

# Handling stiff ordinary differential equations (ODEs) using IMEX methods in JULIA

TKP4555 Process-System Engineering Specialization Module  
TKP11 Advanced Process Simulation

Kjetil Sonerud

Department of Chemical Engineering  
Norwegian University of Science and Technology (NTNU)

December 17, 2014



# Agenda

Outline of the presentation

- 1 Introduction
- 2 What is an IMEX solver?
- 3 Time scale analysis
- 4 Numerical solvers for ODE systems
- 5 Model
- 6 Results
- 7 Discussion

# Introduction

## The role of differential equations

In general, many of the practical problems in science and engineering may be formulated as differential equations. Being able to solve these correctly and efficiently are of great importance to the chemical engineer.

In practise, many PDEs may be approximated – and all higher-order ODEs may be reduced – to a system of first-order ordinary differential equations. Thus, the solution of the latter covers a wide range of applications.

We are seldom able to solve the differential equations that are of practical interest analytically, so approximating the solution through the use of *numerical methods* are often the only viable option.

This is where the issue of *stiffness* arises. Handling this through the use of so-called *IMEX methods* is the main focus of the current work, and will be further explained.

# Introduction

## The role of differential equations

In general, many of the practical problems in science and engineering may be formulated as differential equations. Being able to solve these correctly and efficiently are of great importance to the chemical engineer.

In practise, many PDEs may be approximated – and all higher-order ODEs may be reduced – to a system of first-order ordinary differential equations. Thus, the solution of the latter covers a wide range of applications.

We are seldom able to solve the differential equations that are of practical interest analytically, so approximating the solution through the use of *numerical methods* are often the only viable option.

This is where the issue of *stiffness* arises. Handling this through the use of so-called *IMEX methods* is the main focus of the current work, and will be further explained.

# Introduction

## The role of differential equations

In general, many of the practical problems in science and engineering may be formulated as differential equations. Being able to solve these correctly and efficiently are of great importance to the chemical engineer.

In practise, many PDEs may be approximated – and all higher-order ODEs may be reduced – to a system of first-order ordinary differential equations. Thus, the solution of the latter covers a wide range of applications.

We are seldom able to solve the differential equations that are of practical interest analytically, so approximating the solution through the use of *numerical methods* are often the only viable option.

This is where the issue of *stiffness* arises. Handling this through the use of so-called *IMEX methods* is the main focus of the current work, and will be further explained.

# Introduction

## The role of differential equations

In general, many of the practical problems in science and engineering may be formulated as differential equations. Being able to solve these correctly and efficiently are of great importance to the chemical engineer.

In practise, many PDEs may be approximated – and all higher-order ODEs may be reduced – to a system of first-order ordinary differential equations. Thus, the solution of the latter covers a wide range of applications.

We are seldom able to solve the differential equations that are of practical interest analytically, so approximating the solution through the use of *numerical methods* are often the only viable option.

This is where the issue of *stiffness* arises. Handling this through the use of so-called *IMEX methods* is the main focus of the current work, and will be further explained.

# Introduction

## The role of differential equations

In general, many of the practical problems in science and engineering may be formulated as differential equations. Being able to solve these correctly and efficiently are of great importance to the chemical engineer.

In practise, many PDEs may be approximated – and all higher-order ODEs may be reduced – to a system of first-order ordinary differential equations. Thus, the solution of the latter covers a wide range of applications.

We are seldom able to solve the differential equations that are of practical interest analytically, so approximating the solution through the use of *numerical methods* are often the only viable option.

This is where the issue of *stiffness* arises. Handling this through the use of so-called *IMEX methods* is the main focus of the current work, and will be further explained.

# Introduction

## What is a stiff differential equation?

Simply put, it is (a set of) differential equations that unfold on sufficiently different time scales. The concept of stiffness is largely a practical one – no single theoretical definition applies to all problems normally considered stiff.

*“If a numerical method is forced to use, in a certain interval of integration, a step length which is excessively small in relation to the smoothness of the exact solution in that interval, then the problem is said to be stiff in that interval.”*

– (Lambert, 1991)

It is an observed fact that *implicit methods* normally outperform *explicit methods* when stiff problems arise. However, *explicit methods* are usually faster in the general case. Thus, there is a trade-off. This is where IMEX methods come in, being “the best of both worlds”



# Introduction

What is a stiff differential equation?

Simply put, it is (a set of) differential equations that unfold on sufficiently different time scales. The concept of stiffness is largely a practical one – no single theoretical definition applies to all problems normally considered stiff.

*“If a numerical method is forced to use, in a certain interval of integration, a step length which is excessively small in relation to the smoothness of the exact solution in that interval, then the problem is said to be stiff in that interval.”*

– (Lambert, 1991)

It is an observed fact that *implicit methods* normally outperform *explicit methods* when stiff problems arise. However, *explicit methods* are usually faster in the general case. Thus, there is a trade-off. This is where IMEX methods come in, being “the best of both worlds”

# Introduction

What is a stiff differential equation?

Simply put, it is (a set of) differential equations that unfold on sufficiently different time scales. The concept of stiffness is largely a practical one – no single theoretical definition applies to all problems normally considered stiff.

*“If a numerical method is forced to use, in a certain interval of integration, a step length which is excessively small in relation to the smoothness of the exact solution in that interval, then the problem is said to be stiff in that interval.”*

– (Lambert, 1991)

It is an observed fact that *implicit methods* normally outperform *explicit methods* when stiff problems arise. However, *explicit methods* are usually faster in the general case. Thus, there is a trade-off. This is where IMEX methods come in, being “the best of both worlds”

# Introduction

What is a stiff differential equation?

Simply put, it is (a set of) differential equations that unfold on sufficiently different time scales. The concept of stiffness is largely a practical one – no single theoretical definition applies to all problems normally considered stiff.

*“If a numerical method is forced to use, in a certain interval of integration, a step length which is excessively small in relation to the smoothness of the exact solution in that interval, then the problem is said to be stiff in that interval.”*

– (Lambert, 1991)

It is an observed fact that *implicit methods* normally outperform *explicit methods* when stiff problems arise. However, *explicit methods* are usually faster in the general case. Thus, there is a trade-off. This is where IMEX methods come in, being “the best of both worlds”

# Introduction

## An illustrative example

Comparing IM- and EX Euler with the analytic solution




Figure: Comparison of the analytic solution to those of EX Euler and IM Euler method for the ODE-system described in  $y'' + 101y' + 100y = 0$ . The time step is  $\Delta t = 1$  s. Note that the IM Euler solution is quite close to the analytical, but that the EX Euler method fails spectacularly.

# Introduction

## An illustrative example

Comparing IM- and EX Euler with the analytic solution




Figure: Comparison of the analytic solution to those of EX Euler and IM Euler method for the ODE-system described in  $y'' + 101y' + 100y = 0$ . The time step is  $\Delta t = 1 \times 10^{-3}$  s. Now, both EX Euler and IM Euler are very close to the actual analytic solution.

# What is an IMEX solver?

In many practical applications, large systems of ODEs that contain both stiff and non-stiff elements must be solved. Two straight-forward approaches are:

- Use an explicit method with a short time step
- Solve the whole system using an implicit scheme

The basic premise of an IMEX solver is to provide a compromise – the stiff parts of the system are solved with an implicit solver to ensure stability, while the non-stiff parts are solved with an explicit solver to reduce computational load.

$$\dot{\underline{y}}(t) = F(t, \underline{y}) + G(t, \underline{y}) \quad (1)$$

Here  $F(t, \underline{y})$  represent the non-stiff part and  $G(t, \underline{y})$  represent the stiff part. In general, it is assumed that *a priori* knowledge of the system is exploited to distinguish the stiff and non-stiff part so that the corresponding solver may be employed.

# What is an IMEX solver?

In many practical applications, large systems of ODEs that contain both stiff and non-stiff elements must be solved. Two straight-forward approaches are:

- Use an explicit method with a short time step
- Solve the whole system using an implicit scheme

The basic premise of an IMEX solver is to provide a compromise – the stiff parts of the system are solved with an implicit solver to ensure stability, while the non-stiff parts are solved with an explicit solver to reduce computational load.

$$\dot{\underline{y}}(t) = F(t, \underline{y}) + G(t, \underline{y}) \quad (1)$$

Here  $F(t, \underline{y})$  represent the non-stiff part and  $G(t, \underline{y})$  represent the stiff part. In general, it is assumed that *a priori* knowledge of the system is exploited to distinguish the stiff and non-stiff part so that the corresponding solver may be employed.

# What is an IMEX solver?

In many practical applications, large systems of ODEs that contain both stiff and non-stiff elements must be solved. Two straight-forward approaches are:

- Use an explicit method with a short time step
- Solve the whole system using an implicit scheme

The basic premise of an IMEX solver is to provide a compromise – the stiff parts of the system are solved with an implicit solver to ensure stability, while the non-stiff parts are solved with an explicit solver to reduce computational load.

$$\dot{\underline{y}}(t) = F(t, \underline{y}) + G(t, \underline{y}) \quad (1)$$

Here  $F(t, \underline{y})$  represent the non-stiff part and  $G(t, \underline{y})$  represent the stiff part. In general, it is assumed that *a priori* knowledge of the system is exploited to distinguish the stiff and non-stiff part so that the corresponding solver may be employed.



# Time scale analysis

Alternative strategies for solving stiff equations




Figure: Illustration of time scales: the dynamic time scale is flanked to the left by the small time scales, which may be considered like *events*. To the right, the dynamic time scale is flanked by the large time scales, which may be considered *constant*.

# Numerical solvers for ODE systems

## Solving the general IVP

A general *initial value problem* (IVP) may be described as

$$\frac{d\underline{y}}{dt} = \dot{\underline{y}}(t) = f(\underline{y}(t); \Theta), \quad \underline{y}(t_0) = \underline{y}_0 \quad (2)$$

The exact solution may be found by integration of the right-hand side from  $t_0$  to  $t$ . In general, this is not possible analytically.

Thus, a discrete approach is taken where the integral is approximated over a short interval from  $t_k$  to  $t_k + \Delta_t$  by

$$\underline{y}_{k+1} - \underline{y}_k = (\Delta t) \cdot \underbrace{F[\underline{y}_k, \underline{y}_{k+1}, f(\underline{y}; \Theta)]}_{\text{Defines the numerical method}} \quad (3)$$

Defines the numerical method



# Numerical solvers for ODE systems

## Solving the general IVP

A general *initial value problem* (IVP) may be described as

$$\frac{d\underline{\mathbf{y}}}{dt} = \underline{\dot{\mathbf{y}}}(t) = f(\underline{\mathbf{y}}(t); \Theta), \quad \underline{\mathbf{y}}(t_0) = \underline{\mathbf{y}}_0 \quad (2)$$

The exact solution may be found by integration of the right-hand side from  $t_0$  to  $t$ . In general, this is not possible analytically.

Thus, a discrete approach is taken where the integral is approximated over a short interval from  $t_k$  to  $t_k + \Delta_t$  by

$$\underline{\mathbf{y}}_{k+1} - \underline{\mathbf{y}}_k = (\Delta t) \cdot \underbrace{F[\underline{\mathbf{y}}_k, \underline{\mathbf{y}}_{k+1}, f(\underline{\mathbf{y}}; \Theta)]}_{\text{Defines the numerical method}} \quad (3)$$



# Numerical solvers for ODE systems

## Solving the general IVP

A general *initial value problem* (IVP) may be described as

$$\frac{d\underline{\mathbf{y}}}{dt} = \dot{\underline{\mathbf{y}}}(t) = f(\underline{\mathbf{y}}(t); \Theta), \quad \underline{\mathbf{y}}(t_0) = \underline{\mathbf{y}}_0 \quad (2)$$

The exact solution may be found by integration of the right-hand side from  $t_0$  to  $t$ . In general, this is not possible analytically.

Thus, a discrete approach is taken where the integral is approximated over a short interval from  $t_k$  to  $t_k + \Delta_t$  by

$$\underline{\mathbf{y}}_{k+1} - \underline{\mathbf{y}}_k = (\Delta t) \cdot \underbrace{F[\underline{\mathbf{y}}_k, \underline{\mathbf{y}}_{k+1}, f(\underline{\mathbf{y}}; \Theta)]}_{\text{Defines the numerical method}} \quad (3)$$

Defines the numerical method



# Numerical solvers for ODE systems

Solving the general IVP

A general *initial value problem* (IVP) may be described as

$$\frac{d\mathbf{y}}{dt} = \dot{\mathbf{y}}(t) = f(\mathbf{y}(t); \Theta), \quad \mathbf{y}(t_0) = \mathbf{y}_0 \quad (2)$$

The exact solution may be found by integration of the right-hand side from  $t_0$  to  $t$ . In general, this is not possible analytically.

Thus, a discrete approach is taken where the integral is approximated over a short interval from  $t_k$  to  $t_k + \Delta_t$  by

$$\mathbf{y}_{k+1} - \mathbf{y}_k = (\Delta t) \cdot \underbrace{F[\mathbf{y}_k, \mathbf{y}_{k+1}, f(\mathbf{y}; \Theta)]}_{\text{Defines the numerical method}} \quad (3)$$

Defines the numerical method



# Numerical solvers for ODE systems

EX and IM Euler schemes

The *explicit Euler scheme* (EX Euler, also known as “forward Euler”) may be found as

$$\underline{y}_{k+1} = \underline{y}_k + (\Delta t)f(\underline{y}_k; \Theta) \quad (4)$$

The *implicit Euler scheme* (IM Euler, also known as “backward Euler”) may be found as

$$\underline{y}_{k+1} = \underline{y}_k + (\Delta t)f(\underline{y}_{k+1}; \Theta) \quad (5)$$

From the equation, it is evident that a (possibly non-linear) system of equations must be solved in each time step to find  $\underline{y}_{k+1}$ . In the current implementation, the Newton-Raphson method is used.

# Numerical solvers for ODE systems

EX and IM Euler schemes

The *explicit Euler scheme* (EX Euler, also known as “forward Euler”) may be found as

$$\underline{y}_{k+1} = \underline{y}_k + (\Delta t)f(\underline{y}_k; \Theta) \quad (4)$$

The *implicit Euler scheme* (IM Euler, also known as “backward Euler”) may be found as

$$\underline{y}_{k+1} = \underline{y}_k + (\Delta t)f(\underline{y}_{k+1}; \Theta) \quad (5)$$

From the equation, it is evident that a (possibly non-linear) system of equations must be solved in each time step to find  $\underline{y}_{k+1}$ . In the current implementation, the Newton-Raphson method is used.

# Numerical solvers for ODE systems

EX and IM Euler schemes

The *explicit Euler scheme* (EX Euler, also known as “forward Euler”) may be found as

$$\underline{\mathbf{y}}_{k+1} = \underline{\mathbf{y}}_k + (\Delta t)f(\underline{\mathbf{y}}_k; \Theta) \quad (4)$$

The *implicit Euler scheme* (IM Euler, also known as “backward Euler”) may be found as

$$\underline{\mathbf{y}}_{k+1} = \underline{\mathbf{y}}_k + (\Delta t)f(\underline{\mathbf{y}}_{k+1}; \Theta) \quad (5)$$

From the equation, it is evident that a (possibly non-linear) system of equations must be solved in each time step to find  $\underline{\mathbf{y}}_{k+1}$ . In the current implementation, the Newton-Raphson method is used.



# Numerical solvers for ODE systems

## The IMEX Euler scheme

With the assumption that the system of ODEs are readily split into an implicit part and an explicit part, the IMEX Euler scheme follows

$$\begin{bmatrix} \underline{y}_{k+1}^{\text{IM}} \\ \underline{y}_{k+1}^{\text{EX}} \end{bmatrix} = \begin{bmatrix} \underline{y}_k^{\text{IM}} \\ \underline{y}_k^{\text{EX}} \end{bmatrix} + \Delta t \begin{bmatrix} f(\underline{y}_{k+1}^{\text{IM}}, \underline{y}_{k+1}^{\text{EX}}; \Theta) \\ f(\underline{y}_k^{\text{EX}}; \Theta) \end{bmatrix} \quad (6)$$

Note that to solve for  $\underline{y}_{k+1}^{\text{IM}}$ , it is necessary to first solve for  $\underline{y}_{k+1}^{\text{EX}}$ , as the latter is used in the solution of the former. Thus, in the Newton-Raphson iteration to solve for  $\underline{y}_{k+1}^{\text{IM}}$ , it is assumed that  $\underline{y}_{k+1}^{\text{EX}}$  is constant during the course of the iterations.

# Numerical solvers for ODE systems

## The IMEX Euler scheme

With the assumption that the system of ODEs are readily split into an implicit part and an explicit part, the IMEX Euler scheme follows

$$\begin{bmatrix} \underline{y}_{k+1}^{\text{IM}} \\ \underline{y}_{k+1}^{\text{EX}} \end{bmatrix} = \begin{bmatrix} \underline{y}_k^{\text{IM}} \\ \underline{y}_k^{\text{EX}} \end{bmatrix} + \Delta t \begin{bmatrix} f(\underline{y}_{k+1}^{\text{IM}}, \underline{y}_{k+1}^{\text{EX}}; \Theta) \\ f(\underline{y}_k^{\text{EX}}; \Theta) \end{bmatrix} \quad (6)$$

Note that to solve for  $\underline{y}_{k+1}^{\text{IM}}$ , it is necessary to first solve for  $\underline{y}_{k+1}^{\text{EX}}$ , as the latter is used in the solution of the former. Thus, in the Newton-Raphson iteration to solve for  $\underline{y}_{k+1}^{\text{IM}}$ , it is assumed that  $\underline{y}_{k+1}^{\text{EX}}$  is constant during the course of the iterations.

# Numerical solvers for ODE systems

The IMEX Euler scheme

With the assumption that the system of ODEs are readily split into an implicit part and an explicit part, the IMEX Euler scheme follows

$$\begin{bmatrix} \underline{\mathbf{y}}_{k+1}^{\text{IM}} \\ \underline{\mathbf{y}}_{k+1}^{\text{EX}} \end{bmatrix} = \begin{bmatrix} \underline{\mathbf{y}}_k^{\text{IM}} \\ \underline{\mathbf{y}}_k^{\text{EX}} \end{bmatrix} + \Delta t \begin{bmatrix} f(\underline{\mathbf{y}}_{k+1}^{\text{IM}}, \underline{\mathbf{y}}_{k+1}^{\text{EX}}; \Theta) \\ f(\underline{\mathbf{y}}_k^{\text{EX}}; \Theta) \end{bmatrix} \quad (6)$$

Note that to solve for  $\underline{\mathbf{y}}_{k+1}^{\text{IM}}$ , it is necessary to first solve for  $\underline{\mathbf{y}}_{k+1}^{\text{EX}}$ , as the latter is used in the solution of the former. Thus, in the Newton-Raphson iteration to solve for  $\underline{\mathbf{y}}_{k+1}^{\text{IM}}$ , it is assumed that  $\underline{\mathbf{y}}_{k+1}^{\text{EX}}$  is constant during the course of the iterations.

# Model

## The original Robertson kinetic problem

*“When the equations represent the behaviour of a system containing a number of fast and slow reactions, a forward integration of these equations becomes difficult.”*  
– (Robertson, 1966)

The original problem as stated by Robertson is one of an autocatalytic reaction




Note that the reaction rate constants differ with several orders of magnitude, which we might suspect – correctly, as it turns out – makes the system stiff.

# Model

The original Robertson kinetic problem

*“When the equations represent the behaviour of a system containing a number of fast and slow reactions, a forward integration of these equations becomes difficult.”*  
– (Robertson, 1966)

The original problem as stated by Robertson is one of an autocatalytic reaction



Note that the reaction rate constants differ with several orders of magnitude, which we might suspect – correctly, as it turns out – makes the system stiff.

# Model

The original Robertson kinetic problem

*“When the equations represent the behaviour of a system containing a number of fast and slow reactions, a forward integration of these equations becomes difficult.”*  
– (Robertson, 1966)

The original problem as stated by Robertson is one of an autocatalytic reaction



Note that the reaction rate constants differ with several orders of magnitude, which we might suspect – correctly, as it turns out – makes the system stiff.

# Model

## The expanded Robertson kinetic problem


An expanded version of the original problem is proposed in order to promote solution by the IMEX method central to this work. The idea behind the reaction scheme shown is to expand the non-stiff parts of the system, thus promoting the effectiveness of the IMEX solver.



# Model

## The expanded Robertson kinetic problem

An expanded version of the original problem is proposed in order to promote solution by the IMEX method central to this work. The idea behind the reaction scheme shown is to expand the non-stiff parts of the system, thus promoting the effectiveness of the IMEX solver.





# Results

## Original Robertson problem




Figure: Solving the Robertson kinetics problem – original version – with IMEX Euler. The step size in this case is  $\Delta t = 1$  s.

# Results

## Original Robertson problem

Table: Key performance data using different numerical methods for solving the original Robertson kinetics problem

Case	$\Delta t$	Number of time steps	CPU-time	Total number of N-R iterations
EX Euler (0 s to 40 s)	$6 \times 10^{-4}$ s	66667	0.161 s	–
IM Euler (0 s to 40 s)	1 s	40	0.0239 s	131
IMEX Euler (0 s to 40 s)	1 s	40	0.122 s	131
EX Euler (0 s to 1000 s)	$3 \times 10^{-4}$ s	3333333	7.789 s	–
IM Euler (0 s to 1000 s)	1 s	1000	0.0612 s	2138
IMEX Euler (0 s to 1000 s)	1 s	1000	0.0804 s	2123

Note that the IM method perform consistently better then the IMEX method, but that both these perform better than the EX method. This clearly evident for the 0 s to 1000 s simulation.

# Results

## Original Robertson problem




Figure: Solving the Robertson kinetics problem – original version – with EX Euler. The emerging stability problem when calculating the B-trajectory are evident.

# Results

## Expanded Robertson problem

Solving the expanded Robertson problem with an IMEX Euler solver




Figure: Solving the Robertson kinetics problem – expanded version – with an IMEX Euler for the time interval 0s to 600s. The step size in this case is  $\Delta t = 1$  s.


# Results

## Expanded Robertson problem

Table: Key performance data using different numerical methods for solving the expanded Robertson kinetics problem

Case	$\Delta t$	Number of time steps	CPU-time	Total number of N-R iterations
EX Euler (0 s to 600 s)	$1 \times 10^{-3} \text{ s}$	600000	2.160 s	–
IM Euler (0 s to 600 s)	1 s	600	0.0533 s	1471
IMEX Euler (0 s to 600 s)	1 s	600	0.143 s	2506

Number of N-R iterations at each time step



# Discussion

From the results, it seems evident that the current IMEX Euler implementation is less effective than the pure IM Euler implementation on the current model system. This may be due to several factors:

- Sub-optimal implementation of the IMEX solver
- The system is not sufficiently large (more precisely, the non-stiff part is relatively too small) to fully exploit the IMEX scheme
- The added benefit of solving the whole system implicitly in JULIA outweighs the added cost of solving a larger system of equations

Thus, further work should be done to investigate

- Better IMEX implementations, possibly with non-Euler methods
- Larger systems of ODEs, especially ones with a small stiff and a large non-stiff part

# Discussion

From the results, it seems evident that the current IMEX Euler implementation is less effective than the pure IM Euler implementation on the current model system. This may be due to several factors:

- Sub-optimal implementation of the IMEX solver
- The system is not sufficiently large (more precisely, the non-stiff part is relatively too small) to fully exploit the IMEX scheme
- The added benefit of solving the whole system implicitly in JULIA outweighs the added cost of solving a larger system of equations

Thus, further work should be done to investigate

- Better IMEX implementations, possibly with non-Euler methods
- Larger systems of ODEs, especially ones with a small stiff and a large non-stiff part

# Discussion

From the results, it seems evident that the current IMEX Euler implementation is less effective than the pure IM Euler implementation on the current model system. This may be due to several factors:

- Sub-optimal implementation of the IMEX solver
- The system is not sufficiently large (more precisely, the non-stiff part is relatively too small) to fully exploit the IMEX scheme
- The added benefit of solving the whole system implicitly in JULIA outweighs the added cost of solving a larger system of equations

Thus, further work should be done to investigate

- Better IMEX implementations, possibly with non-Euler methods
- Larger systems of ODEs, especially ones with a small stiff and a large non-stiff part