



# For-loop equivalents and massive parallelization in CasADI

*Citius, Altius, Fortius*

Joris Gillis  
(MECO group, KU Leuven)




# Outline

- Key ingredients of CasADI
- New concepts: Map, MapAccum
- Massive parallelization
- Software architecture



# Core idea

- Computer-aided formulation of optimization problems



Obtain gradients, Jacobians, Hessians automatically

# Basic example

- Find initial condition to achieve a goal

$$\dot{x} = ax + \cos x$$

$$\text{minimize } x_N^2$$

$$x_0$$

```
x0 = optivar()
a = optivar()
dt = 1
```

```
x1 = x0 + dt*(a*x0 + cos(x0))
x2 = x1 + dt*(a*x1 + cos(x1))
x3 = x2 + dt*(a*x2 + cos(x2))
```

...

```
xN = ...
```

```
optimize(xN**2)
```

# Basic example

- Find initial condition to achieve a goal

$$\dot{x} = ax + \cos x$$

minimize  $x_N^2$

$x_0$

```
x0 = optivar()
a = optivar()
dt = 1
```

```
x = x0
for i in range(N):
    x = x + dt*(a*x + cos(x))
```

$xN = x$

```
optimize(xN**2)
```

# Key ingredient: efficient symbolics

Matlab symbolic toolbox (mupad)

```
syms('x0')
syms('a')
dt = 1;

x = x0;

for i=1:N
    x = x+dt*(a*x+cos(x));
end
```

# Key ingredient: efficient symbolics

Matlab: yalmip

```
x0 = sdpvar(1,1)
a = sdpvar(1,1)
dt = 1;

x = x0;

for i=1:N
    x = x+dt*(a*x+cos(x));
end
```

# Key ingredient: efficient symbolics

Python: sympy

```
x0 = symbols('x0')
a = symbols('a')
dt = 1

x = x0

for i in range(N):
    x = x+dt*(a*x+cos(x))
```

# Key ingredient: efficient symbolics


Python: casadi

```
x0 = MX.sym('x0')
a = MX.sym('a')
dt = 1

x = x0

for i in range(N):
    x = x+dt*(a*x+cos(x))
```

# Key ingredient: efficient symbolics



```
x0 = MX.sym('x0')
a = MX.sym('a')
dt = 1

x = x0

for i in range(N):
    x = x+dt*(a*x+cos(x))
```

# Key ingredient: efficient symbolics

$$x_0 + dt * (a * x_0 + \cos(x_0))$$
$$x_2 = x_1 + dt * (a * x_1 + \cos(x_1))$$


```
x0 = optivar()  
a   = optivar()  
dt  = 1
```

```
x1 = x0 + dt*(a*x0 + cos(x0))  
x2 = x1 + dt*(a*x1 + cos(x1))  
x3 = x2 + dt*(a*x2 + cos(x2))  
...  
xN = ...
```

$$x_2 = (x_0 + dt * (a * x_0 + \cos(x_0))) + dt * (a * (x_0 + dt * (a * x_0 + \cos(x_0)))) + \cos(x_0 + dt * (a * x_0 + \cos(x_0)))$$

Expression length  $\sim 3^N$

# Key ingredient: efficient symbolics



graph representation

```
x0 = MX.sym('x0')
a = MX.sym('a')
dt = 1

x = x0

for i in range(N):
    x = x+dt*(a*x+cos(x))
```

#nodes:  $O(N)$


# Key ingredient: algorithmic differentiation

Cost of a (forward/reverse) derivative of:  
small multiple of original

$$\text{minimize}_{x_0} x_N^2$$

$$\frac{d x_N}{d x_0}$$

# Key ingredient: algorithmic differentiation



```
x0 = MX.sym('x0')  
a = MX.sym('a')  
dt = 1  
  
x = x0  
  
for i in range(N):  
    x = x+dt*(a*x+cos(x))  
  
jacobian(x, x0)
```

N


$$\frac{dx_N}{dx_0}$$

# Key ingredient: matrix-valued graphs


$$\dot{x} = Ax + \cos x$$
$$x \in \mathbb{R}^n$$

```
x = MX.sym('x', 2, 2)
A = MX.sym('A', 2)
mul(A, x)
```

```
x = SX.sym('x', 2, 2)
A = SX.sym('A', 2, 2)
mul(A, x)
```



# Key ingredient: function embedding




```
x0 = MX.sym('x0')
a = MX.sym('a')
dt = 1

x = x0


for i in range(N):
    x = x+dt*(a*x+cos(x))
```

# Key ingredient: function embedding



```
x0 = MX.sym('x0')  
a = MX.sym('a')  
dt = 1  
  
x = x0  
  
for i in range(N):  
    x = F(x,a)
```

# Key ingredient: function embedding




```
x = MX.sym('x')
a = MX.sym('a')
dt = 1
```

```
e = x+dt*(a*x+cos(x))
F = Function("F", [x,a], [e])
```

```
x0 = MX.sym('x0')
a = MX.sym('a')
dt = 1
```

```
x = x0
```

```
for i in range(N):
    x = F(x,a)
```



# Key ingredient: function embedding



rootfinder  
linsol  
integrator

...


# Outline

- Key ingredients of CasADi
  - Efficient symbolics
  - Algorithmic differentiation
  - Matrix-valued graphs
  - Function embedding
- New concepts: Map, MapAccum
- Massive parallelization
- Software architecture



# Towards large scale optimal control

$$x_{k+1} = F(x_k, u_k)$$



MX graph (optimized for memory footprint)

SX graph (optimized for speed) ←

$$\text{minimize } x_N^2$$


$$x_0, u_0, u_1 \dots u_{N-1}$$

```
x0 = MX.sym('x0')
U = MX.sym('u', 1, N)
dt = 1


w = horzsplit(U)

x = x0

for i in range(N):
    x = F(x, w[i])
```




# Towards large scale optimal control





# Towards large scale optimal control

Not fast enough!



$N = 10\ 000, \ 100\ 000, \dots$

# New node: MapAccum




# Practical example: sys id


$$\underset{x_0, k, c, d}{\text{minimize}} \quad \|x_i - \bar{x}_i\|_2^2$$

$$\dot{x} = v \\ \dot{v} = u - kx - cv - dx^3$$

$$x_{k+1} = F(x_k, u_k, p)$$



```
Function("F", [x, u, vertcat([k, c, d])], [...])
```



# Practical example: sys id

$$\underset{x_0, k, c, d}{\text{minimize}} \quad \|x_i - \bar{x}_i\|_2^2$$

$$\begin{aligned}\dot{x} &= v \\ \dot{v} &= u - kx - cv - dx^3\end{aligned}$$

Known signals


```
R = F.mapaccum("R", N)  
X_symbolic = R(x0, u_data, repmat(params, 1, N))  
e = y_data-X_symbolic[0, :].T;
```

Measured data


# Practical example: sys id

```
for i in range(N):  
    X = F(X, u_data[:, i], params)
```

```
R = F.mapaccum("R", N)
```




# Construction time $\leftrightarrow$ evaluation time




# Construction time $\leftrightarrow$ evaluation time

So MapAccum is useless to speed up online computations?



# Code-generation




# Code-generation




First you need to run the compiler...

Compilation time [s]




Compilation memory [B]



# Outline



- Key ingredients of CasADi
  - Efficient symbolics
  - Algorithmic differentiation
  - Matrix-valued graphs
  - Function embedding
- New concepts: MapAccum
  - 2 orders of magnitude faster construction
  - compile 2 orders of magnitude larger problems
- Software architecture




# MapAccum ~ single shooting

minimize  $x_N^2$

$x_0, u_0, u_1 \dots u_{N-1}$



# Map ~ multiple shooting



# Map ~ multiple shooting

minimize  $x_N^2$

$x_0, x_1, \dots, x_N, u_0, u_1, \dots, u_{N-1}$

s.t.

$$\begin{aligned}x_1 - F(x_0; u_0) &= 0 \\x_2 - F(x_1; u_1) &= 0 \\\dots \\x_N - F(x_{N-1}; u_{N-1}) &= 0\end{aligned}$$

F is a function that:

- reads a few inputs
- computes a lot (e.g. time integration)
- writes a few outputs

Ideal for parallelism:

- openmp support builtin
- opencl → GPU (proof-of-concept)

```
R = F.map("R", "serial", N)
```

```
X_n = R(X, u_data, repmat(params, 1, N))
```


```
gaps = X[:, 1:] - X_n[:, :-1]
```

# Outline

- Key ingredients of CasADi
  - Efficient symbolics
  - Algorithmic differentiation
  - Matrix-valued graphs
  - Function embedding
- New concepts: Map, MapAccum
  - Speedups
- Massive parallelization
- Software architecture



# Operation on an image



# Operation on an image

$$\sum_{i,j} f(D_{i,j}, X)$$

```
X = MX.sym("x", 2)
D = MX.sym("D", 1024, 768)

R = F.map("mymap", "serial", 1024*768)

Dflat = vec(D).T

terms = R(Dflat, X)

e = sum2(terms)
```

# Operation on an image

$$\sum_{i,j} f(D_{i,j}, X)$$

```
X = MX.sym("x", 2)
D = MX.sym("D", 1024, 768)

R = F.map("mymap", "openmp", 1024*768)

Dflat = vec(D).T

terms = R(Dflat, X)

e = sum2(terms)
```

# Operation on an image

$$\sum_{i,j} f(D_{i,j}, X)$$

```
X = MX.sym("x", 2)
D = MX.sym("D", 1024, 768)

R = F.map("mymap", "opencl", 1024*768)

Dflat = vec(D).T

terms = R(Dflat, X)

e = sum2(terms)
```

# Operation on an image

$$\sum_{i,j} f(D_{i,j}, X)$$

```
X = MX.sym("x", 2)
D = MX.sym("D", 1024, 768)

R = F.map("mymap", "opencl", 1024*768)

Dflat = vec(D).T

terms = R(Dflat, X)

e = sum2(terms)
```

# What is OpenCL?

- C library
- Compiler for computation kernels
- Low-level memory concepts
- Write once, run on
  - CPU
  - GPU
  - FPGA

# Operation on an image

$$\sum_{i,j} f(D_{i,j}, X)$$

```
X = MX.sym("x", 2)
D = MX.sym("D", 1024, 768)

R = F.map("mymap", "opencl")
Dflat = vec(D).T

terms = R(Dflat, X)

e = sum2(terms)
```

## Init:

Allocate buffer for x on GPU  
Allocate buffer for D on GPU  
Allocate buffer for result on GPU  
Compile GPU kernel

## Eval:

Send x  
Send D  
Compute kernel  
Retrieve sol

# Kernelsum node

$$\sum_{i,j} f(D_{i,j}, X)$$

## **Init:**

Allocate buffer for x on GPU  
Allocate buffer for D on GPU  
Allocate buffer for result on GPU

## **Send D**

Compile GPU kernel

## **Eval:**

Send x  
Compute kernel  
**Sum result**  
Retrieve sum

# Kernelsum node

For Flanders Make:

100..1000 x speedup on a Titan Black GPU (2000 cores)



# Outline

- Key ingredients of CasADi
  - Efficient symbolics
  - Algorithmic differentiation
  - Matrix-valued graphs
  - Function embedding
- New concepts: Map, MapAccum
- Massive parallelization
- Software architecture



# Software architecture: interfaces

Python interface source

C++ source

casadi.cpp

casadi.hpp

SWIG interface  
header

casadi.i

casadiPYTHON\_WRAP.cxx

\_casadi.so


Matlab interface source

casadiMATLAB\_WRAP.cxx



casadiMATLAB\_wrap.mexa64

libcasadi.so


# Software architecture: before plugins



# Software architecture: before plugins



# Software architecture: before plugins



Mr. Cheapskate

```
from casadi import *
nlpSolver("solver","ipopt",nlp)
```




Mrs. Spender

```
from casadi import *
nlpSolver("solver","snopt",nlp)
```



Error: solver snopt not compiled

# Software architecture: plugins



```
from casadi import *
nlpsol("ipopt",nlp)
```

Dynamically load (dlopen)  
libcasadi\_nlpsolver\_ipopt.so



```
from casadi import *
nlpsol("snopt",nlp)
```

Dynamically load (dlopen)  
libcasadi\_nlpsolver\_snopt.so

libipopt.so (LGPL)

libsnopt.so (commercial)

# Software architecture: version control



Github

# Software architecture: version control

Activate more clang for unittests

develop + release-2.4.0-rc2 + test-loader

jgillis authored 15 days ago

1 parent 002bbf1 commit 798211fe0d465b7eb2aaee6e5d0800323af5b5f1

Showing 1 changed file with 4 additions and 4 deletions.

Unified Split View

8 .travis.yml

```
@@ -58,7 +58,7 @@ matrix:  
 58     script:  
 59         - mkdir build  
 60         - pushd build  
- 61         - cmake -DWITH_ECOS=ON -DWITH_MOSEK=ON -DWITH_PYTHON=ON -DWITH_WORHP=ON -DWITH_SLICOT=ON  
- DWITH_OOQP=ON -DWITH_PROFILING=ON -DWITH_DOC=ON -DWITH_EXAMPLES=ON -DWITH_COVERAGE=ON  
- DWITH_EXTRA_WARNINGS=ON -DWITH_PYTHON=ON -DWITH_JSON=ON ..  
 62         - make -j2  
  
 58     script:  
 59         - mkdir build  
 60         - pushd build  
+ 61         - cmake -DWITH_CLANG=ON -DWITH_ECOS=ON -DWITH_MOSEK=ON -DWITH_PYTHON=ON -DWITH_WORHP=ON  
- DWITH_SLICOT=ON -DWITH_OOQP=ON -DWITH_PROFILING=ON -DWITH_DOC=ON -DWITH_EXAMPLES=ON  
- DWITH_COVERAGE=ON -DWITH_EXTRA_WARNINGS=ON -DWITH_PYTHON=ON -DWITH_JSON=ON ..  
 62         - make -j2
```

I added a flag here

# Software architecture: continuous-integration

Every change (commit) is untested by travis-ci

casadi / casadi build error

Current Branches Build History Pull Requests > Build #1560

develop Some extra .gitignores

Joris Gillis authored and committed

# 1560 passed Commit 5dcd444 Compare 8d94536..5dcd444 ran for 3 hrs 52 min 9 sec 7 days ago

Build Jobs

#	Compiler	TESTMODE	Duration
1560.1	gcc	full_valgrind	34 min 9 sec
1560.2	gcc	full_valgrind	42 min 18 sec
1560.3	gcc	full_remainder	19 min 17 sec
1560.4	gcc	full_slow	33 min 58 sec
1560.5	gcc	docs	25 min 21 sec

cpu-time

Memory checks

Documentation

Green = good

# Software architecture: continuous-integration

[PROJECTS](#)[ENVIRONMENTS](#)[DOCS](#)[SUPPORT](#)[+ NEW PROJECT](#)

## binaries

automatic test commit c4a8d10

7 hours ago by [casaditestbot](#)

tests-windows ↗ 7418b476

1.0.404

7 hours ago in 5 min 18 sec

## casadi

Merge pull request #1546 from b4be1/patch-1

8 hours ago by [Joel Andersson](#)

develop ↗ c4a8d10a

1.0.1418

8 hours ago in 6 min 37 sec

KU LEUVEN

# Software architecture: continuous-integration

Both appveyor and travis are free for open-source projects

Encryption support → can test commercial plugins/interfaces

e.g. Matlab

# Software architecture: binaries

Linux buildslaves

The screenshot shows a GitHub commit page for the CasADi repository. The URL is [Home / CasADi / commits / 2e60be1](#). The page displays a list of files and folders:

Name	Modified	Size	Downloads	Week
Parent folder				
<a href="#">linux</a>	2015-09-04		18	
<a href="#">windows</a>	2015-09-04		30	
<a href="#">osx</a>	2015-09-04		7	
<a href="#">casadi-docs-2e60be1.zip</a>	2015-09-05	59.2 MB	6	<a href="#">i</a>
<a href="#">casadi-example_pack-2e60be1.zip</a>	2015-09-05	603.7 kB	7	<a href="#">i</a>
<a href="#">README.md</a>	2015-09-05	81 Bytes	6	<a href="#">i</a>
<b>Totals: 6 Items</b>		<b>59.8 MB</b>	<b>19</b>	

A callout bubble highlights the 'linux' folder with the text 'g++' and 'x86\_64-w64-mingw32-g++'. An arrow points from the 'linux' folder to the 'Parent folder' header.

# Software architecture: binaries (python)

Binary = zip folder

e.g. casadi-py27-np1.9.1-v3.0.0.tar.gz

/home/me/software

- casadi-py27-np1.9.1-v2.4.0
  - casadi
  - include
  - lib
- casadi-py27-np1.9.1-v3.0.0
  - casadi
  - include
  - lib

```
import sys
sys.path.append("/home/me/software/casadi-py27-np1.9.1-v3.0.0")

import casadi
```

# Software architecture: binaries (matlab)

Binary = zip folder

e.g. casadi-matlabR2014b-v2.4.0.zip

/home/me/software

- casadi-matlabR2014b-v2.4.0
  - +casadi
  - include
  - lib
- casadi-matlabR2014b-v3.0.0
  - +casadi
  - include
  - lib

```
addpath( '/home/me/software/casadi-matlabR2014b-v3.0.0' )
```

```
import casadi.*
```

# CasADI

SX, MX, Function,  
AD, JIT  
nlp, qpsol, ...

Ipopt plugin,  
Snopt plugin,  
OOQP plugin,  
Ecos plugin,  
Mosek plugin,  
....

Developers: Joel Andersson, Joris Gillis  
Greg Horn, Niels van Duijkeren



**CasADI**

Go write your dynamic optimization tool...

```

from optoy import *

x = state()
y = state()
q = state()
u = control()

T = var(lb=0,init=10)


x.dot = (1-y**2)*x-y+u
y.dot = x
q.dot = x**2+y**2


ocp(T,[u>=-1,u<=1,q.start==0,x.start==1,y.start==0,x.end==0,y.end==0],T=T,N=20)

plot(x.sol)
plot(y.sol)
plot(u.sol)

```

## e.g. Optoy







## optistack

A simple matlab interface to casadi

The goal of this project is to provide a Yalmip-like Matlab interface to casadi.

```
x = optivar();
y = optivar();

nlp = optisolve((1-x)^2+100*(y-x^2)^2, {x^2+y^2<=1, x+y>=0});

optival(x)
optival(y)
```

# CasADI

Go write your dynamic optimization tool...

Prototype-style speed of development  
State-of-the-art performance

# Thanks for your attention

Let's do some coding...

