

Outline

- 1 Background
- 2 MRI Method Development
- 3 Multirate temporal adaptivity
- 4 Results
- 5 Conclusions & future work

Outline


- 1 Background
- 2 MRI Method Development
- 3 Multirate temporal adaptivity
- 4 Results
- 5 Conclusions & future work

Multiphysics Simulations

Multiphysics simulations couple different models either in the bulk or across interfaces.

Climate:


- Atmospheric simulations combine fluid dynamics with local “physics” models for chemistry, condensation, . . .
- Atmosphere is coupled at interfaces to myriad other processes (ocean, land ice, . . .), each using distinct models



Above: <https://e3sm.org>.

Astrophysics/cosmology:

- Dark matter modeled using particles that give rise to large-scale gravitational structures (at right).
- Baryonic matter modeled by combining fluid dynamics, gravity, radiation transport, and reaction networks for chemical ionization states.



Below: <http://svs.gsfc.nasa.gov>.

Multiphysics Challenges

[Keyes et al., 2013]

Since multiphysics simulations combine multiple physical processes, they can challenge textbook methods:

- “Multirate” processes evolve on different time scales but prohibit analytical reformulation.
- Stiff components disallow fully explicit methods (c.f., Prof. Seibold's seminar last week).
- Nonlinearity and low differentiability challenge fully implicit methods.
- Parallel scalability demands optimal algorithms – while robust/scalable algebraic solvers exist for parts (e.g., FMM for particles, multigrid for diffusion), none are optimal for the whole.

We may consider a prototypical problem as having m coupled evolutionary processes:

$$y'(t) = f^{\{1\}}(t, y) + \dots + f^{\{m\}}(t, y), \quad t \in [t_0, t_f], \quad y(t_0) = y_0.$$

Each component $f^{\{k\}}(t, y)$:

- may act on all of y (in the bulk), or on only a subset of y (within a subdomain),
- may evolve on a different characteristic time scale,
- may be “stiff” or “nonstiff,” thereby desiring implicit or explicit treatment.

Implicit-Explicit Additive Runge–Kutta Methods

[Ascher et al., 1997; Kennedy & Carpenter, 2003; ...]

ImEx-ARK methods allow high-order adaptive time integration for additively-split *single rate* simulations:

$$\dot{y}(t) = f^E(t, y) + f^I(t, y), \quad t \in [t_0, t_f], \quad y(t_0) = y_0,$$

- $f^E(t, y)$ contains the nonstiff terms to be treated explicitly,
- $f^I(t, y)$ contains the stiff terms to be treated implicitly.

Combine two s -stage RK methods; denoting $h_n = t_{n+1} - t_n$, $t_{n,j}^E = t_n + c_j^E h_n$, $t_{n,j}^I = t_n + c_j^I h_n$:

$$z_i = y_n + h_n \sum_{j=1}^{i-1} a_{i,j}^E f^E(t_{n,j}^E, z_j) + h_n \sum_{j=1}^i a_{i,j}^I f^I(t_{n,j}^I, z_j), \quad i = 1, \dots, s,$$

$$y_{n+1} = y_n + h_n \sum_{j=1}^s \left[b_j^E f^E(t_{n,j}^E, z_j) + b_j^I f^I(t_{n,j}^I, z_j) \right] \quad (\text{solution})$$

$$\tilde{y}_{n+1} = y_n + h_n \sum_{j=1}^s \left[\tilde{b}_j^E f^E(t_{n,j}^E, z_j) + \tilde{b}_j^I f^I(t_{n,j}^I, z_j) \right] \quad (\text{embedding})$$

Solving each stage z_i , $i = 1, \dots, s$

[Ascher et al., 1997; Kennedy & Carpenter, 2003; ...]

At each stage we must solve a root-finding problem:

$$0 = F_i(z) := \left[z - h_n a_{i,i}^I f^I(t_{n,i}^I, z) \right] - \left[y_n + h_n \sum_{j=1}^{i-1} \left(a_{i,j}^E f^E(t_{n,j}^E, z_j) + a_{i,j}^I f^I(t_{n,j}^I, z_j) \right) \right]$$

- If $f^I(t, y) = J(t)y$ (i.e., f^I is *linear* in y) then this is a large-scale linear system for each z_i :

$$\left(I - h_n a_{i,i}^I J(t_{n,i}) \right) z_i = rhs_i.$$

- Else this requires an iterative solver (e.g., Newton, accelerated fixed-point, or problem-specific), that itself may require solution of multiple linear systems.
- All operators in $f^E(t, y)$ are treated explicitly (do not affect algebraic solver convergence).

ImEx-ARK methods are defined by compatible *explicit* $\{c^E, A^E, b^E, \tilde{b}^E\}$ and *implicit* $\{c^I, A^I, b^I, \tilde{b}^I\}$ tables. These are derived in unison to satisfy order conditions, stability, ...

Multirate Infinitesimal (MRI) methods

[Schlegel et al., 2009; Sandu, 2019; ...]

MRI methods provide a rigorous and highly accurate approach to “subcycling” for multirate problems:

$$\dot{y}(t) = f^S(t, y) + f^F(t, y), \quad t \in [t_0, t_f], \quad y(t_0) = y_0.$$

- $f^S(t, y)$ contains the “slow” dynamics, naturally evolved with large steps H .
- $f^F(t, y)$ contains the “fast” dynamics, that evolves with small steps $h \ll H$.
- It is generally assumed that $f^S(t, y)$ is *considerably more costly to evaluate* than $f^F(t, y)$, and thus a method which evaluates $f^S(t, y)$ infrequently is desirable.
- Extremely efficient – fourth order is attainable with *only a single traversal* of $[t_n, t_{n+1}]$.

MRI Algorithm Outline

Denoting $y_n \approx y(t_n)$, a single explicit MRI step $t_n \rightarrow t_n + H$ proceeds as:

1. Let: $z_1 := y_n$
 2. For $i = 2, \dots, s$:
 - a. Solve: $v_i'(t) = f^F(t, v_i(t)) + r_i(t)$, for $t \in [t_{n,i-1}, t_{n,i}]$ with $v_i(t_{n,i-1}) := z_{i-1}$.
 - b. Let: $z_i := v_i(t_{n,i})$.
 3. Solve: $\tilde{v}_s'(t) = f^F(t, \tilde{v}_s(t)) + \tilde{r}_s(t)$, for $t \in [t_{n,s-1}, t_{n+1}]$ with $\tilde{v}_s(t_{n,s-1}) := z_{s-1}$.
 4. Let: $y_{n+1} := z_s$ (solution) and $\tilde{y}_{n+1} := \tilde{v}_s(t_{n+1})$ (embedding).
- Uniquely defined by: abscissae $0 = c_1 \leq \dots \leq c_s = 1$ ($t_{n,i} := t_n + c_i H$) and forcing functions $r_i(t)$, $\tilde{r}_s(t)$, that are linear combinations of $\{f^S(t_{n,j}, z_j)\}$ and propagate information from slow to fast scales.
 - Implicit versions are possible but require padding for solve-decoupled stages; replace step 2a with a DIRK-like solve during implicit stages.
 - The original MIS methods [Schlegel et al., 2009] used constant $r_i(t)$ and achieved up to $\mathcal{O}(H^3)$.
 - Sandu's MRI-GARK [2019] extended these to $\mathcal{O}(H^4)$ via polynomial $r_i(t)$, and introduced embeddings.

Outline

- 1 Background
- 2 MRI Method Development**
- 3 Multirate temporal adaptivity
- 4 Results
- 5 Conclusions & future work

Implicit-Explicit MRI Methods

We have extended Sandu's MRI-GARK methods to support implicit-explicit treatment of the slow scale:

$$\dot{y}(t) = \underbrace{f^I(t, y) + f^E(t, y)}_{f^S(t, y)} + f^F(t, y), \quad t \in [t_0, t_f], \quad y(t_0) = y_0.$$


ImEx-MRI-GARK methods [Chinomona & R., *SISC*, 2021]:

- Natural extension of MRI-GARK to incorporate f^E into slow time scale.
- Up to $\mathcal{O}(H^4)$ conditions leverage the GARK framework [Sandu & Günther, *SINUM*, 2015].
- *These inherited the requirement for sorted abscissae, thus we struggled to derive embedded and 4th order ImEx-MRI-GARK methods.*



ImEx-MRI-SR methods [Fish, R., & Roberts, *JCAM*, 2024]:

- Circumvent the sorted abscissae requirement by assuming a simpler structure, where each inner solve evolves over $[t_n, t_{n,i}]$.
- Again leverage GARK framework for $\mathcal{O}(H^4)$ conditions.
- There is no “hidden” dependence on $\Delta c_i = 0$ for the stage structure, and no padding is required to derive ImEx-MRI-SR methods from ImEx-ARK.



Multirate Exponential Runge–Kutta (MERK) and Rosenbrock (MERB) Methods

[Luan, Chinomona & R., *SISC*, 2020; Luan, Chinomona & R., *J. Sci. Comput.*, 2022]

To circumvent the explosion in GARK order conditions, we leveraged exponential method theory, replacing the action of the φ_j fcns. with “infinitesimal” IVP solves.



We apply multirate IVP splittings of the form

$$y'(t) = F(t, y) = \mathcal{J}_n y + \mathcal{V}_n t + \mathcal{N}_n(t, y), \quad t \in [t_n, t_n + H], \quad y(t_n) = y_n \in \mathbb{R}^n.$$

- “Fast” scale corresponds to $\mathcal{J}_n y$: for MERK \mathcal{J}_n may be arbitrary, for MERB $\mathcal{J}_n = \frac{\partial F}{\partial y}(t_n, y_n)$.
- “Slow” scale corresponds to $\mathcal{V}_n t + \mathcal{N}_n$ with $\mathcal{V}_n = \frac{\partial F}{\partial t}(t_n, y_n)$ and $\mathcal{N}_n = F - \mathcal{J}_n y - \mathcal{V}_n t$.
- Due to their exponential structure, MERK and MERB only support explicit treatment of f^S .
- Their implementations are nearly identical to ImEx-MRI-SR, albeit without any implicit solves.
- We have embedded MERK methods up to $\mathcal{O}(H^5)$, and non-embedded MERB up to $\mathcal{O}(H^6)$.
- Computational results are available in the ([Appendix](#)).

Outline

- 1 Background
- 2 MRI Method Development
- 3 Multirate temporal adaptivity**
- 4 Results
- 5 Conclusions & future work

Single rate control

[Gustafsson, 1991; Söderlind, 2006]

- Traditional adaptivity controls *local error*, $\ell_n := y(t_n + H_n) - y_{n+1}$, by assuming y_n is exact and adapting H_n to ensure $\varepsilon_n := \|\ell_n\| \leq 1$ where the norm incorporates the user tolerances,

$$\|x\| := \left(\frac{1}{N} \sum_{i=1}^N \left(\frac{x_i}{\text{rtol}|y_{n,i}| + \text{atol}_i} \right)^2 \right)^{1/2}$$

- A “controller” \mathcal{C} typically depends on a few (H_k, ε_k) , $\tilde{H} = \mathcal{C}(H_n, \varepsilon_n, H_{n-1}, \varepsilon_{n-1}, \dots; p)$, where p is the global method order, i.e., $\varepsilon_n \leq c(t) H_n^{p+1}$, for some $c(t)$ independent of H_n .
- The simple *I controller* assumes equality above, and a piecewise constant $c = \frac{\varepsilon_n}{H_n^{p+1}}$, to predict \tilde{H} :

$$1 = c\tilde{H}^{p+1} = \varepsilon_n \frac{\tilde{H}^{p+1}}{H_n^{p+1}} \quad \Leftrightarrow \quad \tilde{H} = \frac{H_n}{\varepsilon_n^{1/(p+1)}}.$$


- More advanced options exist, that typically use additional (H_k, ε_k) values to build higher-degree piecewise polynomial approximations of the principal error function.

For multirate control, we thus require both a strategy to estimate local temporal errors ε_n^S and ε_n^F , and algorithms for selecting step sizes H and h .

MRI time step control – Decoupled (*Dec*) controllers

The simplest approach to MRI adaptivity uses decoupled single-rate controllers:

$$\begin{aligned}\tilde{H} &= \mathcal{C}^S(H_n, \varepsilon_n^S, H_{n-1}, \varepsilon_{n-1}^S, \dots; P), \\ \tilde{h} &= \mathcal{C}^F(h_{n,m}, \varepsilon_{n,m}^F, h_{n,m-1}, \varepsilon_{n,m-1}^F, \dots; p),\end{aligned}$$



where (H_k, ε_k^S) are the stepsize and local error estimates for time step k at the slow time scale, and $(h_{k,l}, \varepsilon_{k,l}^F)$ are the stepsize and local error estimates for the fast substep l within the slow step k .

- \mathcal{C}^S and \mathcal{C}^F are distinct, so selection of \tilde{H} and \tilde{h} occurs independently.
- We expect this to work well for problems with weakly coupled time scales.
- Due to its decoupled nature, this trivially extends to an arbitrary number of time scales, allowing adaptivity for so-called “telescopic” MRI methods.

MRI time step control – Step-tolerance ($H\text{-}Tol$) controllers

For problems with more strongly coupled scales, we may want tighter accuracy, tolfac_n , from the inner solver. When called over the slow step $[t_n, t_n + H]$, we assume the accumulated fast error satisfies

$$\varepsilon_i^F = \chi(t_n) H_n (\text{tolfac}_n),$$

where $\chi(t_n)$ is independent of tolfac_n , but may vary with time.

This fits the asymptotic error assumption $\varepsilon_n = c(t)h^{p+1}$, through identifying $\chi(t_n)H_n$ with $c(t)$, tolfac_n with h , and $p = 0$. Thus a single-rate controller could adjust tolfac_n between slow step attempts, and we construct an “ $H\text{-}Tol$ ” MRI controller from three single-rate controllers:

- $\mathcal{C}^{S,H}(H_n, \varepsilon_n^S, H_{n-1}, \varepsilon_{n-1}^S, \dots; P)$ – adapts H_n to achieve user-requested solution tolerances
- $\mathcal{C}^{S,Tol}(\text{tolfac}_n, \varepsilon_n^F, \text{tolfac}_{n-1}, \varepsilon_{n-1}^F, \dots; 0)$ – adapts tolfac_n using the strategy described above
- $\mathcal{C}^F(h_{n,m}, \varepsilon_{n,m}^F, h_{n,m-1}, \varepsilon_{n,m-1}^F, \dots; p)$ – adapts inner time steps $h_{n,m}$ to achieve tolfac_n

This class of controllers also support telescopic multirate methods.

MRI time step control – Coupled (H - h) controllers[Fish & R., *SISC*, 2023]

This pre-existing class extended the single-rate derivations of [Gustafsson, 1994] by approximating both slow and fast principal error functions using piecewise polynomials. Four MRI controllers were proposed that simultaneously adapt H_n and $M_n = H_n/h_n$:

- *constant-constant*: $H_{n+1} = H_n (\varepsilon_{n+1}^S)^\alpha$, $M_{n+1} = M_n (\varepsilon_{n+1}^S)^{\beta_1} (\varepsilon_{n+1}^F)^{\beta_2}$,

- *linear-linear*: $H_{n+1} = H_n \left(\frac{H_n}{H_{n-1}} \right) (\varepsilon_{n+1}^S)^{\alpha_1} (\varepsilon_n^S)^{\alpha_2}$,

$$M_{n+1} = M_n \left(\frac{M_n}{M_{n-1}} \right) (\varepsilon_{n+1}^S)^{\beta_{11}} (\varepsilon_n^S)^{\beta_{12}} (\varepsilon_{n+1}^F)^{\beta_{21}} (\varepsilon_n^F)^{\beta_{22}}.$$

- *PIMR* (a multirate extension of the PI single-rate controller):

$$H_{n+1} = H_n (\varepsilon_{n+1}^S)^{\alpha_1} (\varepsilon_n^S)^{\alpha_2}, \quad M_{n+1} = M_n (\varepsilon_{n+1}^S)^{\beta_{11}} (\varepsilon_n^S)^{\beta_{12}} (\varepsilon_{n+1}^F)^{\beta_{21}} (\varepsilon_n^F)^{\beta_{22}}.$$

- *PIDMR* (a multirate extension of the PID single-rate controller):

$$H_{n+1} = H_n (\varepsilon_{n+1}^S)^{\alpha_1} (\varepsilon_n^S)^{\alpha_2} (\varepsilon_{n-1}^S)^{\alpha_3},$$

$$M_{n+1} = M_n (\varepsilon_{n+1}^S)^{\beta_{11}} (\varepsilon_n^S)^{\beta_{12}} (\varepsilon_{n-1}^S)^{\beta_{13}} (\varepsilon_{n+1}^F)^{\beta_{21}} (\varepsilon_n^F)^{\beta_{22}} (\varepsilon_{n-1}^F)^{\beta_{23}}.$$

Outline

- 1 Background
- 2 MRI Method Development
- 3 Multirate temporal adaptivity
- 4 Results**
- 5 Conclusions & future work

Kværno-Prothero-Robinson (KPR) test problem

$$\begin{pmatrix} u'(t) \\ v'(t) \end{pmatrix} = \begin{bmatrix} G & e_s \\ e_f & -1 \end{bmatrix} \begin{pmatrix} (u^2 - p - 2)/(2u) \\ (v^2 - q - 2)/(2v) \end{pmatrix} + \begin{pmatrix} p'(t)/(2u) \\ q'(t)/(2v) \end{pmatrix},$$


over $0 \leq t \leq 5$, where $p(t) = \cos(t)$ and $q(t) = \cos(\omega t(1 + e^{-(t-2)^2}))$.

The analytical solution is $u(t) = \sqrt{2 + p(t)}$ and $v(t) = \sqrt{2 + q(t)}$.

- e_s determines the coupling from the fast to the slow time scale,
- e_f determines the coupling from the slow to the fast time scale,
- $G < 0$ determines the stiffness at slow time scale,
- ω that determines the time-scale separation factor.

Top right: analytical solutions with $G = -10$, $e_s = e_f = 1/10$, $\omega = 5$.

Bottom right: internal single-rate time steps.



Stiff Brusselator test problem


$$\begin{pmatrix} u'(t) \\ v'(t) \\ w'(t) \end{pmatrix} = \begin{pmatrix} a + vu^2 - (w+1)u \\ -vu^2 + wu \\ \frac{b-w}{\epsilon} - wu \end{pmatrix}, \quad 0 \leq t \leq 10.$$

Initial condition $(1.2, \quad 3.1, \quad 3)^T$; parameters $a = 1, b = 3.5$.

- ϵ is the stiffness parameter. MRI methods are often used to circumvent implicit solves using explicit substepping.
- Sharp solution transition at $t \approx 6.5$ stresses slow time scale temporal adaptivity.

Top right: example solution with $\epsilon = 1/2500$.

Bottom right: internal single-rate time steps.



Embedded MRI methods

We test MRI adaptivity using 15 embedded MRI methods of varying order and type:

- MRI-GARK methods from [Sandu, 2019]:
 - Explicit ERK22a, ERK22b, ERK33a, and ERK45a methods (orders 2, 2, 3, 4)
 - Implicit IRK21a, ESDIRK34a, and ESDIRK46a methods (orders 2, 3, 4)
- Explicit RALSTON2 MRI-GARK method (order 2) from [Roberts, 2022]
- IMEXSR21, IMEXSR32, and IMEXSR43 ImEx-MRI-SR methods (orders 2, 3, 4) from [Fish et al., 2024]
- MERK21, MERK32, MERK43, and MERK54 explicit MERK methods (orders 2, 3, 4, 5) from [Luan et al., 2020] (custom embeddings)

Each of the above methods include an embedding with order of accuracy one lower.

MRI adaptive accuracy – KPR


We compare the ability of each method to achieve the target accuracy over all components l and time steps n for the KPR test: $\text{accuracy} = \max_{n,l} \left| \frac{y_{n,l} - y_{ref,l}(t_n)}{\text{abstol} + \text{reitol} |y_{ref,l}(t_n)|} \right|$.

Left order 2,
Center order 3,
Right orders 4-5

Shading includes
all results from a
family.

H-Tol and *Dec*
robust across all
tests.

H-h accuracy
varies wildly
across methods.




MRI adaptive accuracy – Stiff Brusselator

We repeat the previous experiment for the stiff Brusselator test:

H-Tol and *Dec*
robust for all
but ESDIRK34a,
ESDIRK46a


H-h fails for all
but MERK21 and
MERK32.



MRI adaptive efficiency

We ran 4200 tests varying problem setup, MRI method, adaptivity controller, and tolerance, to determine the most “efficient” combinations. The optimal embedded MRI methods varied for each test problem and choice of work metric (so we recommend using libraries like ARKODE that allow experimentation), but some controllers rose to the top (H - h controllers never won).

Below: best $\mathcal{O}(H^3)$ methods for stiff Brusselator.



MRI adaptive efficiency statistics


We computed z-scores for each controller's average rank when holding other variables constant (MRI method, multirate ratio, etc.). Ranking these separately for slow time scale work (left) and fast time scale work (right):

Slow Scale		Fast Scale	
Multirate controller	z-score	Multirate controller	z-score
HT-I	-0.691	D-I	-0.558
HT-H0321	-0.350	D-H312	-0.419
HT-H312	-0.276	D-H211	-0.405
HT-H0211	-0.265	D-H0321	-0.180
HT-H211	-0.254	D-H0211	0.035
D-I	0.065	HT-I	0.150
D-H312	0.308	HT-H211	0.190
D-H211	0.422	HT-H312	0.246
D-H0321	0.444	HT-H0321	0.441
D-H0211	0.599	HT-H0211	0.499

MRI adaptive step histories – KPR

We plot the slow and fast step size histories for the HT-H211, D-H211 and MRI-CC controllers using MRI_ERK33a at $\text{reltol} = 10^{-5}$, listing the total numbers of slow and fast time steps, and accuracy ratio.

As expected, HT-H211 and D-H211 achieve reasonable accuracy, with HT-H211 moving some work from the slow to fast scales. MRI-CC handles slow adaptivity relatively well, but does not allow sufficient fast time scale separation, leading to large error.



Use in Real-Time Boltzmann Transport collaboration

Accelerating Real-Time Boltzmann Transport Equation Simulations



Caltech

With the Death Valley SciDAC BES Partnership

Scientific Achievement


- Electronic structure calculations in the time domain provide a deeper understanding of nonequilibrium dynamics in materials.
- The real-time Boltzmann equation (rt-BTE) enables reliable predictions of coupled electron and lattice dynamic however, the different time scales of the system and computational cost of collision integrals make simulating long times intractable.
- Utilizing adaptive and multirate time integrators from the SUNDIALS library leads to at least 10x faster simulations and enables long time simulations of bulk materials.

Significance and Impact

- Coherent control of electronic and lattice degrees of freedom is a novel frontier in materials physics that requires understanding processes across different time scales.
- Efficient computational methods are necessary to unravel the microscopic mechanisms of ultrafast dynamics and quantitatively interpret time-domain spectroscopies.
- Advanced time integration methods make studies of ultrafast electronic and lattice dynamics in materials driven by optical or terahertz pulses possible.

Technical Approach

- Multirate methods improve computational efficiency by advancing slow, expensive processes with large time steps while fast, cheap processes use smaller steps.
- Adaptive methods vary the step size to track the changing dynamics of the problem.



Comparison of the solution accuracy versus run time simulating coupled electron and lattice dynamics in graphene with fixed step sizes (orange), a single rate method with adaptive step sizes (gray), and a multirate method with adaptive fast step sizes. Points closer to the lower left corner are more efficient (lower error and less run time). The multirate methods provides at least a 10x speedup over fixed steps.

PI(s)/Facility Lead(s): Carol Woodward
 Collaborating Institutions: LLNL, Caltech
 Program: BES SciDAC Partnership
 PM: Marco Fornari (ASCR) and Matthias Graf (BES)
 Publication for this work: J. Yao, I. Maliyov, D. J. Gardner, C. S. Woodward, and M. Bernardi, "Adaptive and multirate time integration methods for the real-time Boltzmann transport equation," to appear in Nature Materials, arXiv: 2412.13402



U.S. DEPARTMENT OF
ENERGY

Office of
Science

Use in Real-Time Boltzmann Transport collaboration

Accelerating Real-Time Boltzmann Transport Equation Simulations



Caltech

With the Death Valley SciDAC BES Partnership

Multirate adaptivity

Sci

- Elec
- non
- The
- elec
- cor
- Utili
- leas

More recently, we've applied our *Decoupled* controllers for slow step adaptivity, which improves on the best green dot/cross:


- $\sim 1.5\times$ runtime speedup
- $\sim 1.5\times$ improved carrier (electron) accuracy
- $\sim 20\times$ improved phonon accuracy

Sign

- Coherent control of electronic and lattice degrees of freedom is a novel frontier in materials physics that requires understanding processes across different time scales.
- Efficient computational methods are necessary to unravel the microscopic mechanisms of ultrafast dynamics and quantitatively interpret time-domain spectroscopies.
- Advanced time integration methods make studies of ultrafast electronic and lattice dynamics in materials driven by optical or terahertz pulses possible.

Technical Approach

- Multirate methods improve computational efficiency by advancing slow, expensive processes with large time steps while fast, cheap processes use smaller steps.
- Adaptive methods vary the step size to track the changing dynamics of the problem.



Comparison of the solution accuracy versus run time simulating coupled electron and lattice dynamics in graphene with fixed step sizes (orange), a single rate method with adaptive step sizes (gray), and a multirate method with adaptive fast step sizes. Points closer to the lower left corner are more efficient (lower error and less run time). The multirate methods provides at least a 10x speedup over fixed steps.

PI(s)/Facility Lead(s): Carol Woodward
 Collaborating Institutions: LLNL, Caltech
 Program: BES SciDAC Partnership
 PM: Marco Fornari (ASCR) and Matthias Graf (BES)
 Publication for this work: J. Yao, I. Maliyov, D. J. Gardner, C. S. Woodward, and M. Bernardi, "Adaptive and multirate time integration methods for the real-time Boltzmann transport equation," to appear in Nature Materials, arXiv: 2412.13402



U.S. DEPARTMENT OF
ENERGY

Office of
Science

Outline

- 1 Background
- 2 MRI Method Development
- 3 Multirate temporal adaptivity
- 4 Results
- 5 Conclusions & future work

Conclusions and future work

- Today's rich ecosystem of scientific applications requires an equally rich “toolkit” of solvers.
- Novel time integration methods aim to achieve higher order, increased flexibility, improved stability, and robust temporal adaptivity.
- Here, I focused on multirate adaptivity:
 - Both the *Dec* and *H-Tol* families show robust adaptive control across a wide range of problem types, tolerances, and MRI methods.
 - By shifting some burden to the inner solver, *H-Tol* is optimal for problems where the slow operator dominates the cost.
 - *H-h* controllers struggle since they artificially constrain the step size ratio $M_n = H_n/h_n$.
- All MRI methods in this talk (except MERB), along with the *Dec* and *H-Tol* families, are available within the ARKODE solver from the **SUNDIALS** library.
- We are actively testing and optimizing these methods on large-scale application codes, including *Perturbo* (solid state physics) and *BOUT++* (fusion).



Suite of Nonlinear and Differential-Algebraic Solvers

- Library of time integrators and nonlinear solvers
 - Packages: CVODE(S), ARKODE, IDA(S), and KINSOL
 - Written in C with modern Fortran interfaces
 - Designed to be easily incorporated into existing codes
- Modular implementation
 - Data manipulation and parallelism are encapsulated by vector, matrix, and solver classes and user-supplied callback functions
 - Native class implementations for serial, threaded, distributed, and GPU computing platforms
 - Vector, matrix, and solver classes can all be user-supplied
- Freely available under BSD 3-Clause license
 - 100,000+ downloads per year: github.com/LLNL/sundials
 - Detailed user manuals: sundials.readthedocs.io
 - Active user community supported by email list and GitHub issues

Pele Combustion Simulations on Frontier



Isosurfaces of diesel fuel entering a turbulent methane-air premixture. High temperature pockets (red and yellow) form when local kernels of diesel fuel ignite.

This simulation ran on 7,000 Frontier nodes at OLCF using SUNDIALS to solve the chemistry systems in every grid cell in a 7-layer adaptive mesh hierarchy (60B grid cells and approximately 2.4T degrees of freedom).

Courtesy of Marc Day and Jon Rood (NREL). Animation by Nicholas Brunhart-Lupo (NREL).

SUNDIALS Solvers

Adaptive Methods for Ordinary Differential Equation (ODE) and Differential-Algebraic Equation (DAE) Initial Value Problems (IVPs)

- **CVODE**: adaptive order and step size linear multistep methods for ODEs, $y' = f(t, y)$
 - Adams methods for non-stiff systems and BDF methods for stiff systems
- **IDA**: adaptive order and step size BDF methods for DAEs, $F(t, y, \dot{y}) = 0$
 - Targets implicit ODEs, index-1 DAEs, and Hessenberg index-2 DAEs
- **CVODES** and **IDAS** support forward and adjoint sensitivity analysis (user-supplied adjoint operator)
- **ARKODE**: infrastructure for adaptive step multistage methods

Stepper	Methods	Systems
ARKStep	ERK, DIRK, and IMEX	$M(t) y' = f^E(t, y) + f^I(t, y)$
ERKStep	ERK (streamlined module)	$y' = f(t, y)$
LSRKStep	Low-storage SSP and STS	$y' = f(t, y)$
MRISStep	Multirate infinitesimal (MRI) step	$y' = f^S(t, y) + f^F(t, y)$
SPRKStep	Symplectic partitioned RK	$p' = f^p(t, q), \quad q' = f^q(t, p)$
SplittingStep	Operator splitting	$y' = f^1(t, y) + \dots + f^N(t, y)$

SUNDIALS team

Current SUNDIALS Team

Current Team:



Cody Balos




David Gardner



Alan Hindmarsh



Dan Reynolds



Steven Roberts



Carol Woodward


Postdocs:



Mustafa Aggul



Sylvia Amihere



Yifan Hu

Alumni:



Radu Serban

Scott D. Cohen, Peter N. Brown, George Byrne, Allan G. Taylor, Steven L. Lee, Keith E. Grant, Aaron Collier, Lawrence E. Banks, Steve G. Smith, Cosmin Petra, Homer Walker, Slaven Peles, John Loffeld, Dan Shumaker, Ulrike M. Yang, James Almgren-Bell, Shelby L. Lockhart, Rujeko Chinomona, Daniel McGreer, Hunter Schwartz, Hilari C. Tiedeman, Ting Yan, Jean M. Sexton, and Chris White

6 Appendix

Implicit-Explicit Multirate Infinitesimal GARK Methods

[Chinomona & R., *SISC*, 2021]

Extended Sandu's MRI-GARK methods to support implicit-explicit treatment of the slow time scale:

$$\dot{y}(t) = \underbrace{f^I(t, y) + f^E(t, y)}_{f^S(t, y)} + f^F(t, y), \quad t \in [t_0, t_f], \quad y(t_0) = y_0.$$



These define an ImEx forcing function


$$r_i(t) = \sum_{j=1}^i \gamma_{i,j} \left(\frac{t-t_n}{(c_i-c_{i-1})H} \right) f^I(t_{n,j}, z_j) + \sum_{j=1}^{i-1} \omega_{i,j} \left(\frac{t-t_n}{(c_i-c_{i-1})H} \right) f^E(t_{n,j}, z_j),$$

- Order conditions up to $\mathcal{O}(H^4)$ leverage the GARK framework [Sandu & Günther, *SINUM*, 2015].
- *These inherited the requirement for sorted abscissae, thus we struggled to derive embedded and 4th order ImEx-MRI-GARK methods.*

Begin with an ImEx-ARK pair $\{A^I, b^I, c^I; A^E, b^E, c^E\}$ where $c^I = c^E \equiv c$ with $0 = c_1 \leq \dots \leq c_s \leq 1$.

- Convert to solve-decoupled form: insert redundant stages such that $\Delta c_i A_{ii}^I = 0$ for $i = 1, \dots, s$.
- Extend A^I , A^E and c to ensure “stiffly-accurate” condition: $c_s = 1$, $A_{s,:}^I = b^I$, $A_{s,:}^E = b^E$.
- Generate $\Gamma^{(k)}$ and $\Omega^{(k)}$ for $k = 0, \dots, k_{max}$, to satisfy ARK consistency (s^2 conditions), internal consistency ($2s(k_{max} + 1)$ conditions), plus order conditions:
 - $\mathcal{O}(H^1)$ and $\mathcal{O}(H^2)$: no additional order conditions,
 - $\mathcal{O}(H^3)$: 2 additional order conditions,
 - $\mathcal{O}(H^4)$: 16 additional order conditions.
- With any additional degrees of freedom, we maximized “joint linear stability”.

Note: we found it challenging to construct embedded ImEx-MRI-GARK methods, largely due to our reliance on ImEx-ARK base methods and the “sorted” abscissa requirement.



Dr. David J. Nisbet is an Associate Professor of Biology at the University of North Carolina at Chapel Hill. He is a member of the American Society of Naturalists and the American Society of Systematic Biologists. He is also a past president of the American Society of Systematic Biologists.

- 5/8


Again take an an ImEx-ARK pair $\{A^I, b^I, c^I; A^E, b^E, c^E\}$ where $c^I = c^E \equiv c$ (not necessarily sorted).

- Extend A^I , A^E and c to ensure “stiffly-accurate” condition: $c_s = 1$, $A_{s,:}^I = b^I$, $A_{s,:}^E = b^E$.
- Generate Γ and $\Omega^{(k)}$ for $k = 0, \dots, n_\Omega$, to satisfy ImEx-ARK consistency (s^2 conditions), internal consistency ($s(2 + n_\Omega)$ conditions), plus order conditions:
 - $\mathcal{O}(H^1)$ and $\mathcal{O}(H^2)$: no additional order conditions,
 - $\mathcal{O}(H^3)$: 1 additional order condition,
 - $\mathcal{O}(H^4)$: 6 additional order conditions.
- With remaining degrees of freedom, maximize joint linear stability for the method and minimize leading order error for embedding.


ImEx-MRI-SR Convergence/Efficiency – Stiff Brusselator PDE

[Fish, R., & Roberts, *JCAM*, 2024]

Fixed step size runtime efficiency: ImEx-MRI-SR, ImEx-MRI-GARK, and ImEx-MRI versions of Lie–Trotter and Strang–Marchuk methods.



Adaptive ImEx-MRI-SR efficiency: modified problem with time-dependent advection, diffusion and reaction coefficients, using tolerances 10^{-k} with $k = 1, \dots, 9$:



MERK and MERB Convergence/Efficiency – Reaction-Diffusion PDE

[Luan, Chinomona & R., *SISC*, 2020; Luan, Chinomona & R., *J. Sci. Comput.*, 2022]]

Problem: $u_t = \epsilon u_{xx} + \gamma u^2(1 - u)$, $x \in (0, 5)$, $t \in [0, 5]$,

with $\gamma = 0.1$, $\epsilon = 10^{-2}$, $\lambda = \sqrt{5}$, $u(x, 0) = (1 + \exp(\lambda(x - 1)))^{-1}$, and $u_x(0, t) = u_x(5, t) = 0$.

Efficiency plots (runtime, RHS calls, \mathcal{N}_n calls):

- Top-right: $\mathcal{O}(H^3)$ methods
- Bottom-left: $\mathcal{O}(H^4)$ methods
- Bottom-right: $\mathcal{O}(H^5)$ and $\mathcal{O}(H^6)$ methods
- Methods* use a “natural” splitting; others use dynamic linearization.

