

Lawrence Livermore National Laboratory

# SUNDIALS: Suite of Nonlinear and Differential/Algebraic Equation Solvers



**Carol S. Woodward**

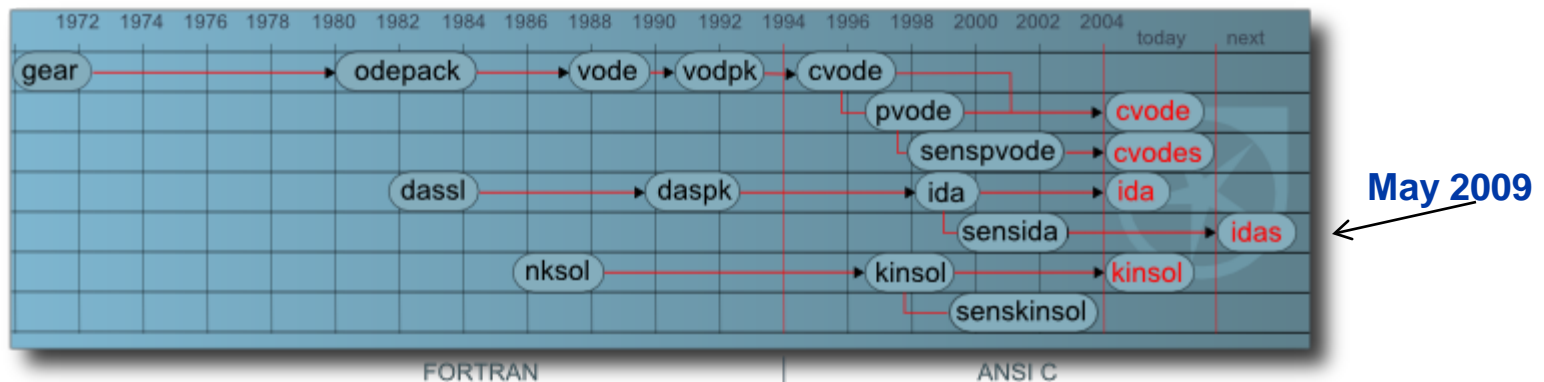
# Outline

- SUNDIALS Overview
- ODE and DAE integration
  - Initial value problems
  - Implicit integration methods
- Nonlinear Systems
  - Newton's method and inexact Newton's method
  - Preconditioning
- SUNDIALS: usage, applications, and availability
- Upcoming additions



# LLNL has a long history of R&D in ODE/DAE methods and software

- Fortran solvers written at LLNL:
  - VODE: stiff/nonstiff ODE systems, with direct linear solvers
  - VODPK: with Krylov linear solver (GMRES)
  - NKSOL: Newton-Krylov solver - nonlinear algebraic systems
  - DASPK: DAE system solver (from DASSL)
- Recent focus has been on sensitivity analysis
- Organized into a single suite, **SUNDIALS**, written in C and including CVODE and CVODES, IDA, IDAS, and KINSOL



# Push to solve large, parallel systems motivated rewrites in C

- **CVODE**: rewrite of VODE/VODPK [Cohen, Hindmarsh, 94]
- **PVODE**: parallel CVODE [Byrne and Hindmarsh, 98]
- **KINSOL**: rewrite of NKSOL [Taylor and Hindmarsh, 98]
- **IDA**: rewrite of DASPK [Hindmarsh and Taylor, 99]
- Sensitivity variants: **SensPVODE**, **SensIDA**, **SensKINSOL** [Brown, Grant, Hindmarsh, Lee, 00-01]
- New sensitivity-capable solvers:
  - **CVODES** [Hindmarsh and Serban, 02]
  - **IDAS** [Serban, Petra, and Hindmarsh, 09]
- Organized into a single suite, **SUNDIALS**, including CVODE and CVODES, IDA, IDAS, and KINSOL



# The SUNDIALS package offers Newton solvers, time integration, and sensitivity solvers

- **CVODE: implicit ODE solver,  $y' = f(y, t)$** 
  - Variable-order, variable step BDF (stiff) or implicit Adams (nonstiff)
  - Nonlinear systems solved by Newton or functional iteration
  - Linear systems by direct (dense or band) or iterative solvers
- **IDA: implicit DAE solver,  $F(t, y, y') = 0$** 
  - Variable-order, variable step BDF
  - Nonlinear system solved by Newton iteration
  - Linear systems by direct (dense or band) or iterative solvers
- **KINSOL: Newton solver,  $F(u) = 0$** 
  - Inexact and Modified (with dense solve) Newton
  - Linear systems by iterative or dense direct solvers
- **CVODES: sensitivity-capable (forward & adjoint) CVODE**
- **IDAS: sensitivity-capable (forward & adjoint) IDA**
- **Iterative linear Krylov solvers: GMRES, BiCGStab, TFQMR**



# SUNDIALS was designed to easily interface with legacy codes

- **Philosophy: *Keep codes simple to use***
- **Written in C**
  - **Fortran interfaces: FCVODE, FIDA, and FKINSOL**
  - **Matlab interfaces: sundialsTB (CVODES, IDA, & KINSOL)**
- **Written in a **data structure neutral** manner**
  - **No specific assumptions about data**
  - **Application-specific data representations can be used**
- **Modular implementation**
  - **Vector modules**
  - **Linear solver modules**
- **Require minimal problem information, but offer user control over most parameters**



# Initial value problems (IVPs) come in the form of ODEs and DAEs

- The general form of an IVP is given by

$$\begin{aligned} F(t, \dot{\mathbf{x}}, \mathbf{x}) &= 0 \\ \mathbf{x}(t_0) &= \mathbf{x}_0 \end{aligned}$$

- If  $\partial F / \partial \dot{\mathbf{x}}$  is invertible, we solve for  $\dot{\mathbf{x}}$  to obtain an **ordinary differential equation (ODE)**, but this is not always the best approach
- Else, the IVP is a **differential algebraic equation (DAE)**
- A DAE has **differentiation index  $i$**  if  $i$  is the minimal number of analytical differentiations needed to extract an explicit ODE

# Stiffness of an equation can significantly impact whether implicit methods are needed

- (Ascher and Petzold, 1998): If the system has widely varying time scales, and the phenomena that change on fast scales are *stable*, then the problem is **stiff**
- Stiffness depends on
  - Jacobian eigenvalues,  $\lambda_j$
  - System dimension
  - Accuracy requirements
  - Length of simulation
- In general a problem is stiff on  $[t_0, t_1]$  if

$$(t_1 - t_0) \min_j \Re(\lambda_j) \ll -1$$



# Dalquist test problem shows impact of stability on step sizes for explicit and implicit methods

Dalquist test equation:  $\dot{\mathbf{y}} = \lambda \mathbf{y}, \mathbf{y}(0) = \mathbf{y}_0$

Exact solution:  $\mathbf{y}(t_n) = \mathbf{y}_0 e^{\lambda t_n}$

Absolute stability requirement

$$|\mathbf{y}_n| \leq |\mathbf{y}_{n-1}|, \quad n = 1, 2, \dots$$

If  $\text{Re}(\lambda) < 0$ , then  $|\mathbf{y}(t_n)|$  decays exponentially; we cannot tolerate growth in the approximate solution  $\mathbf{y}_n$

Region of absolute stability of an integrator written as:

$\mathbf{y}_n = R(z)\mathbf{y}_{n-1}$ , with time step  $z = h\lambda$

$$\mathbf{S} = \{z \in \mathbf{C}; |R(z)| \leq 1\}$$

# Forward and backward Euler show different stability restrictions

- Forward Euler:  $\mathbf{y}_n = \mathbf{y}_{n-1} + h(\lambda \mathbf{y}_{n-1}) \Rightarrow R(\mathbf{z}) = |1 + h\lambda|$

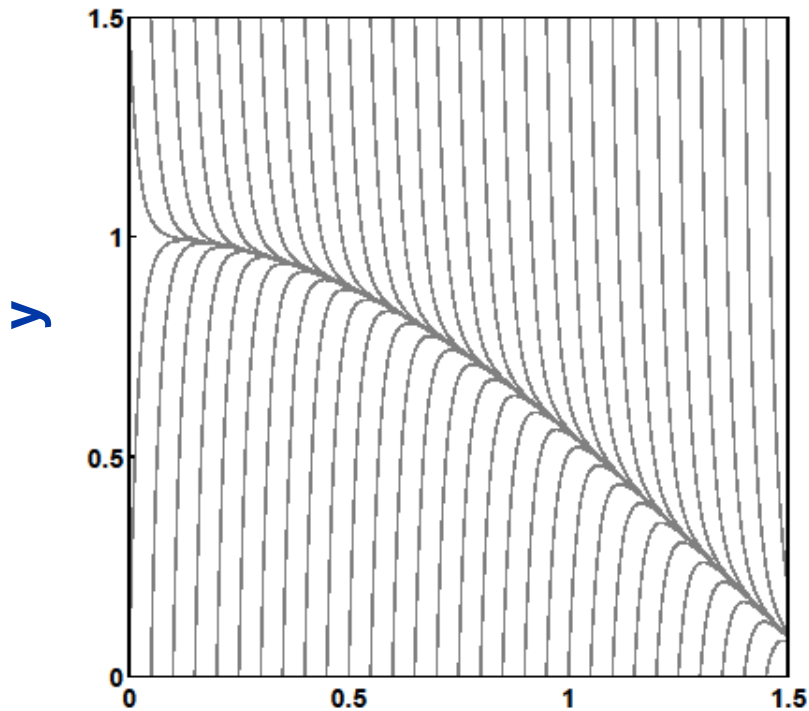
So, if  $\lambda < 0$ , FE has the step size restriction:  $h \leq \frac{2}{|\lambda|}$

- Backward Euler:  $\mathbf{y}_n = \mathbf{y}_{n-1} + h(\lambda \mathbf{y}_n) \Rightarrow R(\mathbf{z}) = \left| \frac{1}{1 - h\lambda} \right|$

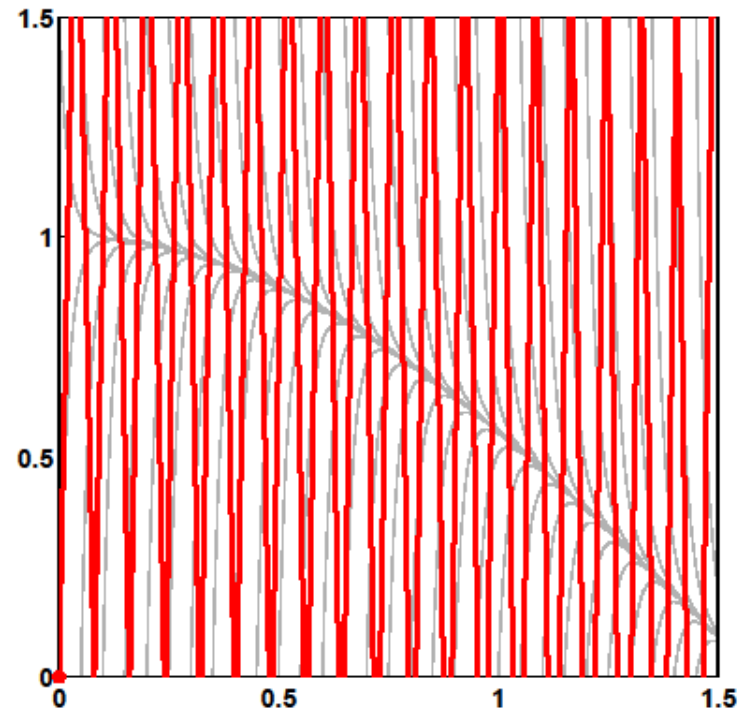
So, if  $\lambda < 0$ , BE has the step size restriction:  $h > 0$

# Curtiss and Hirschfelder example

$$\dot{y} = -50(y - \cos(t)) \quad \lambda = -50$$



**Solution curves**

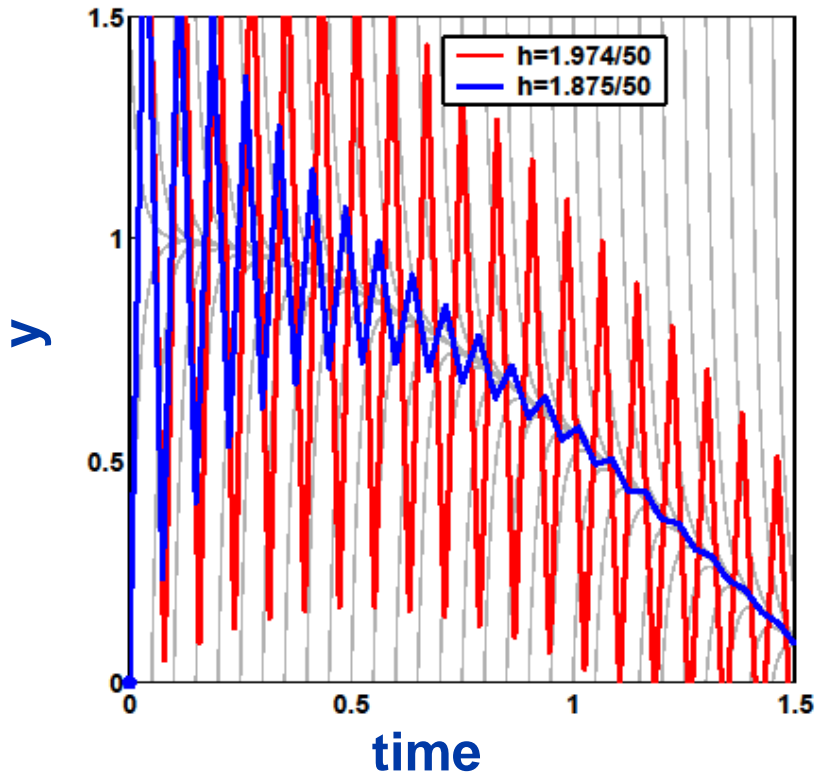


**Forward Euler**

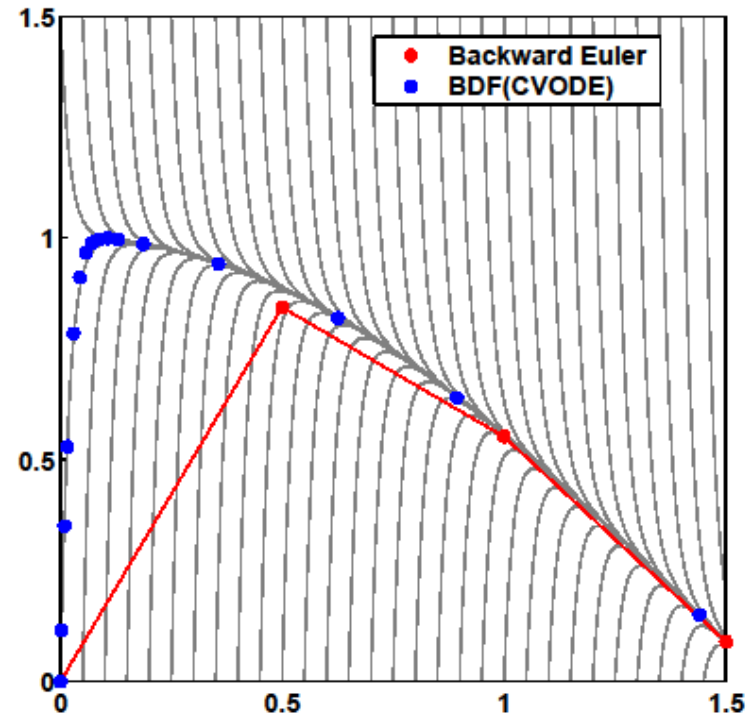
**$h=2.01/50$**

# Curtiss and Hirschfelder example

$$\dot{y} = -50(y - \cos(t)) \quad \lambda = -50$$



**Forward Euler**



**Implicit schemes**

**$h=0.5$  for BE**



# SUNDIALS has implementations of Linear Multistep Methods (LMM)

General form of LMM: 
$$\sum_{i=0}^{K_1} \alpha_{n,i} \mathbf{y}_{n-i} + h_n \sum_{i=0}^{K_2} \beta_{n,i} \dot{\mathbf{y}}_{n-i} = \mathbf{0}$$

- Two methods:
  - Adams-Moulton (nonstiff);  $K_1 = 1, K_2 = k, k = 1, \dots, 12$
  - BDF (stiff);  $K_1 = k, K_2 = 0, k = 1, \dots, 5$

- Nonlinear systems (BDF)

- ODE:

$$\dot{\mathbf{y}} = \mathbf{f}(\mathbf{y}) \quad \mathbf{G}(\mathbf{y}_n) \equiv \mathbf{y}_n - \beta_0 h_n \mathbf{f}(t, \mathbf{y}_n) - \sum_{i=1}^k \alpha_{n,i} \mathbf{y}_{n-i} = \mathbf{0}$$

- DAE:

$$\mathbf{F}(\dot{\mathbf{y}}, \mathbf{y}) = \mathbf{0} \quad \mathbf{G}(\mathbf{y}_n) \equiv \mathbf{F}\left(t, (\beta_0 h_n)^{-1} \sum_{i=1}^k \alpha_{n,i} \mathbf{y}_{n-i}, \mathbf{y}_n\right) = \mathbf{0}$$

# Stability is very restricted for higher orders of BDF methods

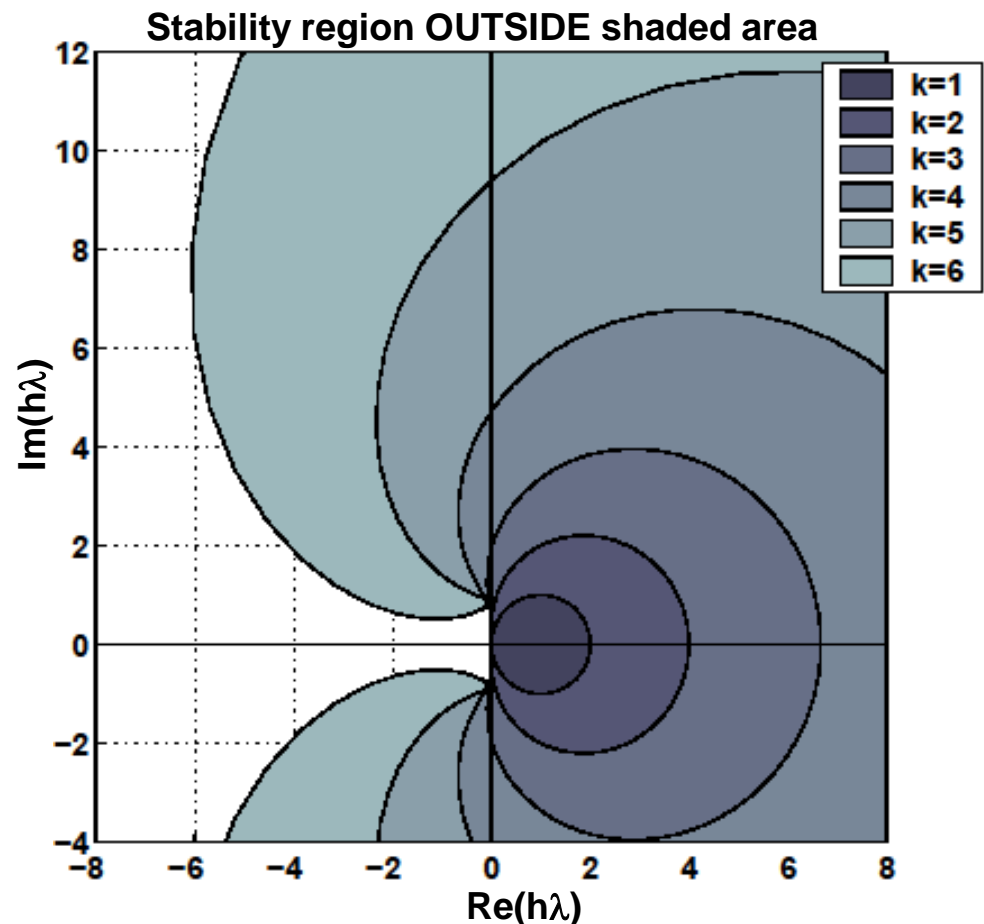
$$\mathbf{y}_n - \beta_0 \mathbf{h}_n \dot{\mathbf{y}}_n = \sum_{i=1}^k \alpha_{n,i} \mathbf{y}_{n-i}$$

Regions of instability grow with the order

CVODE and IDA allow up to order 5

CVODE includes an optional stability limit detection algorithm:

- Based on linear analysis
- Limits step if it detects a potential stability problem



# CVODE solves $\dot{\mathbf{y}} = \mathbf{f}(\mathbf{t}, \mathbf{y})$

- Variable order and variable step size methods:
  - BDF (backward differentiation formulas) for stiff systems
  - Implicit Adams for nonstiff systems
- (Stiff case) Solves time step for the system  $\dot{\mathbf{y}} = \mathbf{f}(\mathbf{t}, \mathbf{y})$ 
  - applies an explicit predictor to give  $\mathbf{y}_{n(0)}$

$$\mathbf{y}_{n(0)} = \sum_{j=1}^q \alpha_j^p \mathbf{y}_{n-j} + \Delta t \beta_1^p \dot{\mathbf{y}}_{n-1}$$

- applies an implicit corrector with  $\mathbf{y}_{n(0)}$  as the initial guess

$$\mathbf{y}_n = \sum_{j=1}^q \alpha_j \mathbf{y}_{n-j} + \Delta t \beta_0 \mathbf{f}_n(\mathbf{y}_n)$$

# Time steps and order are chosen to minimize the local truncation error

- Time steps are chosen by:
  - Estimate the error:  $E(\Delta t) = C(y_n - y_{n(0)})$ 
    - Accept step if  $\|E(\Delta t)\|_{WRMS} < 1$
    - Reject step otherwise
  - Estimate error at the next step,  $\Delta t'$ , as

$$E(\Delta t') \approx (\Delta t' / \Delta t)^{q+1} E(\Delta t)$$

- Choose next step so that  $\|E(\Delta t')\|_{WRMS} < 1$
- Choose method order by:
  - Estimate error for next higher and lower orders
  - Choose the order that gives the largest time step meeting the error condition



# Computations weighted so no component disproportionately impacts convergence

- An absolute tolerance is specified for each solution component,  $ATOL^i$
- A relative tolerance is specified for all solution components,  $RTOL$
- Norm calculations are weighted by:

$$ewt^i = \frac{1}{RTOL \cdot |y^i| + ATOL^i} \quad \|y\|_{WRMS} = \sqrt{\frac{1}{N} \sum_{i=1}^N (ewt^i \cdot y^i)^2}$$

- Bound time integration error with:

$$\|y_n - y_{n(0)}\| < \frac{1}{6}$$

The  $1/6$  factor tries to account for estimation errors

# Nonlinear system will require nonlinear solves

- Use predicted value as the initial iterate for the nonlinear solver
- Nonstiff systems: Functional iteration

$$\mathbf{y}_{n(m+1)} = \beta_0 \mathbf{h}_n \mathbf{f}(\mathbf{y}_{n(m)}) + \sum_{i=1}^q \alpha_{n,i} \mathbf{y}_{n-i}$$

- Stiff systems: Newton iteration

$$\mathbf{M}(\mathbf{y}_{n(m+1)} - \mathbf{y}_{n(m)}) = -\mathbf{G}(\mathbf{y}_{n(m)})$$

- ODE:  $\mathbf{M} \approx \mathbf{I} - \gamma \partial \mathbf{f} / \partial \mathbf{y}, \quad \gamma = \beta_0 \mathbf{h}_n$
- DAE:  $\mathbf{M} \approx \partial \mathbf{F} / \partial \mathbf{y} + \gamma \partial \mathbf{F} / \partial \dot{\mathbf{y}}, \quad \gamma = \mathbf{1} / (\beta_0 \mathbf{h}_n)$

# SUNDIALS provides many options for linear solvers

- Iterative linear solvers
  - Result in inexact Newton solver
  - Scaled preconditioned solvers: GMRES, Bi-CGStab, TFQMR
  - Only require matrix-vector products
  - Require preconditioner for the Newton matrix,  $M$
- Jacobian information (matrix or matrix-vector product) can be supplied by the user or estimated with finite difference quotients
- Two options require serial environments and some pre-defined structure to the data
  - Direct dense
  - Direct band



# An inexact Newton-Krylov method can be used to solve the implicit systems

- Krylov iterative methods find the linear system solution in a Krylov subspace:  $K(J, r) = \{r, Jr, J^2r, \dots\}$
- Only require matrix-vector products
- Difference approximations to the matrix-vector product are used,

$$J(x)v \approx \frac{F(x + \theta v) - F(x)}{\theta}$$

- Matrix entries need never be formed, and memory savings can be used for a better preconditioner

# IDA solves $F(t, y, y') = 0$

- C rewrite of DASPK [Brown, Hindmarsh, Petzold]
- Variable order / variable coefficient form of BDF
- Targets: implicit ODEs, index-1 DAEs, and Hessenberg index-2 DAEs
- Optional routine solves for consistent values of  $y_0$  and  $y_0'$ 
  - Semi-explicit index-1 DAEs, differential components known, algebraic unknown OR all of  $y_0'$  specified,  $y_0$  unknown
- Nonlinear systems solved by Newton-Krylov method
- Optional constraints:  $y^i > 0$ ,  $y^i < 0$ ,  $y^i \geq 0$ ,  $y^i \leq 0$

# KINSOL solves $F(\mathbf{u}) = 0$

- C rewrite of Fortran NKSOL (Brown and Saad)
- Inexact Newton solver: solves  $\mathbf{J} \Delta \mathbf{u}^n = -\mathbf{F}(\mathbf{u}^n)$  approximately
- Modified Newton option (with direct solves) – this freezes the Newton matrix over a number of iterations
- Krylov solver: scaled preconditioned GMRES, TFQMR, Bi-CGStab
  - Optional restarts for GMRES
  - Preconditioning on the right:  $(\mathbf{J} \mathbf{P}^{-1})(\mathbf{P}\mathbf{s}) = -\mathbf{F}$
- Direct solvers: dense and band (serial & special structure)
- Optional constraints:  $u_i > 0$ ,  $u_i < 0$ ,  $u_i \geq 0$  or  $u_i \leq 0$
- Can scale equations and/or unknowns
- Dynamic linear tolerance selection

# An inexact Newton's method is used to solve the nonlinear problem

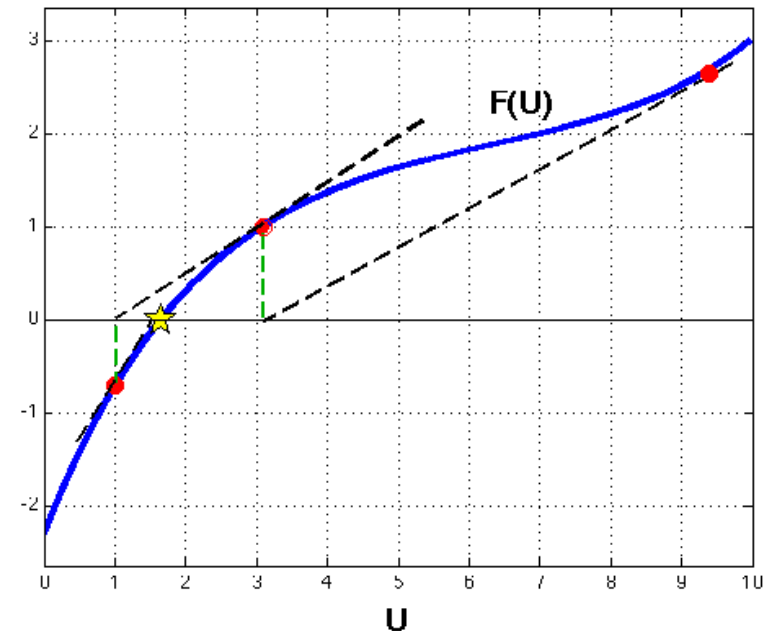
1. Starting with  $x^0$ , want  $x^*$  such that  $F(x^*) = 0$
2. Repeat for each  $k$  until  $\|\mathbf{F}(\mathbf{x}^{k+1})\| \leq \mathbf{tol}$

a. Solve (approximately)

$$\mathbf{J}(\mathbf{x}^k)\mathbf{s}^k = -\mathbf{F}(\mathbf{x}^k)$$

b. Update,  $x^{k+1} = x^k + \lambda s^k$

- $\mathbf{tol}$  may be chosen adaptively based on accuracy requirements
- $\lambda$  is a search parameter
- $\|\cdot\|$  is a weighted L-2 norm



courtesy of D. Reynolds (SMU)

# Linear stopping tolerances must be chosen to prevent “oversolves”

The linear system is solved to a given tolerance:

$$\left\| \mathbf{F}(\mathbf{x}^k) + \mathbf{J}(\mathbf{x}^k) \mathbf{s}^{k+1} \right\| \leq \eta^k \left\| \mathbf{F}(\mathbf{x}^k) \right\|$$

- Newton method assumes a linear model
  - Bad approximation far from solution, loose tol.
  - Good approximation close to solution, tight tol.
- Eisenstat and Walker (SISC 96)
  - Choice 1  $\eta^k = \frac{\left\| \mathbf{F}^k \right\| - \left\| \mathbf{F}^{k-1} - \mathbf{J}^{k-1} \mathbf{s}^{k-1} \right\|}{\left\| \mathbf{F}^{k-1} \right\|}$
  - Choice 2  $\eta^k = 0.9 \left( \frac{\left\| \mathbf{F}^{(k)} \right\|}{\left\| \mathbf{F}^{(k-1)} \right\|} \right)^2$
- ODE literature  $\eta^k = 0.05$



# Inexact methods maintain the fast rate of convergence of Newton's method

- Convergence of Newton's method is *q-quadratic* locally, for some constant C

$$\|\mathbf{x}^{k+1} - \mathbf{x}^*\| \leq C \|\mathbf{x}^k - \mathbf{x}^*\|^2$$

- Convergence of an inexact Newton method is
  - *q-linear* if  $\eta^k$  is constant in  $k$
  - *q-super-linear* if  $\lim_{k \rightarrow \infty} \eta^k = 0$
  - *q-quadratic* if for some constant C

$$\|\mathbf{F}(\mathbf{x}^k) + \mathbf{J}(\mathbf{x}^k)\mathbf{s}^{k+1}\| \leq C \|\mathbf{F}(\mathbf{x}^k)\|^2$$

- Eisenstat and Walker methods are *q-quadratic*

# Line-search globalization for Newton's method can enhance robustness

- User can select:
  - Inexact Newton
  - Inexact Newton with line search
- Line searches can provide more flexibility in the initial guess (larger time steps)
- Take,  $x^{k+1} = x^k + \lambda s^{k+1}$ , for  $\lambda$  chosen appropriately (to satisfy the Goldstein-Armijo conditions):
  - sufficient decrease in  $F$  relative to the step length
  - minimum step length relative to the initial rate of decrease
  - full Newton step when close to the solution

# Preconditioning is essential for large problems as Krylov methods can stagnate

- Preconditioner  $P$  must approximate Newton matrix, yet be reasonably efficient to evaluate and solve.
- Typical  $P$  (for time-dep. ODE problem) is  $I - \gamma\tilde{J}$ ,  $\tilde{J} \approx J$
- The user must supply two routines for treatment of  $P$ :
  - Setup: evaluate and preprocess  $P$  (infrequently)
  - Solve: solve systems  $Px=b$  (frequently)
- User can save and reuse approximation to  $J$ , as directed by the solver
- SUNDIALS offers hooks for user-supplied preconditioning
- Band and block-banded preconditioners are supplied for use with the supplied vector structure

# Sensitivity Analysis

- Sensitivity Analysis (SA) is the study of how the variation in the output of a model (**numerical** or otherwise) can be apportioned, qualitatively or **quantitatively**, to different sources of variation in inputs.
- Applications:
  - Model evaluation (most and/or least influential parameters), Model reduction, Data assimilation, Uncertainty quantification, Optimization (parameter estimation, design optimization, optimal control, ...)
- Approaches:
  - Forward sensitivity analysis
  - Adjoint sensitivity analysis

# The SUNDIALS vector module is generic

- Data vector structures can be user-supplied
- The generic NVECTOR module defines:
  - A `content` structure (void \*)
  - An `ops` structure – pointers to actual vector operations supplied by a vector definition
- Each implementation of NVECTOR defines:
  - Content structure specifying the actual vector data and any information needed to make new vectors (problem or grid data)
  - Implemented vector operations
  - Routines to clone vectors
- Note that all parallel communication resides in reduction operations: dot products, norms, mins, etc.

# SUNDIALS provides serial and parallel NVECTOR implementations

- *Use is optional*
- Vectors are laid out as an array of doubles (or floats)
- Appropriate lengths (local, global) are specified
- Operations are fast since stride is always 1
- All vector operations are provided for both serial and parallel cases
- For the parallel vector, MPI is used for global reductions
- These serve as good templates for creating a user-supplied vector structure around a user's own existing structures



# SUNDIALS provides Fortran interfaces

- CVODE, IDA, and KINSOL
- Cross-language calls go in both directions:
- Fortran user code  $\leftrightarrow$  interfaces  $\leftrightarrow$  CVODE/KINSOL/IDA
  
- Fortran main  $\rightarrow$  interfaces to solver routines
- Solver routines  $\rightarrow$  interface to user's problem-defining routine and preconditioning routines
  
- For portability, all user routines have fixed names
- Examples are provided



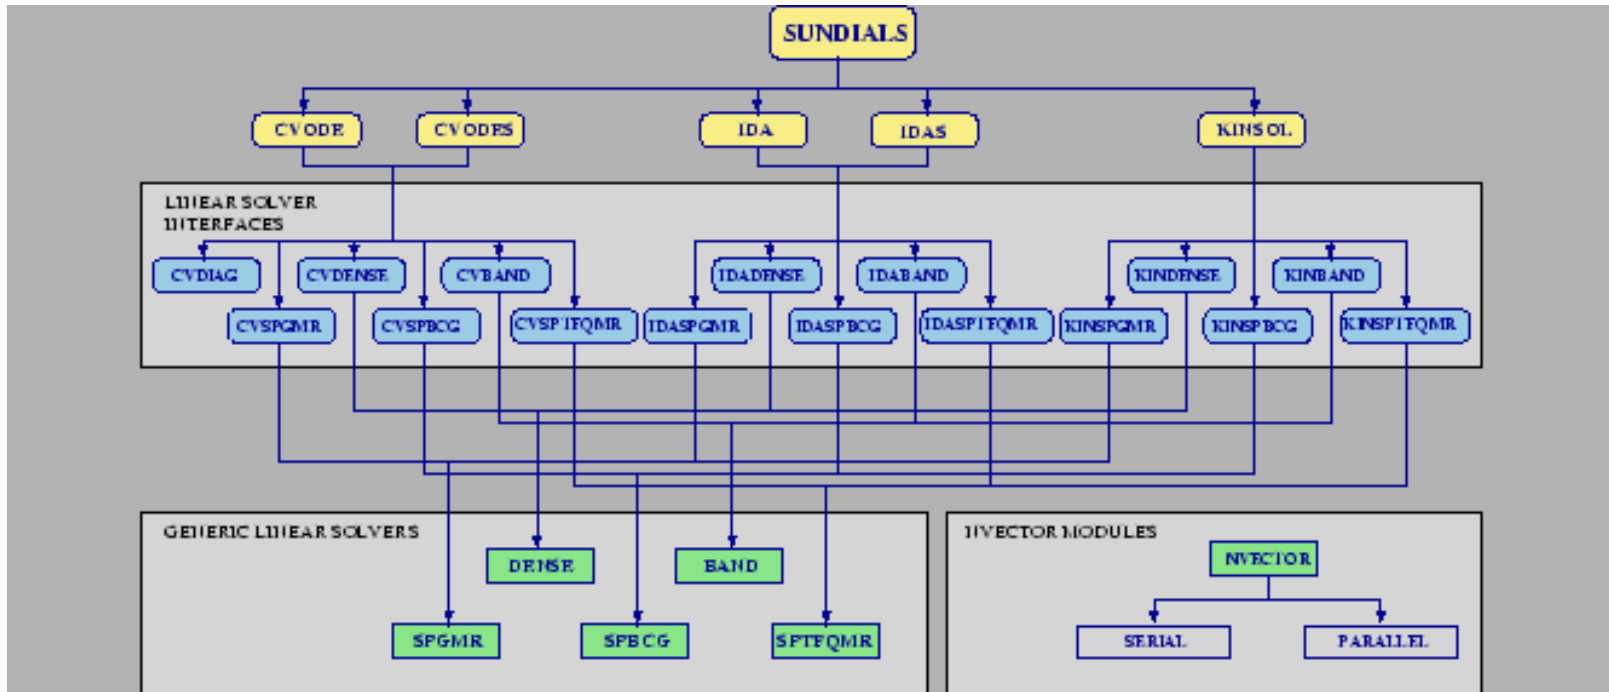
# SUNDIALS provides Matlab interfaces

- CVODES, KINSOL, and IDAS
- The core of each interface is a single MEX file which interfaces to solver-specific user-callable functions
- Guiding design philosophy: make interfaces equally familiar to both SUNDIALS and Matlab users
  - all user-provided functions are Matlab m-files
  - all user-callable functions have the same names as the corresponding C functions
  - unlike the Matlab ODE solvers, we provide the more flexible SUNDIALS approach in which the 'Solve' function only returns the solution at the next requested output time.
- Includes complete documentation (including through the Matlab help system) and several examples





# Structure of SUNDIALS



High-level diagram (note that none of the Lapack-based linear solver modules are represented.)

# SUNDIALS code usage is similar across the suite

- Have a series of Set/Get routines to set options
- For CVODE with parallel vector implementation:

```
#include "cvode.h"
#include "cvode_spgmr.h"
#include "nvector_*.h"

y = NV_New_(n,...);
cvmem = CVodeCreate(CV_BDF,CV_NEWTON);
flag = CVodeSet*(...);
flag = CVodeInit(cvmem,rhs,t0,y,...);
flag = CVSpgmr(cvmem,...);
for(tout = ...) {
    flag = CVode(cvmem, ...,y,...); }

NV_Destroy(y);
CVodeFree(&cvmem);
```

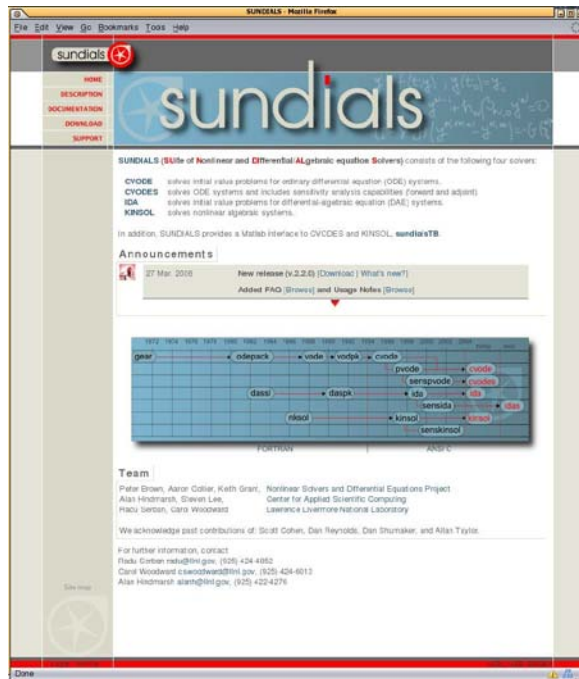
# Availability

## Open source BSD license

<https://computation.llnl.gov/casc/sundials>

## Publications

<https://computation.llnl.gov/casc/sundials/documentation/documentation.html>



## Web site:

Individual codes download  
SUNDIALS suite download  
User manuals  
User group email list

## The SUNDIALS Team:

Alan Hindmarsh, Radu Serban, and  
Carol Woodward