Task graph-based performance analysis of PinT methods

 ${\rm Jens} \; {\rm Hahne}^1 \\ {\rm In \; collaboration \; with: \; Stephanie \; Friedhoff^1 \; and \; Matthias \; Bolten^2} \\$ 

<sup>1</sup>Bergische Universität Wuppertal, Germany



MATHEMATICAL MODELLING, ANALYSIS AND COMPUTATIONAL MATHEMATICS



BERGISCHE UNIVERSITÄT WUPPERTAL

#### ► Goals:

- Performance model for PinT methods
- Static load balacing
- Dynamic load balancing



#### ► Goals:

- Performance model for PinT methods
- Static load balacing
- Dynamic load balancing



#### Goals:

- Performance model for PinT methods
- Static load balacing
- Dynamic load balancing
- Challenges:
  - More frameworks than methods (variations with same idea)
  - Large parameter space for each method



#### Goals:

- Performance model for PinT methods
- Static load balacing
- Dynamic load balancing
- Challenges:
  - More frameworks than methods (variations with same idea)
  - Large parameter space for each method
- Approach:
  - Task-graph based on data-driven formulation of algorithms
  - Allows to cover any implementation using schedules of the tasks
  - Well known tool for load balancing techniques
  - Theoretical lower runtime bound based entirely on the algorithm



#### Parareal

• Decompose [0,T] into N non-overlapping intervals

$$0 = T_0 < T_1 < \dots < T_N = T$$

Two propagation operators:

- *F* expensive and high accurate
- G cheap but less accurate
- Initalization:

$$\mathbf{u}_n^0 = \mathcal{G}(T_n, T_{n-1}, \mathbf{u}_{n-1}^0)$$
 for  $n = 1, ..., N$ ,

Parareal iteration:

 $\mathbf{u}_n^{k+1} = \mathcal{F}(T_n, T_{n-1}, \mathbf{u}_{n-1}^k) + \mathcal{G}(T_n, T_{n-1}, \mathbf{u}_{n-1}^{k+1}) - \mathcal{G}(T_n, T_{n-1}, \mathbf{u}_{n-1}^k)$ 

for n = 1, ..., N and k = 0, ...

Task graph-based performance analysis of PinT methods Jens Hahne



#### Algorithm: Parareal

1  $\mathbf{u}_0^0 \leftarrow \mathbf{u}_0$ 2 for  $i \leftarrow 1$  to N ▷ Compute initial guess  $\mathbf{\tilde{u}}_{i}^{0} \leftarrow \mathcal{G}(T_{i}, T_{i-1}, \mathbf{u}_{i-1}^{0})$ 4  $\mathbf{u}_{i}^{0} \leftarrow \tilde{\mathbf{u}}_{i}^{0}$ 5 for  $k \leftarrow 1$  to N > Parareal iterations foreach  $i \in \{k, k + 1, ..., N\}$  do 6  $\hat{\mathbf{u}}_{i}^{k-1} \leftarrow \mathcal{F}(T_{i}, T_{i-1}, \mathbf{u}_{i-1}^{k-1})$ 7 for  $i \leftarrow k$  to N 8  $\tilde{\mathbf{u}}_{i}^{k} \leftarrow \mathcal{G}(T_{i}, T_{i-1}, \mathbf{u}_{i-1}^{\min(k, i-1)})$ q  $\mathbf{u}_i^k \leftarrow \mathbf{\tilde{u}}_i^k + \mathbf{\hat{u}}_i^{k-1} - \mathbf{\tilde{u}}_i^{k-1}$ 10

• Typically only  $K \ll N$  iterations required



#### Parallel Full Approximation Scheme in Space and Time (PFASST)



## PFASST

- Parallel Full Approximation Scheme in Space and Time (PFASST)
- Sort of Parareal approach, but ...
  - Deferred correction approach instead of fine and coarse propagator
  - Arbitrary hierarchy of levels
  - Coarse level problems are modified using a space-time FAS correction



## PFASST

- Parallel Full Approximation Scheme in Space and Time (PFASST)
- Sort of Parareal approach, but ...
  - Deferred correction approach instead of fine and coarse propagator
  - Arbitrary hierarchy of levels
  - Coarse level problems are modified using a space-time FAS correction
- Multigrid method
  - Type of space-time multigrid method
  - "Smoother" in the time direction is a spectral deferred correction (SDC) sweep



## PFASST

- Parallel Full Approximation Scheme in Space and Time (PFASST)
- Sort of Parareal approach, but ...
  - Deferred correction approach instead of fine and coarse propagator
  - Arbitrary hierarchy of levels
  - Coarse level problems are modified using a space-time FAS correction
- Multigrid method
  - Type of space-time multigrid method
  - "Smoother" in the time direction is a spectral deferred correction (SDC) sweep
- Spectral deferred correction methods
  - Pipelined version of multi-level SDC with each time slice performing SDC sweeps in parallel



## Multigrid-reduction-in-time (MGRIT)



- Reduction-based time-multigrid method
- In specialized two-level setting: MGRIT  $\equiv$  Parareal



Stopping criterion required for all three iterative methods



- Stopping criterion required for all three iterative methods
- Measures the quality of the solution



- Stopping criterion required for all three iterative methods
- Measures the quality of the solution
- Various criteria exist:

► ...

- Jump of the approximation between two iterations
- Finest-level information
- Residual at single time points
- Space-time residual





- Stopping criterion required for all three iterative methods
- Measures the quality of the solution
- Various criteria exist:
  - Jump of the approximation between two iterations
  - Finest-level information
  - Residual at single time points
  - Space-time residual
  - ► ...
- Criteria hard to compare

- Stopping criterion required for all three iterative methods
- Measures the quality of the solution
- Various criteria exist:
  - Jump of the approximation between two iterations
  - Finest-level information
  - Residual at single time points
  - Space-time residual
  - ► ...
- Criteria hard to compare
- High level categorization into two groups:
  - Local: Measurement of solution quality at individual points in time
  - Global: Global measurement of solution quality



- Stopping criterion required for all three iterative methods
- Measures the quality of the solution
- Various criteria exist:
  - Jump of the approximation between two iterations
  - Finest-level information
  - Residual at single time points
  - Space-time residual
  - ► ...
- Criteria hard to compare
- High level categorization into two groups:
  - Local: Measurement of solution quality at individual points in time
  - Global: Global measurement of solution quality
- Both can be formulated in a data-driven formulation



• Directed acyclic graph 
$$G = (V, E, \omega, c)$$



- Directed acyclic graph  $G = (V, E, \omega, c)$
- Vertices  $V = \{v_1, ... v_n\}$ 
  - Represent tasks





- ▶ Directed acyclic graph  $G = (V, E, \omega, c)$
- Vertices  $V = \{v_1, ... v_n\}$ 
  - Represent tasks
- $\blacktriangleright \text{ Directed edges } E \subseteq V \times V$ 
  - Represent dependencies of tasks





- $\blacktriangleright$  Directed acyclic graph  $G = (V, E, \omega, c)$
- Vertices  $V = \{v_1, \dots, v_n\}$ 
  - Represent tasks
- $\blacktriangleright$  Directed edges  $E \subseteq V \times V$ 
  - Represent dependencies of tasks
- ▶ Node weights  $\omega: V \to \mathbb{R}_0^+$ 
  - Represent computational cost of tasks

(	$v_1$ $\omega(v_1)$	
$v_2$ $\omega(v_2)$		$v_3 \\ \omega(v_3)$
(	$\left( \begin{array}{c} v_4 \\ \overline{\omega(v_4)} \end{array} \right)$	



7/21

- Directed acyclic graph  $G = (V, E, \omega, c)$
- Vertices  $V = \{v_1, ... v_n\}$ 
  - Represent tasks
- $\blacktriangleright \text{ Directed edges } E \subseteq V \times V$ 
  - Represent dependencies of tasks
- ▶ Node weights  $\omega: V \to \mathbb{R}_0^+$ 
  - Represent computational cost of tasks
- Edge weights  $c: E \to \mathbb{R}^+_0$ 
  - Represent communication cost between tasks





• 
$$P = \{p_1, \dots, p_{N_P}\}$$
  
• Set of  $N_P$  processes



- $\blacktriangleright P = \{p_1, \ldots, p_{N_P}\}$ 
  - Set of N<sub>P</sub> processes
- Allocation function  $A: V \to P$ 
  - Assigns each task in V to a process

- $\blacktriangleright P = \{p_1, \ldots, p_{N_P}\}$ 
  - Set of N<sub>P</sub> processes
- Allocation function  $A: V \to P$ 
  - Assigns each task in V to a process
- $\blacktriangleright \text{ Schedule } S: V \to \mathbb{R}_0^+$ 
  - Assigns a starting point to each task, subject to the constraints:

$$\forall (v_i, v_j) \in E, S(v_j) \ge S(v_i) + \omega(v_i) + c(v_i, v_j) \\ \forall v_i, v_j \in V, v_i \neq v_j, A(v_i) = A(v_j) \Rightarrow S(v_i) \ge \\ S(v_j) + \omega(v_j) \lor S(v_j) \ge S(v_i) + \omega(v_i)$$



- $\blacktriangleright P = \{p_1, \ldots, p_{N_P}\}$ 
  - Set of  $N_P$  processes
- Allocation function  $A: V \to P$ 
  - Assigns each task in V to a process
- $\blacktriangleright \text{ Schedule } S: V \to \mathbb{R}_0^+$ 
  - Assigns a starting point to each task, subject to the constraints:

$$\forall (v_i, v_j) \in E, S(v_j) \ge S(v_i) + \omega(v_i) + c(v_i, v_j) \\ \forall v_i, v_j \in V, v_i \neq v_j, A(v_i) = A(v_j) \Rightarrow S(v_i) \ge \\ S(v_j) + \omega(v_j) \lor S(v_j) \ge S(v_i) + \omega(v_i)$$

Makespan or runtime of a given allocation and schedule:

$$\max_{v \in V} (S(v) + \omega(v))$$



- $\blacktriangleright P = \{p_1, \ldots, p_{N_P}\}$ 
  - $\blacktriangleright$  Set of  $N_P$  processes
- ▶ Allocation function  $A: V \to P$ 
  - Assigns each task in V to a process
- $\blacktriangleright$  Schedule  $S: V \to \mathbb{R}^+_0$ 
  - Assigns a starting point to each task, subject to the constraints:

$$\forall (v_i, v_j) \in E, S(v_j) \ge S(v_i) + \omega(v_i) + c(v_i, v_j) \\ \forall v_i, v_j \in V, v_i \neq v_j, A(v_i) = A(v_j) \Rightarrow S(v_i) \ge \\ S(v_j) + \omega(v_j) \lor S(v_j) \ge S(v_i) + \omega(v_i)$$

Makespan or runtime of a given allocation and schedule:

$$\max_{v \in V} (S(v) + \omega(v))$$

▶ If  $c(e) = 0 \forall e \in E$ : Minimum possible makespan for  $N_P = \infty$  is longest path within the graph

Task graph-based performance analysis of PinT methods Jens Hahne



8/21







$$N_P = 2$$

$$A(v)$$

$$v_1 \rightarrow p_1$$

$$v_2 \rightarrow p_1$$

$$v_3 \rightarrow p_2$$

$$v_4 \rightarrow p_2$$

$$v_5 \rightarrow p_1$$









Runtime

9/21

8







Task graph-based performance analysis of  $\mathsf{PinT}$  methods  $\mathsf{Jens}$  Hahne













Task graph-based performance analysis of PinT methods Jens Hahne










## Schedule example





Makespan of schedule: 6 



## Schedule example



- Makespan of schedule: 6
- ▶ Longest path  $\{v_1 \rightarrow v_2 \rightarrow v_5\} \Rightarrow$  minimal parallel runtime: 4



#### Algorithm: Parareal

 $\begin{array}{c|c|c} \mathbf{1} & \mathbf{u}_0^0 \leftarrow \mathbf{u}_0 \\ \mathbf{2} & \text{for } i \leftarrow 1 \text{ to } N \\ \mathbf{3} & & \mathbf{\tilde{u}}_i^0 \leftarrow \mathcal{G}(T_i, T_{i-1}, \mathbf{u}_{i-1}^0) \\ \mathbf{4} & & & \mathbf{u}_i^0 \leftarrow \mathbf{\tilde{u}}_i^0 \\ \mathbf{5} & \dots \end{array}$ 

▷ Initial guess







#### Algorithm: Parareal

$$\begin{array}{c|c|c} 1 & \mathbf{u}_0^0 \leftarrow \mathbf{u}_0 \\ 2 & \text{for } i \leftarrow 1 \text{ to } N \\ 3 & \mid & \mathbf{\tilde{u}}_i^0 \leftarrow \mathcal{G}(T_i, T_{i-1}, \mathbf{u}_{i-1}^0) \\ 4 & \mid & \mathbf{u}_i^0 \leftarrow \mathbf{\tilde{u}}_i^0 \end{array}$$

▷ Initial guess







#### Algorithm: Parareal





#### Algorithm: Parareal





#### Algorithm: Parareal





#### Algorithm: Parareal





#### Algorithm: Parareal

$$\begin{array}{c|c} \mathbf{1} & \mathbf{u}_0^0 \leftarrow \mathbf{u}_0 \\ \mathbf{2} & \text{for } i \leftarrow 1 \text{ to } N \\ \mathbf{3} & | & \mathbf{\tilde{u}}_i^0 \leftarrow \mathcal{G}(T_i, T_{i-1}, \mathbf{u}_{i-1}^0) \\ \mathbf{4} & | & \mathbf{u}_i^0 \leftarrow \mathbf{\tilde{u}}_i^0 \\ \mathbf{5} & \dots \end{array} \right.$$





#### Algorithm: Parareal







Task graph-based performance analysis of  $\mathsf{PinT}$  methods  $\mathsf{Jens}$  Hahne

 $\triangleright k = 1$ 































Typically static allocation of time points to processes



- Typically static allocation of time points to processes
- Scheduling based on this allocation



- Typically static allocation of time points to processes
- Scheduling based on this allocation
- Two types of allocation:



- Typically static allocation of time points to processes
- Scheduling based on this allocation
- Two types of allocation:
  - Block-by-block basis (each block roughly same number of time points)





- Typically static allocation of time points to processes
- Scheduling based on this allocation
- Two types of allocation:
  - Block-by-block basis (each block roughly same number of time points)
  - Windowing with block-by-block (applying method multiple times)

• Example: 
$$N = 9$$
,  $N_P = 3$ 





Task graph-based performance analysis of  $\mathsf{PinT}$  methods  $\mathsf{Jens}$  Hahne





12/21

How to set the costs?







- How to set the costs?
- Two ways:
  - Theoretical model:
    - + Exact and theoretical analysis
    - + Independent of implementation
    - Huge effort
    - Not very flexible, models one problem
  - Measure runtimes of operations:
    - Measure any type of operation
    - + Easy and flexible
    - Requires implementation





- How to set the costs?
- Two ways:
  - Theoretical model:
    - + Exact and theoretical analysis
    - + Independent of implementation
    - Huge effort
    - Not very flexible, models one problem
  - Measure runtimes of operations:
    - Measure any type of operation
    - + Easy and flexible
    - Requires implementation
- Decision depends on use case





Compare model with four Pint libraries:

- LibPFASST<sup>1</sup>: Fortran implementation of PFASST
- PySDC<sup>2</sup>: Python implementation of PFASST
- XBRAID<sup>3</sup>: C implementation of MGRIT
- PyMGRIT<sup>4</sup>: Python implementation of MGRIT

<sup>1</sup>https://github.com/libpfasst/LibPFASST

<sup>2</sup>https://github.com/Parallel-in-Time/pySDC

<sup>3</sup>https://github.com/XBraid/xbraid

<sup>4</sup>https://github.com/pymgrit/pymgrit



Compare model with four Pint libraries:

- ► LibPFASST<sup>1</sup>: Fortran implementation of PFASST
- PySDC<sup>2</sup>: Python implementation of PFASST
- XBRAID<sup>3</sup>: C implementation of MGRIT
- PyMGRIT<sup>4</sup>: Python implementation of MGRIT
- Each library has a large set of pre-implemented features

<sup>1</sup>https://github.com/libpfasst/LibPFASST

<sup>2</sup>https://github.com/Parallel-in-Time/pySDC

<sup>3</sup>https://github.com/XBraid/xbraid

<sup>4</sup>https://github.com/pymgrit/pymgrit



Compare model with four Pint libraries:

- LibPFASST<sup>1</sup>: Fortran implementation of PFASST
- PySDC<sup>2</sup>: Python implementation of PFASST
- XBRAID<sup>3</sup>: C implementation of MGRIT
- PyMGRIT<sup>4</sup>: Python implementation of MGRIT
- Each library has a large set of pre-implemented features
- Each library requires problem-dependent functions

<sup>1</sup>https://github.com/libpfasst/LibPFASST

<sup>2</sup>https://github.com/Parallel-in-Time/pySDC

<sup>3</sup>https://github.com/XBraid/xbraid

<sup>4</sup>https://github.com/pymgrit/pymgrit



- Compare model with four Pint libraries:
  - LibPFASST<sup>1</sup>: Fortran implementation of PFASST
  - PySDC<sup>2</sup>: Python implementation of PFASST
  - XBRAID<sup>3</sup>: C implementation of MGRIT
  - PyMGRIT<sup>4</sup>: Python implementation of MGRIT
- Each library has a large set of pre-implemented features
- Each library requires problem-dependent functions
- Challenging to choose problem
  - Different languages
  - Methods with different strengths

<sup>1</sup>https://github.com/libpfasst/LibPFASST

<sup>2</sup>https://github.com/Parallel-in-Time/pySDC

<sup>3</sup>https://github.com/XBraid/xbraid

<sup>4</sup>https://github.com/pymgrit/pymgrit



- Compare model with four Pint libraries:
  - ► LibPFASST<sup>1</sup>: Fortran implementation of PFASST
  - PySDC<sup>2</sup>: Python implementation of PFASST
  - XBRAID<sup>3</sup>: C implementation of MGRIT
  - PyMGRIT<sup>4</sup>: Python implementation of MGRIT
- Each library has a large set of pre-implemented features
- Each library requires problem-dependent functions
- Challenging to choose problem
  - Different languages
  - Methods with different strengths
- Pseudo spatial problem:
  - Each required function sleeps only for a controllable time
  - Enables large test sets in terms of task costs

<sup>1</sup>https://github.com/libpfasst/LibPFASST

<sup>2</sup>https://github.com/Parallel-in-Time/pySDC

<sup>3</sup>https://github.com/XBraid/xbraid

<sup>4</sup>https://github.com/pymgrit/pymgrit



# PFASST - setting

LibPFASST	Parameters				
	Number of time intervals N				
	Number processes time dimension $N_P$				
Burn-in prediction phase	Number of levels L				
	Number of iterations K				
► PySDC	Convergence criterion				
Multigrid view	Two "views"				
	Number of sweeps on level $\ell$				
Fine-sweep prediction phase	Collocation nodes on level $\ell$				
	Predictor type				
Both:	Skip fine-level sweep at start/end				
Local convergence criterion	Runtime communication in time				
	Runtime convergence criterion				
Implicit sweeper (pre-implemented)	Runtime SDCSweep at level $\ell$				
Cost per implicit solve: 0.05 s	Runtime FEvalAll at level $\ell$				
	Runtime FEvalSingle at level $\ell$				
All other costs: U s	Runtime RestrictAll at level $\ell$				
Runtime of operations is measured	Runtime RestrictSingle at level $\ell$				
(reuseble for other problems)	Runtime InterpolateAll at level $\ell$				
(reusable for other problems)	Runtime InterpolateSingle at level $\ell$				
Communication cost: 0	Runtime FAS at level $\ell$				



## LibPFASST - model vs. runtime

Ν	$N_P$	L	K	# Sweeps	collocation nodes	LibPFASST runtime	Model	L. b. $(N_P = \infty)$
16	16	2	10	(1, 1)	(5,3)	5.06	5.05	5.04
16	16	2	10	(1, 1)	(7, 5)	8.88	8.87	8.86
16	16	2	10	(2, 1)	(5,3)	7.07	7.04	7.03
16	16	2	10	(2, 1)	(7, 5)	11.9	11.89	11.88
32	32	4	10	$(1 \ 1 \ 1 \ 1)$	(7 5 3 2)	12 16	12 12	12 11
52	52	-	10	(1, 1, 1, 1)	(1, 0, 0, 2)	12.10	12.12	12.11
32	32	4	10	(2, 1, 1, 1)	(7, 5, 3, 2)	15.17	15.13	15.12
32	32	4	10	(1, 2, 1, 1)	(7,5,3,2)	16.17	16.14	16.13
32	32	4	10	(1, 1, 2, 1)	(7, 5, 3, 2)	14.17	14.14	14.13
32	32	4	10	(1, 2, 2, 1)	(7,5,3,2)	18.18	18.14	18.13



### LibPFASST - model vs. runtime

N	$N_P$	L	K	# Sweeps	collocation nodes	LibPFASST runtime	Model	L. b. $(N_P = \infty)$
16	16	2	10	(1, 1)	(5,3)	5.06	5.05	5.04
16	16	2	10	(1, 1)	(7, 5)	8.88	8.87	8.86
16	16	2	10	(2, 1)	(5,3)	7.07	7.04	7.03
16	16	2	10	(2, 1)	(7, 5)	11.9	11.89	11.88
32	32	4	10	(1, 1, 1, 1)	(7, 5, 3, 2)	12.16	12.12	12.11
32	32	4	10	(2, 1, 1, 1)	(7, 5, 3, 2)	15.17	15.13	15.12
32	32	4	10	(1, 2, 1, 1)	(7, 5, 3, 2)	16.17	16.14	16.13
32	32	4	10	(1, 1, 2, 1)	(7, 5, 3, 2)	14.17	14.14	14.13
32	32	4	10	(1, 2, 2, 1)	(7,5,3,2)	18.18	18.14	18.13

#### Actual runtime very close to predictions



## LibPFASST - model vs. runtime

N	$N_P$	L	Κ	# Sweeps	collocation nodes	LibPFASST runtime	Model	L. b. $(N_P = \infty)$
16	16	2	10	(1, 1)	(5,3)	5.06	5.05	5.04
16	16	2	10	(1, 1)	(7, 5)	8.88	8.87	8.86
16	16	2	10	(2, 1)	(5,3)	7.07	7.04	7.03
16	16	2	10	(2, 1)	(7, 5)	11.9	11.89	11.88
32	32	4	10	(1, 1, 1, 1)	(7, 5, 3, 2)	12 16	12 12	12 11
02	02	÷	10	(1, 1, 1, 1)	(1,0,0,2)	12.10	12.12	15.10
32	32	4	10	(2, 1, 1, 1)	(7, 5, 3, 2)	15.17	15.13	15.12
32	32	4	10	(1, 2, 1, 1)	(7, 5, 3, 2)	16.17	16.14	16.13
32	32	4	10	(1, 1, 2, 1)	(7,5,3,2)	14.17	14.14	14.13
32	32	4	10	(1, 2, 2, 1)	(7,5,3,2)	18.18	18.14	18.13

- Actual runtime very close to predictions
- Lower bound (L.b.) and predictions very similar


N	$M_{\rm D}$	T	K	# Swoons	collocation	pySDC	Model	L. b.
11	IVP	L	Π	# Sweeps	nodes	runtime	Model	$(N_P = \infty)$
16	16	2	10	(1, 1)	(5,3)	5.22	5.00	5.00
16	16	2	10	(1, 1)	(7, 5)	9.45	8.99	8.99
16	16	2	10	(2, 1)	(5,3)	7.38	7.07	7.07
16	16	2	10	(2, 1)	(7, 5)	12.78	12.23	12.23
32	32	4	10	(1, 1, 1, 1)	(7, 5, 3, 2)	12.82	12.34	12.34
32	32	4	10	(2, 1, 1, 1)	(7, 5, 3, 2)	16.23	15.57	15.56
32	32	4	10	(1, 2, 1, 1)	(7, 5, 3, 2)	17.16	16.41	16.30
32	32	4	10	(1, 1, 2, 1)	(7, 5, 3, 2)	14.96	14.26	13.22
32	32	4	10	(1, 2, 2, 1)	(7, 5, 3, 2)	19.38	18.57	15.43



N	N <sub>n</sub>	T	K	# Swoons	collocation	pySDC	Model	L. b.
11	IVP	L	Π	# Sweeps	nodes	runtime	Model	$(N_P = \infty)$
16	16	2	10	(1, 1)	(5,3)	5.22	5.00	5.00
16	16	2	10	(1, 1)	(7, 5)	9.45	8.99	8.99
16	16	2	10	(2, 1)	(5,3)	7.38	7.07	7.07
16	16	2	10	(2, 1)	(7, 5)	12.78	12.23	12.23
32	32	4	10	(1, 1, 1, 1)	(7, 5, 3, 2)	12.82	12.34	12.34
32	32	4	10	(2, 1, 1, 1)	(7, 5, 3, 2)	16.23	15.57	15.56
32	32	4	10	(1, 2, 1, 1)	(7, 5, 3, 2)	17.16	16.41	16.30
32	32	4	10	(1, 1, 2, 1)	(7, 5, 3, 2)	14.96	14.26	13.22
32	32	4	10	(1, 2, 2, 1)	(7, 5, 3, 2)	19.38	18.57	15.43

 $\blacktriangleright$  Actual runtime differs by no more than 5%



N	$M_{\rm D}$	T	K	# Swoons	collocation	pySDC	Model	L. b.
1 V	IVP	L	Λ	# Sweeps	nodes	runtime	Model	$(N_P = \infty)$
16	16	2	10	(1, 1)	(5,3)	5.22	5.00	5.00
16	16	2	10	(1, 1)	(7, 5)	9.45	8.99	8.99
16	16	2	10	(2, 1)	(5,3)	7.38	7.07	7.07
16	16	2	10	(2, 1)	(7, 5)	12.78	12.23	12.23
32	32	4	10	(1, 1, 1, 1)	(7, 5, 3, 2)	12.82	12.34	12.34
32	32	4	10	(2, 1, 1, 1)	(7, 5, 3, 2)	16.23	15.57	15.56
32	32	4	10	(1, 2, 1, 1)	(7, 5, 3, 2)	17.16	16.41	16.30
32	32	4	10	(1, 1, 2, 1)	(7, 5, 3, 2)	14.96	14.26	13.22
32	32	4	10	(1, 2, 2, 1)	(7, 5, 3, 2)	19.38	18.57	15.43

• Actual runtime differs by no more than 5%

Lower bound and predictions similiar for most settings



N	M-	т	K	# Sweens	collocation	pySDC	Model	L. b.
1 V	IVP	L	Λ	# Sweeps	nodes	runtime	model	$(N_P = \infty)$
16	16	2	10	(1, 1)	(5,3)	5.22	5.00	5.00
16	16	2	10	(1, 1)	(7, 5)	9.45	8.99	8.99
16	16	2	10	(2, 1)	(5,3)	7.38	7.07	7.07
16	16	2	10	(2, 1)	(7, 5)	12.78	12.23	12.23
32	32	4	10	(1, 1, 1, 1)	(7, 5, 3, 2)	12.82	12.34	12.34
32	32	4	10	(2, 1, 1, 1)	(7, 5, 3, 2)	16.23	15.57	15.56
32	32	4	10	(1, 2, 1, 1)	(7, 5, 3, 2)	17.16	16.41	16.30
32	32	4	10	(1, 1, 2, 1)	(7, 5, 3, 2)	14.96	14.26	13.22
32	32	4	10	(1, 2, 2, 1)	(7, 5, 3, 2)	19.38	18.57	15.43

• Actual runtime differs by no more than 5%

- Lower bound and predictions similiar for most settings
- Small derivations for some four-level settings (PFASST multigrid view)



VDDAID		
ADRAID	Parameters	
Global convergence criterion	Number of time intervals	N
clobal convergence criterion	Number processes time dimension	$N_P$
PyMGRIT	Number of levels	L
	Number of iterations	K
Local and global convergence	Convergence criterion	
criteria	Coarsening factor from level $\ell$ to level $\ell+1$	
citeria	Cycle type	
Both:	Nested iterations	
	Skip down	
Cost per time integration: 0.05 s	Number of $CF$ -relaxations on level $\ell$	
All other costs: 0 s	Runtime communication in time	
	Runtime convergence criterion	
Runtime of operations is measured	Runtime time integrator at level $\ell$	
(reusable for other problems)	Runtime spatial restriction at level $\ell$	
Communication cost: 0	Runtime spatial interpolation at level $\ell$	
Communication cost: 0		



## XBRAID - model vs. runtime

N	N.	т	K	C) (C	coorconing	# CF-	skip.	XBRAID	Model	L.b.
1 V	NP	L	Π	CyC.	coarsening	relax.	down	runtime	Model	$(N_P = \infty)$
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	True	19.02	18.93	15.61
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	False	22.79	21.83	17.91
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	True	21.59	21.51	13.63
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	False	23.90	23.01	14.99
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	True	22.81	22.66	11.88
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	False	26.38	26.16	13.43
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	True	30.71	30.64	15.17
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	False	32.61	32.49	16.04



N	M <sub>n</sub>	т	K	010	coorconing	# CF-	skip.	XBRAID	Model	L.b.
1 V	IVP	L	п	CyC.	coarsening	relax.	down	runtime	Model	$(N_P = \infty)$
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	True	19.02	18.93	15.61
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	False	22.79	21.83	17.91
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	True	21.59	21.51	13.63
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	False	23.90	23.01	14.99
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	True	22.81	22.66	11.88
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	False	26.38	26.16	13.43
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	True	30.71	30.64	15.17
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	False	32.61	32.49	16.04



N	M <sub>n</sub>	т	K	010	coorconing	# CF-	skip.	XBRAID	Model	L.b.
1 V	IVP	L	п	CyC.	coarsening	relax.	down	runtime	Model	$(N_P = \infty)$
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	True	19.02	18.93	15.61
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	False	22.79	21.83	17.91
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	True	21.59	21.51	13.63
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	False	23.90	23.01	14.99
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	True	22.81	22.66	11.88
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	False	26.38	26.16	13.43
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	True	30.71	30.64	15.17
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	False	32.61	32.49	16.04

Cluster based problem where the first communication between nodes is especially expensive



N	M <sub>n</sub>	т	K	0.40	coorconing	# CF-	skip.	XBRAID	Model	L.b.
1.	IVP	L	Π	CyC.	coarsening	relax.	down	runtime	Model	$(N_P = \infty)$
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	True	19.02	18.93	15.61
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	False	22.79	21.83	17.91
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	True	21.59	21.51	13.63
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	False	23.90	23.01	14.99
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	True	22.81	22.66	11.88
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	False	26.38	26.16	13.43
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	True	30.71	30.64	15.17
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	False	32.61	32.49	16.04

- Actual runtime very close to predictions
  - Cluster based problem where the first communication between nodes is especially expensive
- Theoretically further potential
  - Largely due to global convergence criterion (blocking communication)



N	M <sub>n</sub>	т	K	010	coorconing	# CF-	skip.	XBRAID	Model	L.b.
1.	IVP	L	Π	CyC.	coarsening	relax.	down	runtime	Model	$(N_P = \infty)$
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	True	19.02	18.93	15.61
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	False	22.79	21.83	17.91
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	True	21.59	21.51	13.63
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	False	23.90	23.01	14.99
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	True	22.81	22.66	11.88
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	False	26.38	26.16	13.43
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	True	30.71	30.64	15.17
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	False	32.61	32.49	16.04

- Actual runtime very close to predictions
  - Cluster based problem where the first communication between nodes is especially expensive
- Theoretically further potential
  - Largely due to global convergence criterion (blocking communication)
  - Comparison between 256 and  $\infty$  processes



N	N <sub>n</sub>	т	K	CV/C	coarconing	# CF-	nest.	conv.	PyMGRIT	Model	L.b.
14	$\mathbf{N}\mathbf{P}$	L	п	cyc.	coarsening	relax.	iter.	crit.	runtime	Model	$(N_P = \infty)$
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	Т	local	19.99	19.53	18.01
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	F	local	18.87	18.43	18.02
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	Т	local	17.69	17.17	15.09
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	F	local	16.98	16.66	15.09
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	Т	global	24.66	23.91	18.06
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	F	global	23.28	21.94	18.01
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	Т	global	25.08	24.04	15.07
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	F	global	24.24	23.15	15.08
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	т	local	25.82	25.36	13.52
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	Т	local	26.06	25.43	16.13
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	Т	global	30.91	29.97	13.49
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	Т	global	35.93	34.82	16.14



N	M <sub>n</sub>	т	K	010	coorsoning	# CF-	nest.	conv.	PyMGRIT	Model	L.b.
1.4	P	L	11	cyc.	coarsening	relax.	iter.	crit.	runtime	Model	$(N_P = \infty)$
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	Т	local	19.99	19.53	18.01
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	F	local	18.87	18.43	18.02
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	Т	local	17.69	17.17	15.09
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	F	local	16.98	16.66	15.09
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	Т	global	24.66	23.91	18.06
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	F	global	23.28	21.94	18.01
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	Т	global	25.08	24.04	15.07
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	F	global	24.24	23.15	15.08
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	т	local	25.82	25.36	13.52
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	Т	local	26.06	25.43	16.13
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	Т	global	30.91	29.97	13.49
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	Т	global	35.93	34.82	16.14

Same cluster based problem



Ν	$N_P$	L	K	cyc.	coarsening	# CF-	nest.	conv.	PyMGRIT	Model	L.b.
						relax.	iter.	crit.	runtime		$(N_P = \infty)$
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	Т	local	19.99	19.53	18.01
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	F	local	18.87	18.43	18.02
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	Т	local	17.69	17.17	15.09
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	F	local	16.98	16.66	15.09
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	Т	global	24.66	23.91	18.06
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	F	global	23.28	21.94	18.01
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	Т	global	25.08	24.04	15.07
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	F	global	24.24	23.15	15.08
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	Т	local	25.82	25.36	13.52
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	Т	local	26.06	25.43	16.13
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	Т	global	30.91	29.97	13.49
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	Т	global	35.93	34.82	16.14

- Same cluster based problem
- Model approaches the theoretical limit when local criterion is used



Ν	$N_P$	L	K	cyc.	coarsening	# CF-	nest.	conv.	PyMGRIT	Model	L.b.
						relax.	iter.	crit.	runtime		$(N_P = \infty)$
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	Т	local	19.99	19.53	18.01
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	F	local	18.87	18.43	18.02
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	Т	local	17.69	17.17	15.09
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	F	local	16.98	16.66	15.09
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	Т	global	24.66	23.91	18.06
4096	256	5	6	V	(16, 4, 4, 4)	(1, 1, 1, 1)	F	global	23.28	21.94	18.01
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	Т	global	25.08	24.04	15.07
4096	256	5	6	F	(16, 4, 4, 4)	(0, 0, 0, 0)	F	global	24.24	23.15	15.08
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	Т	local	25.82	25.36	13.52
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	Т	local	26.06	25.43	16.13
4096	256	6	6	V	(4, 4, 4, 4, 4)	(1, 1, 1, 1, 1)	Т	global	30.91	29.97	13.49
4096	256	6	6	F	(4, 4, 4, 4, 4)	(0, 0, 0, 0, 0)	Т	global	35.93	34.82	16.14

- Same cluster based problem
- Model approaches the theoretical limit when local criterion is used
- Preceding nested iteration theoretically with zero cost



### Outlook: Use task graph for scheduling

▶ Parareal example: N = 6, K = 5, Cost  $\mathcal{G}$ : 0.5 s, Cost  $\mathcal{F}$ : 2 s



Block-by-block:

$$\blacktriangleright N_P = 6 = N$$

Runtime: 15.5 s



#### Outlook: Use task graph for scheduling

▶ Parareal example: N = 6, K = 5, Cost  $\mathcal{G}$ : 0.5 s, Cost  $\mathcal{F}$ : 2 s



Task graph-based performance analysis of PinT methods Jens Hahne



Task-based analysis of PinT methods







This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955701. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Belgium, France, Germany, Switzerland. This work is supported by the BMBF (project TIME-X; grant no. 16HPC046K).

- Task-based analysis of PinT methods
- Covers large parameter space of PinT methods







This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955701. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Belgium, France, Germany, Switzerland. This work is supported by the BMBF (project TIME-X; grant no. 16HPC046K).

- Task-based analysis of PinT methods
- Covers large parameter space of PinT methods
- Theoretical lower runtime bound based on graph







This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955701. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Belgium, France, Germany, Switzerland. This work is supported by the BMBF (project TIME-X; grant no. 16HPC046K).

- Task-based analysis of PinT methods
- Covers large parameter space of PinT methods
- Theoretical lower runtime bound based on graph
- Runtime prediction using schedules







This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955701. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Belgium, France, Germany, Switzerland. This work is supported by the BMBF (project TIME-X; grant no. 16HPC046K).

- Task-based analysis of PinT methods
- Covers large parameter space of PinT methods
- Theoretical lower runtime bound based on graph
- Runtime prediction using schedules
- Future work:
  - Load balancing strategies
  - Combine model with convergence predictions



This project has received funding from the European High-Performance Computing Joint Undertaking (JU) under grant agreement No 955701. The JU receives support from the European Union's Horizon 2020 research and innovation programme and Belgium, France, Germany, Switzerland. This work is supported by the BMBF (project TIME-X; grant no. 16HPC046K).

Task graph-based performance analysis of PinT methods Jens Hahne



