



Sandia  
National  
Laboratories

# Parallel Solution of Optimal Gas Network Control Under Uncertainty

Michael Bynum, Larry Biegler, Carl Laird, Sakshi Naik, Robert Parker, and John Siirola






Sandia National Laboratories is a multimission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.









**Goal:** Minimize compression costs of a natural gas distribution system while satisfying demands

# Optimal Compressor Boost Pressure



Optimal compressor boost pressures depend heavily on demands.



# A Stochastic Programming Solution





**A 2-node network with 32 scenarios takes  
1.7 minutes to solve with Ipopt**

# Turn to Parallel Computing



Parallel approaches have demonstrated tremendous success on many problems when tailored to the problem class


**Need software frameworks to support rapid innovation in serial and parallel algorithm development**


- Rapid development in high-level languages
- Support for block-based representations
- Strong serial and parallel computational performance

# PyNumero Overview




PyNumero: A high-level Python framework for rapid development of truly performant serial and parallel optimization algorithms on distributed HPC





$$\begin{bmatrix} K_1 & & B_1 \\ & K_2 & B_2 \\ & & K_3 \\ & & & \ddots \\ & & & & K_N & B_N \\ B_1^T & B_2^T & B_3^T & \cdots & B_N^T & Q \end{bmatrix} \cdot \begin{bmatrix} \Delta v_1 \\ \Delta v_2 \\ \Delta v_3 \\ \vdots \\ \Delta v_N \\ \Delta d \end{bmatrix} = \begin{bmatrix} r_1 \\ r_2 \\ r_3 \\ \vdots \\ r_N \\ r_d \end{bmatrix}$$

# Originally Developed for Dynamic Optimization



After full-space discretization:

$$\begin{aligned} \min_z \quad & f(z) \\ \text{s.t.} \quad & c(z, q) = 0 \\ & \underline{z} \leq z \leq \bar{z} \end{aligned}$$

$$c(z, q) = \begin{bmatrix} \bar{G}z_1 - x_0 \\ R(z_1) \\ \underline{G}z_1 + q_1 \\ \bar{G}z_2 - q_1 \\ R(z_2) \\ \underline{G}z_2 + q_2 \\ \vdots \\ \bar{G}z_{n_e} - q_{n_e-1} \\ R(z_{n_e}) \end{bmatrix}.$$

Collocation equations



Linking constraints

[Adapted from L.T. Biegler (2007)]

# Schur Complements for PinT Optimization



$$z = (u, x, y)$$



$$c(z, q) = \begin{bmatrix} \bar{G}z_1 - x_0 \\ R(z_1) \\ \bar{G}z_1 + q_1 \\ \bar{G}z_2 - q_1 \\ R(z_2) \\ \bar{G}z_2 + q_2 \\ \vdots \\ \bar{G}z_{n_e} - q_{n_e-1} \\ R(z_{n_e}) \end{bmatrix}.$$

Collocation equations



Linking constraints

[Adapted from L.T. Biegler (2007)]

# Schur Complements for PinT Optimization



$$z = (u, x, y)$$




- Factor all  $K_i$  matrices ( $\downarrow$ )
- Form Schur-complement ( $\approx$  or  $\downarrow$ )
  - Backsolve of each  $K_i$  for each nonzero column in  $A_i$  to form local Schur-complement  $S_i$
- Compute  $S = \sum_i S_i$  (MPI Reduction) ( $\uparrow$  communication)
- Solve  $S\Delta q = r_s$  ( $\uparrow$ )
- Solve  $K_i\Delta z_i = r_i$  ( $\downarrow$ )

Based on Kang, Cao, Word and Laird (2014).

[Adapted from L.T. Biegler (2007)]

# Scalability



Approximately 360x  
speedup on 1024 cores!

# A Note on Communication



	Dynamic Optimization	Stochastic Optimization	Multi-stage stochastic optimization	Network-Based Decomposition
Schur- Complement Structure	Sparse	Dense	Sparse	Sparse

Analyze the structure once during symbolic factorization, and exploit it for performant communication during numeric factorization.

# BlockVector/Matrix, MPIBlockVector/Matrix




Efficient block structures are core to PyNumero

- Many algorithms intuitively represented with blocks
- Critical for parallel, block decomposition approaches

PyNumero BlockVector and BlockMatrix classes

- Support creation of block-based vectors and matrices without data duplication
- BlockVector and BlockMatrix store C++ pointers to underlying data structures
- Convenient for performing block operations (e.g., formation and solution of KKT system)
- Support (almost all) standard operations for Numpy vectors and matrices
- Full support to build and interrogate these structures in both serial and parallel

$$\begin{bmatrix} W_k + \Sigma_k + \delta_w I & \nabla c(x_k) \\ \nabla c(x_k)^T & -\delta_c I \end{bmatrix} \begin{bmatrix} \Delta x \\ \Delta \lambda \end{bmatrix} = K v$$



# What does the code look like...



Working with matrices and vectors looks just like Numpy, and (almost all) Numpy / SciPy operations are supported – including solvers

```
20  # ...
21  # working with matrices and vectors
22  y = A.dot(x)          # dot product
23  z = x + y            # addition
24  zm = abs(z).max()    # infinity norm
25
26  # this code looks the same even if these are
27  # a dense numpy vector or Pynumero BlockVector
```

```
2  # ...
3  # working with nlp interfaces
4  nlp.set_primals(x)
5  nlp.set_duals(lam)
6
7  grad_lag = (nlp.evaluate_grad_objective() +
8      nlp.evaluate_jacobian().transpose() * lam)
9
10 residuals = nlp.evaluate_constraints()
11
12 jacobian = nlp.evaluate_jacobian()
13 hessian_lag = nlp.evaluate_hessian_lag()
14
```

Interfaces to the NLP classes are clean, easy to use, and easy to integrate with Numpy and SciPy

# What does parallel code look like...



BlockVector and BlockMatrix are intuitive and performant – references stored internally

```

4  # create vector x
5  x = BlockVector(3)
6  x.set_block(0, np.random.normal(size=3))
7  x.set_block(1, np.random.normal(size=3))
8  x.set_block(2, np.random.normal(size=3))

9
10 # create vector y
11 y = BlockVector(3)
12 y.set_block(0, np.random.normal(size=3))
13 y.set_block(1, np.random.normal(size=3))
14 y.set_block(3, np.random.normal(size=3))

15
16 # perform operations
17 z1 = x + y          # add x and y
18 z2 = x.dot(y)        # dot product
19 z3 = np.abs(x).max() # infinity norm

```

```

5  # get communicator and process #
6  comm = MPI.COMM_WORLD
7  rank = comm.Get_rank()

8
9  # specify the parallel ownership
10 owners = [2, 0, 1]

11
12 # create vector x (each processor does local block)
13 x = MPIBlockVector(3, rank_owner=owners, mpi_comm=comm)
14 x.set_block(owners.index(rank), np.random.normal(size=3))

15
16 # create vector y (each processor does local block)
17 y = MPIBlockVector(3, rank_owner=owners, mpi_comm=comm)
18 y.set_block(owners.index(rank), np.random.normal(size=3))

19
20 # perform parallel operations
21 z1 = x + y          # add x and y
22 z2 = x.dot(y)        # dot product
23 z3 = np.abs(x).max() # infinity norm

```

Parallel code can be written at a high level with mpi4py and the PyNumero MPIBlockVector and MPIBlockMatrix classes

# Writing Algorithms: Equality-Constrained SQP Example



```

1  def sqp(nlp, max_iter=100, tol=1e-8):
2      # setup KKT matrix
3      kkt = BlockMatrix(2, 2)
4      rhs = BlockVector(2)
5
6      # create and initialize the iteration vector
7      x = BlockVector(2)
8      x.set_block(0, nlp.init_primals().copy())
9      x.set_block(1, nlp.init_duals().copy())
10
11     # create the linear solver
12     linear_solver = MA27Interface()
13     linear_solver.set_cntl(1, 1e-6) # pivot tolerance
14
15     # main iteration loop
16     for _iter in range(max_iter):
17         nlp.set_primals(x.get_block(0))
18         nlp.set_duals(x.get_block(1))
19
20         grad_lag = (nlp.evaluate_grad_objective() +
21                     nlp.evaluate_jacobian().transpose() * x.get_block(1))
22         residuals = nlp.evaluate_constraints()
23
24         if np.abs(grad_lag).max() <= tol and np.abs(residuals).max() <= tol:
25             break
26
27         kkt.set_block(0, 0, nlp.evaluate_hessian_lag())
28         kkt.set_block(1, 0, nlp.evaluate_jacobian())
29         kkt.set_block(0, 1, nlp.evaluate_jacobian().transpose())
30
31         rhs.set_block(0, grad_lag)
32         rhs.set_block(1, residuals)
33
34         linear_solver.do_symbolic_factorization(kkt)
35         linear_solver.do_numeric_factorization(kkt)
36         delta = linear_solver.do_back_solve(-rhs)
37         x += delta
38

```

Equality-constrained SQP method

- Makes use of BlockVector / BlockMatrix to form KKT, RHS, and iteration vector
- BlockMatrix sent directly to the linear solver
- BlockMatrix efficiently stores pointers to sub-blocks

# What about performance?

```

1  def sqp(nlp, max_iter=100, tol=1e-8):
2      # setup KKT matrix
3      kkt = BlockMatrix(2, 2)
4      rhs = BlockVector(2)
5
6      # create and initialize the iteration vector
7      x = BlockVector(2)
8      x.set_block(0, nlp.init_primals().copy())
9      x.set_block(1, nlp.init_duals().copy())
10
11     # create the linear solver
12     linear_solver = MA27Interface()
13     linear_solver.set_cntl(1, 1e-6) # pivot tolerance
14
15     # main iteration loop
16     for _iter in range(max_iter):
17         nlp.set_primals(x.get_block(0))
18         nlp.set_duals(x.get_block(1))
19
20         grad_lag = (nlp.evaluate_grad_objective() +
21                     nlp.evaluate_jacobian().transpose() * x.get_block(1))
22         residuals = nlp.evaluate_constraints()
23
24         if np.abs(grad_lag).max() <= tol and np.abs(residuals).max() <= tol:
25             break
26
27         kkt.set_block(0, 0, nlp.evaluate_hessian_lag())
28         kkt.set_block(1, 0, nlp.evaluate_jacobian())
29         kkt.set_block(0, 1, nlp.evaluate_jacobian().transpose())
30
31         rhs.set_block(0, grad_lag)
32         rhs.set_block(1, residuals)
33
34         linear_solver.do_symbolic_factorization(kkt)
35         linear_solver.do_numeric_factorization(kkt)
36         delta = linear_solver.do_back_solve(-rhs)
37         x += delta
38

```

## Equality-constrained SQP method

- Makes use of BlockVector / BlockMatrix to form KKT, RHS, and iteration vector
- BlockMatrix sent directly to the linear solver
- BlockMatrix efficiently stores pointers to sub-blocks

## Performance compared with IPOPT (fully compiled C++)

- Discretized PDE control problem (Burgers)
- PyNumero ESQP within 15% of IPOPT on reasonable problems (100,000 vars, ~2 sec)
- PyNumero directly comparable to compiled C++ when the problem is sufficiently large.



**A 2-node network with 32 scenarios takes 6.7 seconds to solve with Parapint and 16 processes**

# Remarks and Future Work



## Remarks:

- Schur-Complement decomposition provides a scalable solution approach for stochastic optimal control of gas pipeline networks
- PyNumero provides a flexible, high-level Python framework for creating nonlinear optimization solvers
- PyNumero designed to facilitate rapid innovation in parallel algorithm development
- Full support for block representations of linear algebra
- Intuitive interface for NLP and algorithm creation with support for Numpy and SciPy
- Interfaces to commonly used linear solvers for NLP algorithms (SciPy, Mumps, HSL)
- Computationally expensive operations performed with compiled code
- Great scalability to over 1000 cores!

## Future work

- Improved robustness of Parapint's interior point algorithm
- More algorithms!
- Combined scenario and time decomposition for stochastic optimal control via nested Schur-Complement

# Acknowledgements



This work was conducted as part of the Institute for the Design of Advanced Energy Systems (IDAES) with funding from the Office of Fossil Energy, Cross-Cutting Research, U.S. Department of Energy.

**Disclaimer:** This presentation was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

Sandia National Laboratories is a multi-mission laboratory managed and operated by National Technology & Engineering Solutions of Sandia, LLC, a wholly owned subsidiary of Honeywell International Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA0003525.