A. DAHL, AND R. E. ALCOUFFE, *Homogeneous infinite media time-dependent analytical benchmarks*, Tech. Report LA-UR-01-1854, Los Alamos National Laboratory, Los Alamos, New Mexico, 2001.
- [82] B. D. GANAPOL AND V. C. BOFFI, The multiple collision method in solving the Boltzmann equation for time-dependent test particle transport, Meccanica, 15 (1980), pp. 75–81, doi:10.1007/bf02133788.
- [83] B. D. GANAPOL AND L. M. GROSSMAN, *The collided flux expansion method for time-dependent neutron transport*, Nuclear Science and Engineering, 52 (1973), pp. 454–460, doi:10.13182/nse73-a23312.

- [84] B. D. GANAPOL, P. W. MCKENTY, AND K. L. PEDDICORD, *The generation of time-dependent neutron transport solutions in infinite media*, Nuclear Science and Engineering, 64 (1977), pp. 317–331, doi:10.13182/nse77-a27373.
- [85] C. K. Garrett and C. D. Hauck, *A comparison of moment closures for linear kinetic transport equations: The line source benchmark*, Transport Theory and Statistical Physics, 42 (2013), pp. 203–235, doi:10.1080/00411450.2014.910226.
- [86] N. GENTILE, Implicit Monte Carlo diffusion— an acceleration method for Monte Carlo time-dependent radiative transfer simulations, Journal of Computational Physics, 172 (2001), pp. 543–571, doi:10.1006/jcph.2001.6836.
- [87] N. GENTILE, *Including the effects of temperature-dependent opacities in the implicit Monte Carlo algorithm*, Journal of Computational Physics, 230 (2011), pp. 5100–5114, doi:10.1016/j.jcp.2011.03.029.
- [88] L. M. GOMES AND P. N. STEVENS, *Use of Monte Carlo procedures to estimate first and second collision sources for discrete ordinates transport codes*, Nuclear Science and Engineering, 114 (1993), pp. 228–237, doi:10.13182/nse93-a24036.
- [89] A. N. GORBAN AND I. V. KARLIN, *Invariant Manifolds for Physical and Chemical Kinetics*, Springer, Berlin, Heidelberg, 2005, doi:10.1007/b98103.
- [90] G. J. HABETLER AND B. J. MATKOWSKY, *Uniform asymptotic expansions in transport theory with small mean free paths, and the diffusion approximation,* Journal of Mathematical Physics, 16 (1975), pp. 846–854, doi:10.1063/1.522618.
- [91] E. HAIRER, S. P. NØRSETT, AND G. WANNER, *Solving Ordinary Differential Equations I: Nonstiff Problems*, vol. 8 of Springer Series in Computational Mathematics, Springer, Berlin, Heidelberg, second ed., 1993, doi:10.1007/978-3-540-78862-1.
- [92] E. HAIRER AND G. WANNER, Solving Ordinary Differential Equations II: Stiff and Differential-Algebraic Problems, vol. 14 of Springer Series in Computational Mathematics, Springer, Berlin, Heidelberg, second ed., 1996, doi:10.1007/978-3-642-05221-7.
- [93] E. HAIRER AND G. WANNER, *Stiff differential equations solved by Radau methods*, Journal of Computational and Applied Mathematics, 111 (1999), pp. 93–111, doi:10.1016/S0377-0427(99)00134-X.
- [94] T. HANGSTROM AND R. ZHOU, On the spectral deferred correction of splitting methods for initial value problems, Communications in Applied Mathematics and Computational Science, 1 (2006), pp. 169–205, doi:10.2140/camcos.2006.1.169.
- [95] C. HAUCK AND R. MCCLARREN, *Positive P*_N closures, SIAM Journal on Scientific Computing, 32 (2010), pp. 2603–2626, doi:10.1137/090764918.
- [96] C. D. HAUCK, *High-order entropy-based closures for linear transport in slab geometry*, Communications in Mathematical Sciences, 9 (2011), pp. 187–205, doi:10.4310/cms.2011.v9.n1.a9.
- [97] C. D. HAUCK AND R. G. McClarren, *A collision-based hybrid method for time dependent, linear, kinetic transport equations,* Multiscale Modeling and Simulation, 11 (2013), pp. 1197–1227, doi:10.1137/110846610.

- [98] H. V. HENDERSON AND S. R. SEARLE, *On deriving the inverse of a sum of matrices*, SIAM Review, 23 (1981), pp. 53–60, doi:10.1137/1023004.
- [99] A. HENNINK AND D. LATHOUWERS, *A discontinuous Galerkin method for the mono-energetic Fokker-Planck equation based on a spherical interior penalty formulation*, Journal of Computational and Applied Mathematics, 330 (2018), pp. 253–267, doi:10.1016/j.cam.2017.08.019.
- [100] N. J. HIGHAM, Accuracy and Stability of Numerical Algorithms, SIAM, 2002, doi:10.1137/1.9780898718027.
- [101] A. HOISIE, O. LUBECK, AND H. WASSERMAN, *Performance and scalability analysis of teraflop-scale parallel architectures using multidimensional wavefront applications*, The International Journal of High Performance Computing Applications, 14 (2000), pp. 330–346, doi:10.1177/109434200001400405.
- [102] Y.-C. Hua and B.-Y. Cao, *Ballistic-diffusive heat conduction in multiply-constrained nanostructures*, International Journal of Thermal Sciences, 101 (2016), pp. 126–132, doi:10.1016/j.ijthermalsci.2015.10.037.
- [103] Y.-C. Hua and B.-Y. Cao, *The effective thermal conductivity of ballistic-diffusive heat conduction in nanostructures with internal heat source*, International Journal of Heat and Mass Transfer, 92 (2016), pp. 995–1003, doi:10.1016/j.ijheatmasstransfer.2015.09.068.
- [104] J. HUANG, J. JIA, AND M. MINION, *Accelerating the convergence of spectral deferred correction methods*, Journal of Computational Physics, 214 (2006), pp. 633–656, doi:10.1016/j.jcp.2005.10.004.
- [105] J. Huang, J. Jia, and M. Minion, *Arbitrary order Krylov deferred correction methods for differential algebraic equations*, Journal of Computational Physics, 221 (2007), pp. 739–760, doi:10.1016/j.jcp.2006.06.040.
- [106] M. JANKOWSKI AND H. WOŹNIAKOWSKI, *Iterative refinement implies numerical stability*, BIT, 17 (1977), pp. 303–311, doi:10.1007/bf01932150.
- [107] J. J. JARRELL, An adaptive angular discretization method for neutral-particle transport in three-dimensional geometries, PhD thesis, Texas A&M University, 2010.
- [108] C. JOHNSON AND J. PITKÄRANTA, *An analysis of the discontinuous Galerkin method for a scalar hyperbolic equation*, Mathematics of Computation, 46 (1986), pp. 1–26, doi:10.2307/2008211.
- [109] G. Kanschat, *Solution of radiative transfer problems with finite elements*, in Numerical Methods in Multidimensional Radiative Transfer, G. Kanschat, E. Meinköhn, R. Rannacher, and R. Wehrse, eds., Springer, Berlin, Heidelberg, 2009, pp. 49–98, doi:10.1007/978-3-540-85369-5_5.
- [110] C. A. KENNEDY AND M. H. CARPENTER, *Additive Runge-Kutta schemes for convection-diffusion-reaction equations*, Applied Numerical Mathematics, 44 (2003), pp. 139–181, doi:10.1016/s0168-9274(02)00138-1.
- [111] C. A. KENNEDY, M. H. CARPENTER, AND R. LEWIS, *Low-storage, explicit Runge-Kutta schemes for the compressible Navier-Stokes equations*, Applied Numerical Mathematics, 35 (2000), pp. 177–219, doi:10.1016/s0168-9274(99)00141-5.

- [112] D. I. KETCHESON, Highly efficient strong stability-preserving Runge-Kutta methods with low-storage implementations, SIAM Journal on Scientific Computing, 30 (2008), pp. 2113–2136, doi:10.1137/07070485x.
- [113] D. I. KETCHESON, *Runge-Kutta methods with minimum storage implementations*, Journal of Computational Physics, 229 (2010), pp. 1763–1773, doi:10.1016/j.jcp.2009.11.006.
- [114] S. Kholin, *Certain exact solutions of the nonstationary kinetic equation without taking retardation into account*, USSR Computational Mathematics and Mathematical Physics, 4 (1964), pp. 213–221, doi:10.1016/0041-5553(64)90094-1.
- [115] K. R. Koch, R. S. Baker, and R. E. Alcouffe, *Solution of the first-order form of three-dimensional discrete ordinates equations on a massively parallel machine*, Transactions of the Americal Nuclear Society, 65 (1992), pp. 198–199.
- [116] R. KOCH, W. KREBS, S. WITTIG, AND R. VISKANTA, *Discrete ordinates quadrature schemes for multidimensional radiative transfer*, Journal of Quantitative Spectroscopy & Radiative Transfer, 53 (1995), pp. 353–372, doi:10.1016/0022-4073(95)90012-8.
- [117] J. KÓPHÁZI AND D. LATHOUWERS, *A space-angle DGFEM approach for the Boltzmann radiation transport equation with local angular refinement*, Journal of Computational Physics, 297 (2015), pp. 637–668, doi:10.1016/j.jcp.2015.05.031.
- [118] A. KVÆRNØ, Singly diagonally implicit Runge-Kutta methods with an explicit first stage, BIT Numerical Mathematics, 44 (2004), pp. 489–502, doi:10.1023/b:bitn.0000046811.70614.38.
- [119] M. P. LAIU, C. D. HAUCK, R. G. McClarren, D. P. O'Leary, and A. L. Tits, *Positive filtered* P_N moment closures for linear kinetic equations, SIAM Journal on Numerical Analysis, 54 (2016), pp. 3214–3238, doi:10.1137/15m1052871.
- [120] E. W. LARSEN AND J. B. KELLER, *Asymptotic solution of neutron transport problems for small mean free paths*, Journal of Mathematical Physics, 15 (1974), pp. 75–81, doi:10.1063/1.1666510.
- [121] E. W. LARSEN, A. KUMAR, AND J. E. MOREL, *Properties of the implicitly time-differenced equations of thermal radiation transport*, Journal of Computational Physics, 238 (2013), pp. 82–96, doi:10.1016/j.jcp.2012.11.034.
- [122] E. W. LARSEN AND W. F. MILLER JR., Convergence rates of spatial difference equations for the discrete-ordinates neutron transport equations in slab geometry, Nuclear Science and Engineering, 73 (1980), pp. 76–83, doi:10.13182/NSE80-3.
- [123] E. W. LARSEN AND J. E. MOREL, *Advances in discrete-ordinates methodology*, in Nuclear Computational Science, Springer, Dordrecht, Dec. 2009, pp. 1–84, doi:10.1007/978-90-481-3411-3_1.
- [124] E. W. LARSEN AND P. NELSON, Finite-difference approximations and superconvergence for the discrete-ordinate equations in slab geometry, SIAM Journal on Numerical Analysis, 19 (1982), pp. 334–348, doi:10.1137/0719020.
- [125] P. LASAINT AND P. A. RAVIART, *On a finite element method for solving the neutron transport equation,* in Mathematical Aspects of Finite Elements in Partial Differential Equations, C. de Boor, ed., Academic Press, Inc., 1974, pp. 89–123, doi:10.1016/B978-0-12-208350-1.50008-X.

- [126] K. D. LATHROP, *Ray effects in discrete ordinates equations*, Nuclear Science and Engineering, 32 (1968), pp. 357–369, doi:10.13182/NSE68-4.
- [127] K. D. LATHROP, *Remedies for ray effects*, Nuclear Science and Engineering, 45 (1971), pp. 255–268, doi:10.13182/NSE45-03-255.
- [128] K. D. LATHROP AND B. G. CARLSON, *Discrete ordinates angular quadrature of the neutron transport equation*, Tech. Report LA-3186, Los Alamos Scientific Laboratory, Los Alamos, New Mexico, 1964.
- [129] C. Y. LAU, Adaptive discrete-ordinates quadratures based on discontinuous finite elements over spherical quadrilaterals, ph.D., Texas A&M University, 2016.
- [130] A. T. LAYTON AND M. L. MINION, *Implications of the choice of quadrature nodes for Picard integral deferred corrections methods for ordinary differential equations*, BIT Numerical Mathematics, 45 (2005), pp. 341–373, doi:10.1007/s10543-005-0016-1.
- [131] V. LEBEDEV, *Values of the nodes and weights of ninth to seventeenth order Gauss-Markov quadrature formulae invariant under the octahedron group with inversion*, USSR Computational Mathematics and Mathematical Physics, 15 (1975), pp. 44–51, doi:10.1016/0041-5553(75)90133-0.
- [132] V. LEBEDEV, *Quadratures on a sphere*, USSR Computational Mathematics and Mathematical Physics, 16 (1976), pp. 10–24, doi:10.1016/0041-5553(76)90100-2.
- [133] V. I. LEBEDEV, *Spherical quadrature formulas exact to orders 25–29*, Siberian Mathematical Journal, 18 (1977), pp. 99–107, doi:10.1007/bf00966954.
- [134] C. D. LEVERMORE AND G. C. POMRANING, *A flux-limited diffusion theory*, The Astrophysical Journal, 248 (1981), p. 321, doi:10.1086/159157.
- [135] E. E. LEWIS AND W. F. MILLER JR, Computational Methods of Neutron Transport, Wiley, 1993.
- [136] M. LIEBENDÖRFER, S. C. WHITEHOUSE, AND T. FISCHER, *The isotropic diffusion source approximation for supernova neutrino transport*, The Astrophysical Journal, 698 (2009), pp. 1174–1190, doi:10.1088/0004-637x/698/2/1174.
- [137] R. G. McClarren and C. D. Hauck, *Robust and accurate filtered spherical harmonics expansions for radiative transfer*, Journal of Computational Physics, 229 (2010), pp. 5597–5614, doi:10.1016/j.jcp.2010.03.043.
- [138] R. G. McClarren, J. P. Holloway, and T. A. Brunner, On solutions to the P_n equations for thermal radiative transfer, Journal of Computational Physics, 227 (2008), pp. 2864–2885, doi:10.1016/j.jcp.2007.11.027.
- [139] R. G. McClarren and T. J. Urbatsch, *A modified implicit Monte Carlo method for time-dependent radiative transfer with adaptive material coupling*, Journal of Computational Physics, 228 (2009), pp. 5669–5686, doi:10.1016/j.jcp.2009.04.028.
- [140] C. D. MEYER, Matrix Analysis and Applied Linear Algebra, The SIAM, 2000.
- [141] A. MEZZACAPPA AND O. MESSER, *Neutrino transport in core collapse supernovae*, Journal of Computational and Applied Mathematics, 109 (1999), pp. 281–319, doi:10.1016/s0377-0427(99)00162-4.

- [142] J. MICHAUD, *The IDSA and the homogeneous sphere: Issues and possible improvements*, Discrete and Continuous Dynamical Systems - Series S, 9 (2016), pp. 1351–1375, doi:10.3934/dcdss.2016054.
- [143] D. MIHALAS AND B. W. MIHALAS, *Foundations of Radiation Hydrodynamics*, Oxford University Press, New York, 1984.
- [144] M. L. MINION, Semi-implicit spectral deferred correction methods for ordinary differential equations, Communications in Mathematical Sciences, 1 (2003), pp. 471–500, doi:10.4310/CMS.2003.v1.n3.a6.
- [145] M. L. MINION, Semi-implicit projection methods for incompressible flow based on spectral deferred corrections, Applied Numerical Mathematics, 48 (2004), pp. 369–387, doi:10.1016/j.apnum.2003.11.005.
- [146] A. MITTAL AND S. MAZUMDER, Generalized ballistic-diffusive formulation and hybrid S_N - P_N solution of the Boltzmann transport equation for phonons for nonequilibrium heat conduction, Journal of Heat Transfer, 133 (2011), p. 092402, doi:10.1115/1.4003961.
- [147] A. MITTAL AND S. MAZUMDER, *Hybrid discrete ordinates–spherical harmonics solution to the Boltz-mann transport equation for phonons for non-equilibrium heat conduction*, Journal of Computational Physics, 230 (2011), pp. 6977–7001, doi:10.1016/j.jcp.2011.05.024.
- [148] M. F. Modest, *The method of spherical harmonics (P_N-approximation)*, in Radiative Heat Transfer, Elsevier, 2013, pp. 495–540, doi:10.1016/b978-0-12-386944-9.50016-9.
- [149] J. MOREL, Diffusion-limit asymptotics of the transport equation, the $P_{1/3}$ equations, and two flux-limited diffusion theories, Journal of Quantitative Spectroscopy and Radiative Transfer, 65 (2000), pp. 769–778, doi:10.1016/S0022-4073(99)00148-X.
- [150] S. MOUSTAFA, I. DUTKA-MALEN, L. PLAGNE, A. PONÇOT, AND P. RAMET, *Shared memory parallelism for 3D Cartesian discrete ordinates solver*, Annals of Nuclear Energy, 82 (2015), pp. 179–187, doi:10.1016/j.anucene.2014.08.034.
- [151] E. J. NYSTRÖM, Über die praktische auflösung von integralgleichungen mit anwendungen auf randwertaufgaben, Acta Mathematica, 54 (1930), pp. 185–204, doi:10.1007/bf02547521.
- [152] G. L. OLSON, L. H. AUER, AND M. L. HALL, *Diffusion, P*₁, and other approximate forms of radiation transport, Journal of Quantitative Spectroscopy and Radiative Transfer, 64 (2000), pp. 619–634, doi:10.1016/s0022-4073(99)00150-8.
- [153] K.-C. PAN, M. LIEBENDÖRFER, M. HEMPEL, AND F.-K. THIELEMANN, *Two-dimensional core-collapse supernova simulations with the isotropic diffusion source approximation for neutrino transport*, The Astrophysical Journal, 817 (2016), p. 72, doi:10.3847/0004-637x/817/1/72.
- [154] N. Papmehl and H. J. Zech, *The order of neutron-scattering method in plane geometry-one-velocity problems*, Nuclear Science and Engineering, 47 (1972), pp. 435–448, doi:10.13182/nse72-a22435.
- [155] S. D. PAUTZ, An algorithm for parallel S_n sweeps on unstructured meshes, Nuclear Science and Engineering, 140 (2002), pp. 111–136, doi:10.13182/NSE02-1.

- [156] V. Pereyna, *Iterated deferred corrections for nonlinear boundary value problems*, Numerische Mathematik, 11 (1968), pp. 111–125, doi:10.1007/bf02165307.
- [157] F. Petrini, G. Fossum, J. Fernandez, A. L. Varbanescu, M. Kistler, and M. Perrone, *Multicore surprises: Lessons learned from optimizing Sweep3D on the cell broadband engine*, in 2007 IEEE International Parallel and Distributed Processing Symposium, IEEE, 2007, doi:10.1109/ipdps.2007.370252.
- [158] P. PICCA, R. FURFARO, AND B. D. GANAPOL, *A hybrid transport point-kinetic method for simulating source transients in subcritical systems*, Annals of Nuclear Energy, 38 (2011), pp. 2680–2688, doi:10.1016/j.anucene.2011.08.005.
- [159] P. PICCA, R. FURFARO, AND B. D. GANAPOL, *Derivation of a physically based hybrid technique for the solution of source-driven time-dependent linear Boltzmann equations*, Transport Theory and Statistical Physics, 41 (2012), pp. 23–39, doi:10.1080/00411450.2012.671219.
- [160] G. C. POMRANING, The Equations of Radiation Hydrodynamics, Pergamon Press, 1973.
- [161] R. W. Preisendorfer, Radiative Transfer on Discrete Spaces, Pergamon Press, 1965.
- [162] W. Qu, N. Brandon, D. Chen, J. Huang, and T. Kress, *A numerical framework for integrating deferred correction methods to solve high order collocation formulations of ODEs*, Journal of Scientific Computing, 68 (2015), pp. 484–520, doi:10.1007/s10915-015-0146-9.
- [163] D. RADICE, E. ABDIKAMALOV, L. REZZOLLA, AND C. D. OTT, *A new spherical harmonics scheme for multi-dimensional radiation transport I. static matter configurations*, Journal of Computational Physics, 242 (2013), pp. 648–669, doi:10.1016/j.jcp.2013.01.048.
- [164] J. RAGUSA, J.-L. GUERMOND, AND G. KANSCHAT, A robust S_N -DG-approximation for radiation transport in optically thick and diffusive regimes, Journal of Computational Physics, 231 (2012), pp. 1947–1962, doi:10.1016/j.jcp.2011.11.017.
- [165] M. RAMPP AND H.-T. JANKA, *Radiation hydrodynamics with neutrinos*, Astronomy & Astrophysics, 396 (2002), pp. 361–392, doi:10.1051/0004-6361:20021398.
- [166] W. H. REED AND T. R. HILL, *Triangular mesh methods for the neutron transport equation*, Tech. Report LA-UR-73-479, Los Alamos Scientific Laboratory, Los Alamos, New Mexico, 1973.
- [167] S. A. RUKOLAINE AND V. S. YUFEREV, *Discrete ordinates quadrature schemes based on the angular interpolation of radiation intensity*, Journal of Quantitative Spectroscopy & Radiative Transfer, 69 (2001), pp. 257–275, doi:10.1016/S0022-4073(00)00079-0.
- [168] Y. SAAD AND M. H. SCHULTZ, *GMRES: A generalized minimal residual algorithm for solving non-symmetric linear systems*, SIAM Journal on Scientific and Statistical Computing, 7 (1986), pp. 856–869, doi:10.1137/0907058.
- [169] R. D. SKEEL, *Iterative refinement implies numerical stability for Gaussian elimination*, Mathematics of Computation, 35 (1980), pp. 817–817, doi:10.1090/s0025-5718-1980-0572859-4.
- [170] I. SLOAN, *Polynomial interpolation and hyperinterpolation over general regions*, Journal of Approximation Theory, 83 (1995), pp. 238–254, doi:10.1006/jath.1995.1119.

- [171] I. H. SLOAN AND R. S. WOMERSLEY, *Constructive polynomial approximation on the sphere*, Journal of Approximation Theory, 103 (2000), pp. 91–118, doi:10.1006/jath.1999.3426.
- [172] J. SMIT, J. CERNOHORSKY, AND C. DULLEMOND, *Hyperbolicity and critical points in two-moment approximate radiative transfer*, Astronomy & Astrophysics, 325 (1997), pp. 203–211.
- [173] R. SPECK, D. RUPRECHT, M. EMMETT, M. MINION, M. BOLTEN, AND R. KRAUSE, *A multi-level spectral deferred correction method*, BIT Numerical Mathematics, 55 (2015), pp. 843–867, doi:10.1007/s10543-014-0517-x.
- [174] R. SPECK, D. RUPRECHT, M. MINION, M. EMMET, AND R. KRAUSE, *Inexact spectral deferred corrections*, in Domain Decomposition Methods in Science and Engineering XXII, T. Dickopf, M. J. Gander, L. Halpern, R. Krause, and L. F. Pavarino, eds., vol. 104 of Lecture Notes in Computational Science and Engineering, Springer International Publishing, 2016, doi:10.1007/978-3-319-18827-0_39.
- [175] C. Syros, A generalization of the multiple collision method, Nukleonik, 8 (1966), pp. 467–471.
- [176] C. Syros, *Solution of the neutron slowing down problem via a multiple collision approach*, Nuclear Science and Engineering, 28 (1967), pp. 203–214, doi:10.13182/nse67-a17470.
- [177] C. Syros, *Time-dependent transport solution in plane geometry*, Atomkernenergie, 16 (1970), pp. 273–280.
- [178] C. P. THURGOOD, A critical evaluation of the discrete ordinates method using HEART and T_N quadrature., PhD thesis, Queen's University at Kingston, 1993.
- [179] C. P. THURGOOD, A. POLLARD, AND H. A. BECKER, *The* T_N *quadrature set for the discrete ordinates method*, Journal of Heat Transfer, 117 (1995), p. 1068, doi:10.1115/1.2836285.
- [180] C.-L. TIEN, A. MAJUMDAR, AND F. M. GERNER, eds., *Microscale Energy Transport*, Taylor & Francis, 1998.
- [181] W. J. WALTERS AND A. HAGHIGHAT, *The adaptive collision source method for discrete ordinates radiation transport*, Annals of Nuclear Energy, 105 (2017), pp. 45–58, doi:10.1016/j.anucene.2017.02.013.
- [182] T. A. WAREING, J. M. MCGHEE, J. E. MOREL, AND S. D. PAUTZ, Discontinuous finite element S_N methods on three-dimensional unstructured grids, Nuclear Science and Engineering, 138 (2001), pp. 256–268, doi:10.13182/nse138-256.
- [183] J. S. Warsa, T. A. Wareing, and J. E. Morel, Krylov iterative methods and the degraded effectiveness of diffusion synthetic acceleration for multidimensional S_N calculations in problems with material discontinuities, Nuclear Science and Engineering, 147 (2004), pp. 218–248, doi:10.13182/nse02-14.
- [184] M. WEISER, *Faster SDC convergence on non-equidistant grids by DIRK sweeps*, BIT Numerical Mathematics, 55 (2014), pp. 1219–1241, doi:10.1007/s10543-014-0540-y.
- [185] G. WIDMER, R. HIPTMAIR, AND C. SCHWAB, *Sparse adaptive finite elements for radiative transfer*, Journal of Computational Physics, 227 (2008), pp. 6071–6105, doi:10.1016/j.jcp.2008.02.025.

- [186] A. B. WOLLABER, *Four decades of implicit Monte Carlo*, Journal of Computational and Theoretical Transport, 45 (2016), pp. 1–70, doi:10.1080/23324309.2016.1138132.
- [187] R. T. WOLLAEGER, D. R. VAN ROSSUM, C. GRAZIANI, S. M. COUCH, G. C. J. IV, D. Q. LAMB, AND G. A. MOSES, *Radiation transport for explosive outflows: A multigroup hybrid Monte Carlo method*, The Astrophysical Journal Supplement Series, 209 (2013), p. 36, doi:10.1088/0067-0049/209/2/36.
- [188] R. YANG, G. CHEN, M. LAROCHE, AND Y. TAUR, Simulation of nanoscale multidimensional transient heat conduction problems using ballistic-diffusive equations and phonon Boltzmann equation, Journal of Heat Transfer, 127 (2005), p. 298, doi:10.1115/1.1857941.
- [189] M. YAVUZ AND E. W. LARSEN, *Spatial domain decomposition for neutron transport problems*, Transport Theory and Statistical Physics, 18 (1989), pp. 205–219, doi:10.1080/00411458908204321.
- [190] W. Zheng and R. G. McClarren, Moment closures based on minimizing the residual of the P_N angular expansion in radiation transport, Journal of Computational Physics, 314 (2016), pp. 682–699, doi:10.1016/j.jcp.2016.03.030.