1st International Workshop on **S**calable **W**orkflow **E**nactment **E**ngines and **T**echnologies

SWEET'12

EVALUATING PARAMETER SWEEP WORKFLOWS IN HIGH PERFORMANCE COMPUTING

Fernando Chirigati*, Vítor Silva, Eduardo Ogasawara, Daniel de Oliveira, Jonas Dias, Fábio Porto, Patrick Valduriez and Marta Mattoso

COPPE / Federal University of Rio de Janeiro, Rio de Janeiro, Brazil

LNCC, Rio de Janeiro, Brazil
INRIA & LIRMM, Montpellier, France

* Currently at Polytechnic Institute of NYU

Motivation

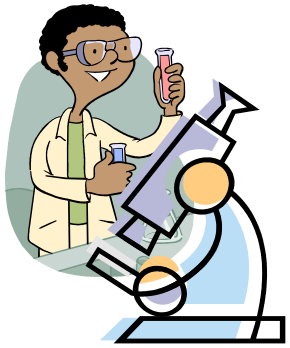

Scientists collect data

1

Data is analyzed

2

New set of data is produced

3

4

Data is processed

5

Results are analyzed

Parameter Sweep

Parameter Sweep (PS)

- Each iteration is often time-consuming
- Experiments have a large space of parameter values
- Candidate for High Performance Computing (HPC)
  - But it is not that simple…
    - There are different execution models
    - It significantly depends on the experiment (e.g., different workloads)

Simulation Framework

- Goal: Which execution model is more suitable for my experiment?
- Focus
  - PS workflows
    - Different patterns
    - Different activities
    - Different workload configurations
  - Measure and compare different execution models
  - Measure scalability

Characterization of PS Workflows

A Model of Scientific Workflow

- Data-centric workflow
  - Set of activities
  - Set of data
- Dependency order
  - Output of an activity is consumed as input by another activity
  - Execution needs to follow this dependency

PS Workflows

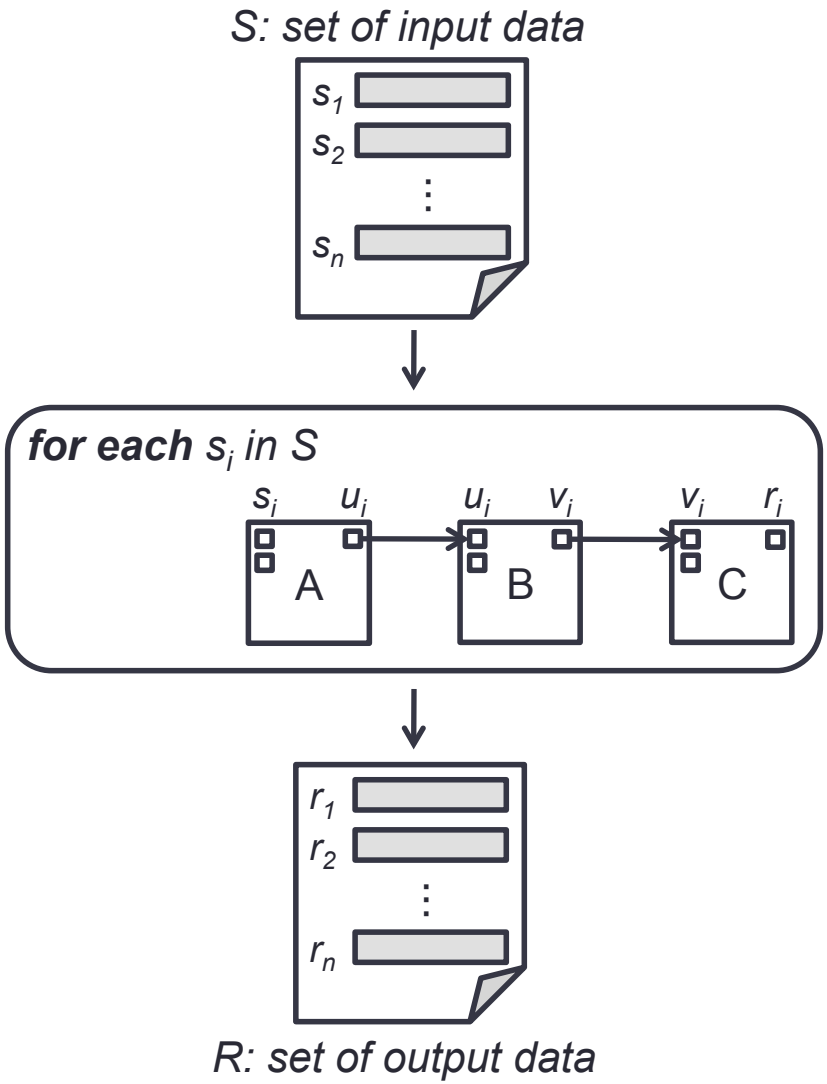

PS Workflow Patterns

- Related to *workflow patterns* (Russell et al. 2006)

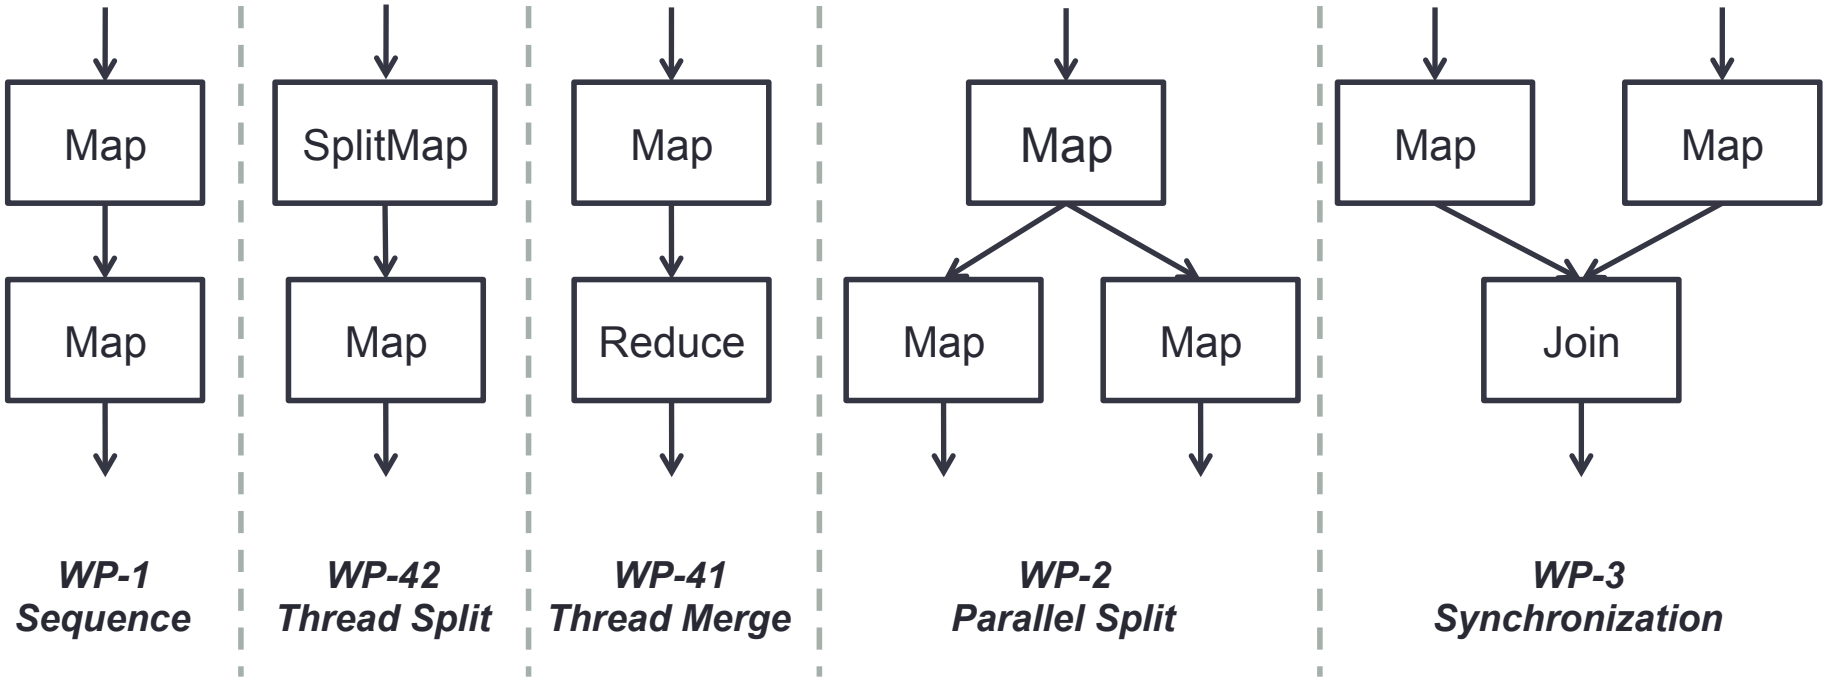

Evaluation Framework

Evaluation Framework

- Three main steps

1. Workload Configuration
2. Definition of PS Workflows
3. Performance Metrics

1 Workload Configuration

- Definition of *scaling factors*
- Instance Scale Factor (ISF)
  - Cardinality of input parameter space
- Activity Cost Factor (ACF)
  - Duration of activities

2 PS Workflows

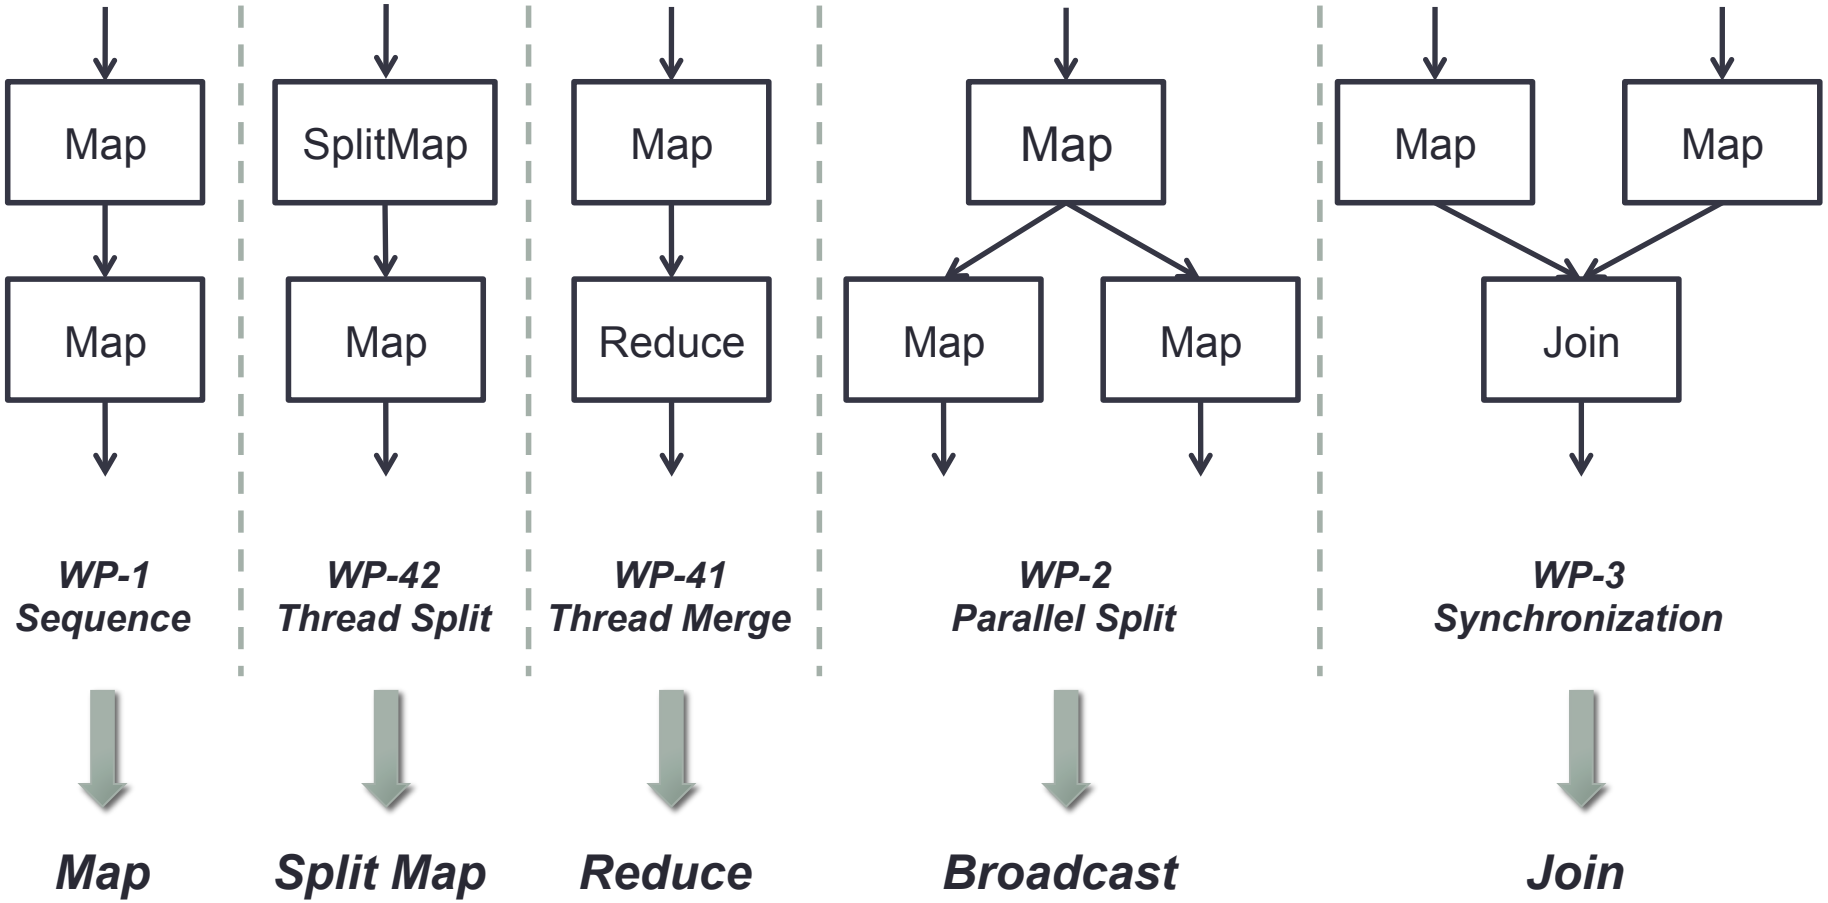

3 Performance Metrics

- Related to execution time
  - Elapsed Time (T_E)
  - Speedup (S)

$$S = T_1 / T_E$$

  - Efficiency (E)

$$E = S / p$$

  - Score (S_E)

$$S_E = \frac{1}{x} \sum_{i=1}^{x} E_i$$

Experimental Evaluation

Experimental Evaluation

- Total of 4 analysis
  - SGI Altix ICE 8200 distributed memory shared-disk cluster (32 nodes with 2 quadcore processors and 8 GB of memory each)
- Comparison of four execution models
  - Combination of two characteristics
    - Task dispatching strategy: STA and DYN
    - Data transfer strategy: FAF (blocking) and FTF (pipeline)

	Static	Dynamic
First Activity First	STA_FAF	DYN_FAF
First Tuple First	STA_FTF	DYN_FTF

First Analysis

ACF = 1 / ISF = 1

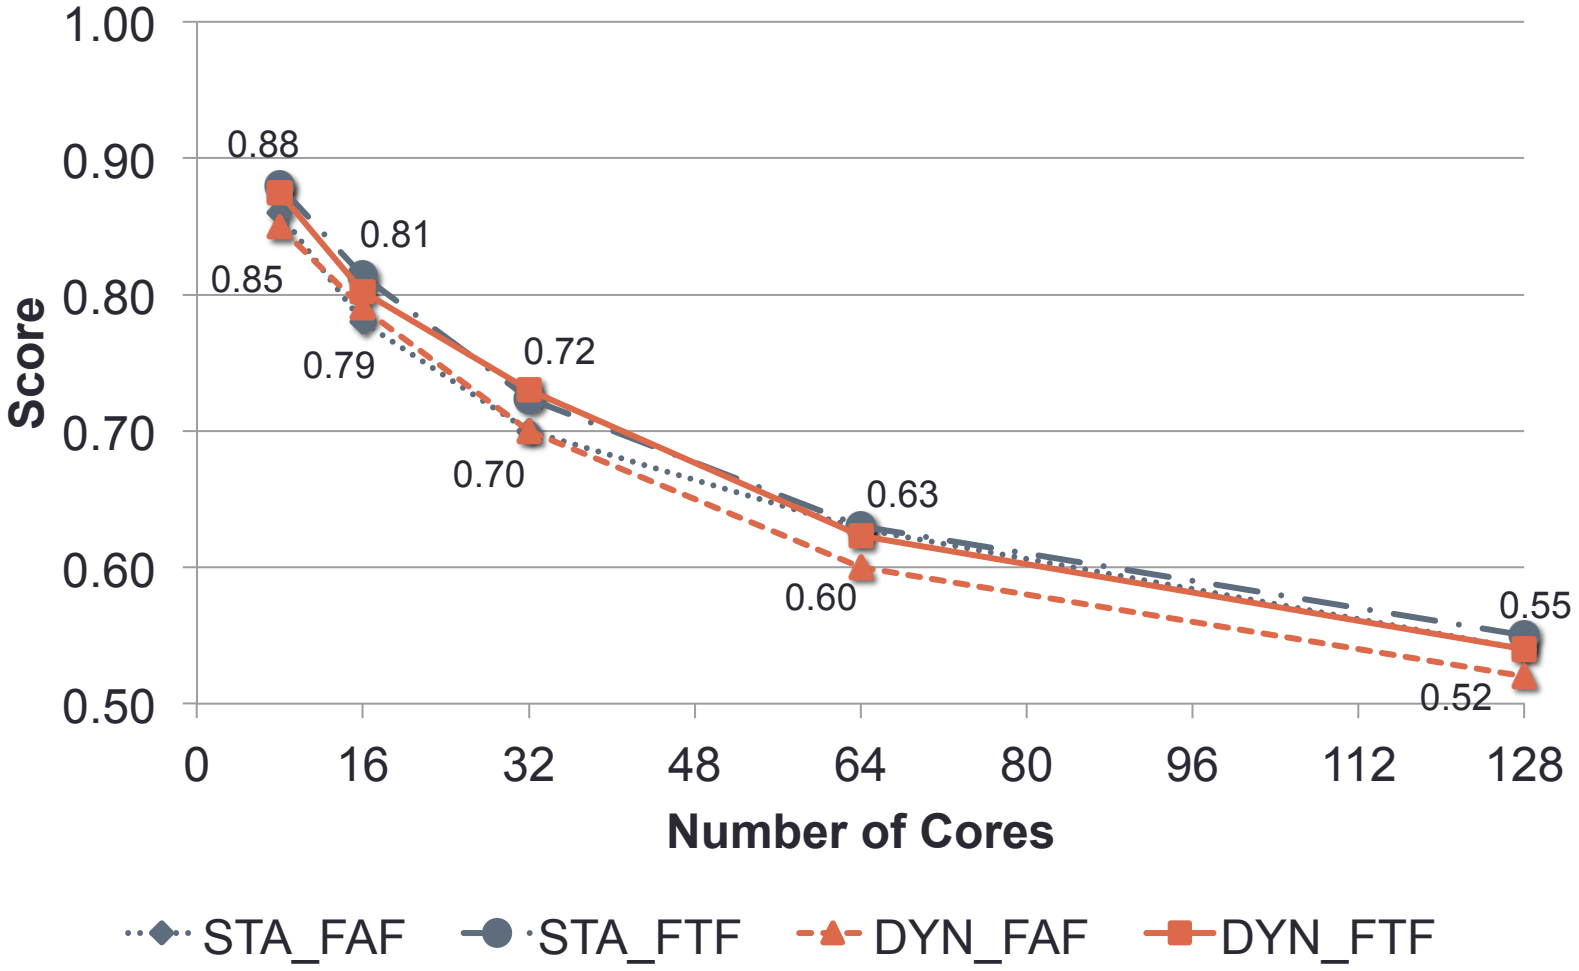

Second Analysis

ACF = 6 / ISF = 1

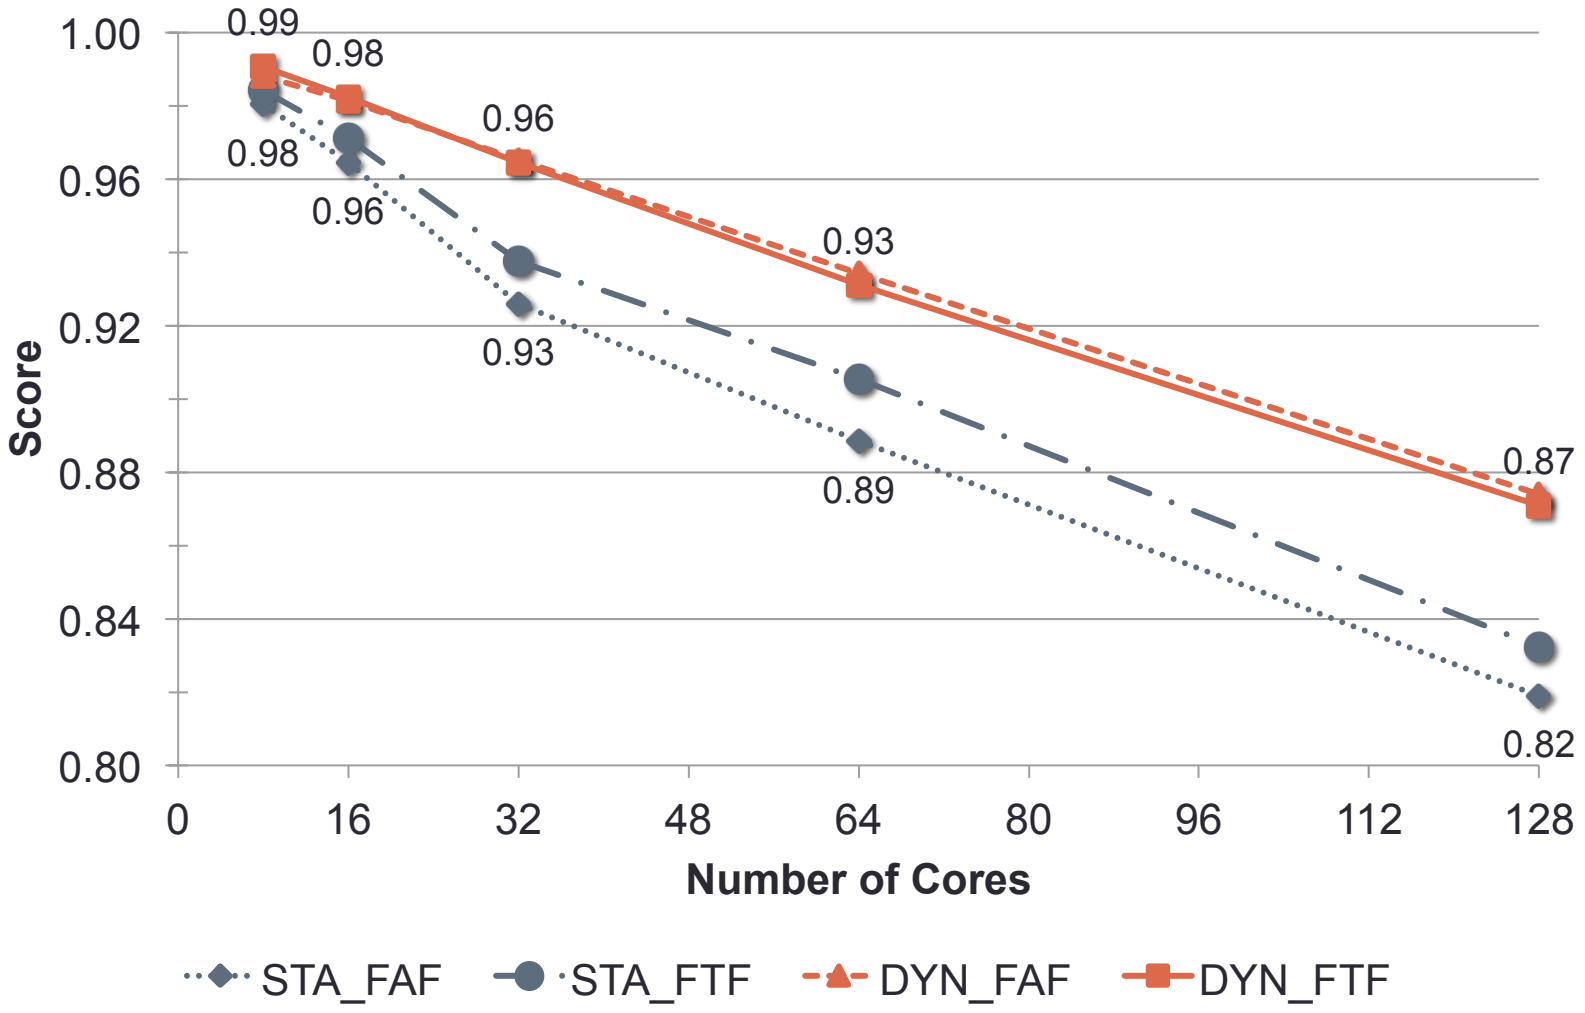

Third Analysis

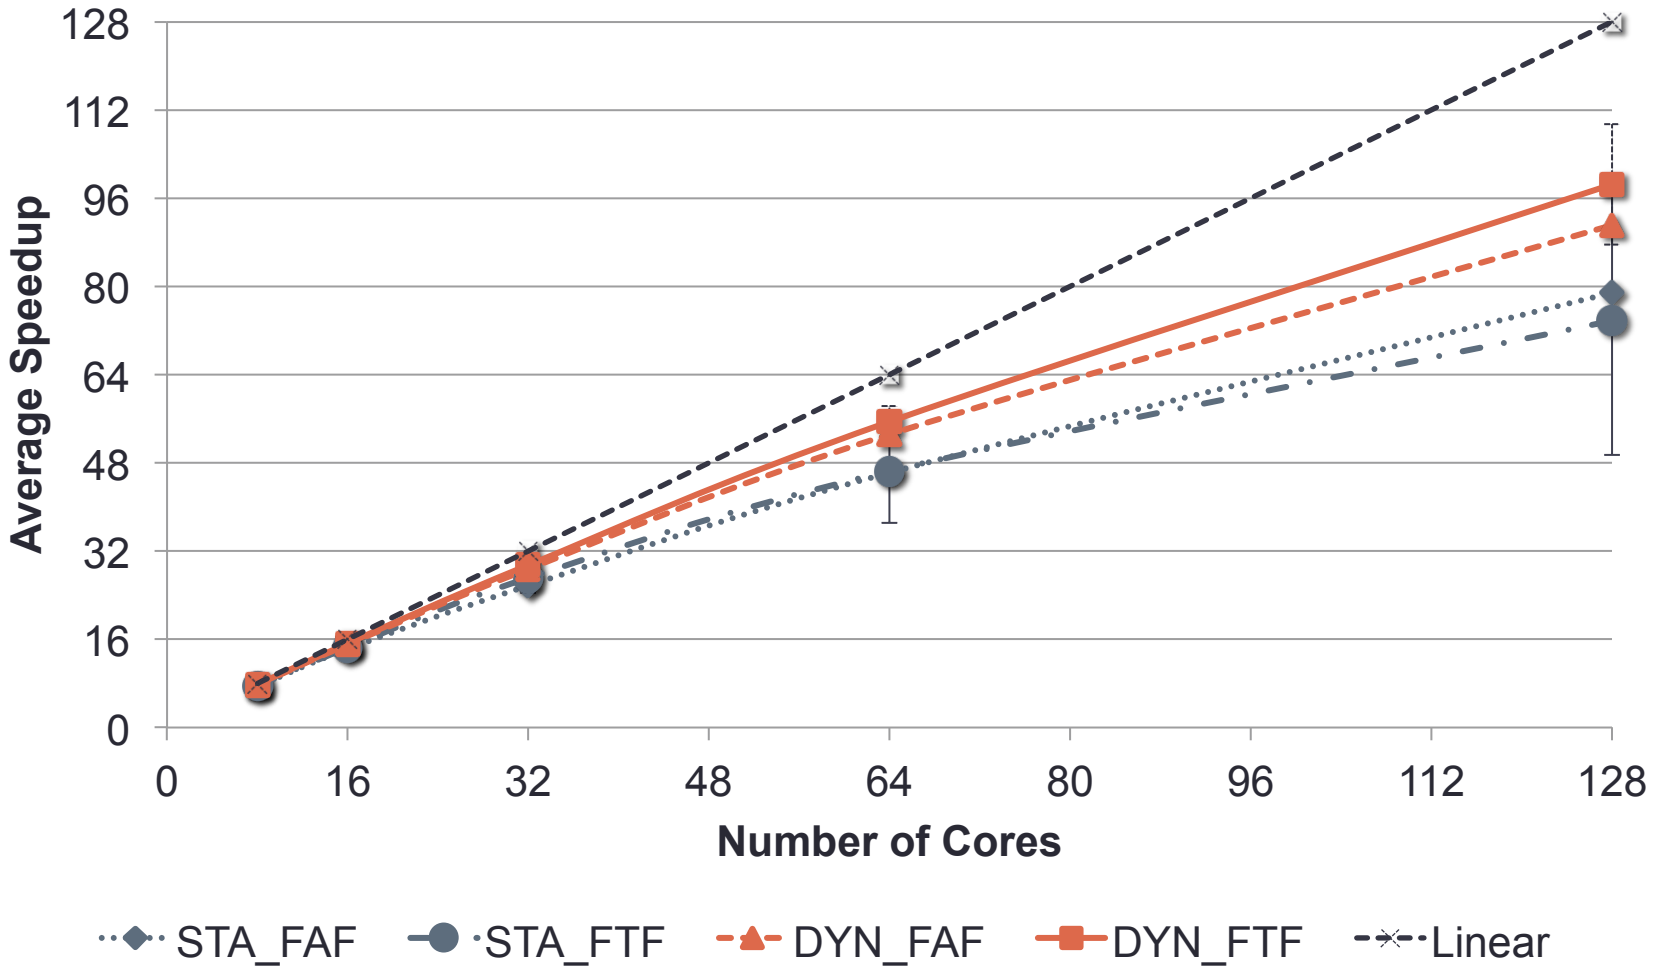

Fourth Analysis

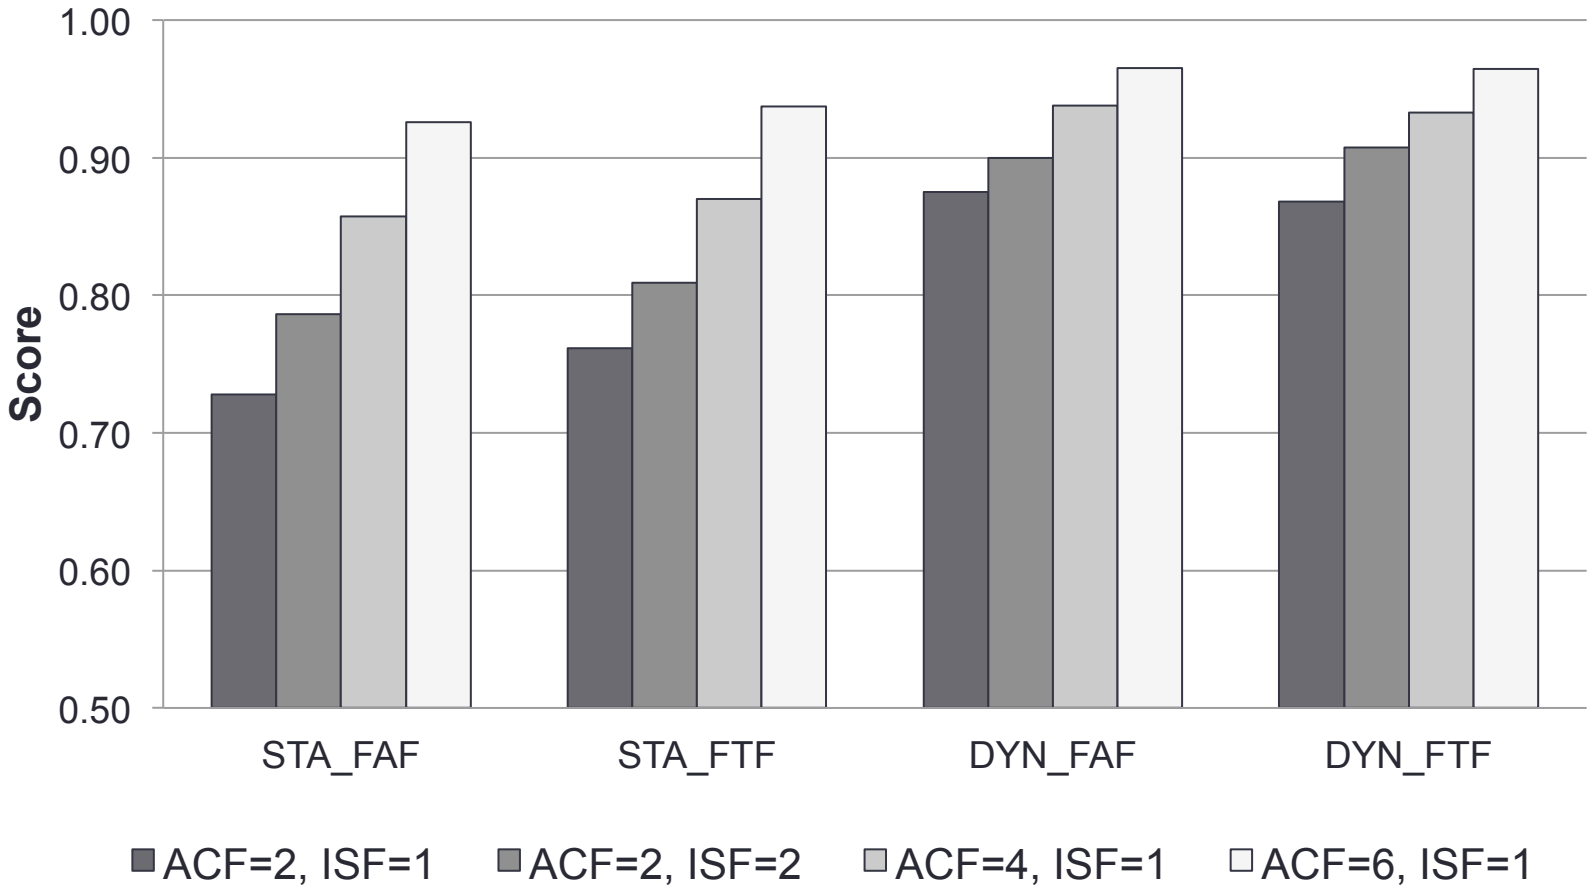

Related Work

Related Work

- Characterization of workflows
  - *Bharathi et al. 2008*
  - *Thain et al. 2003* and *Ostermann et al. 2008*
- Benchmarks
  - *Gillman et al. 2000*
  - *Goderis et al. 2005*
  - Yahoo! Cloud Serving Benchmark

Conclusion and Current Work

Conclusion

- Performance comparison for PS workflows
  - Characterization of PS Workflows
  - Evaluation Framework
    - PS workflow patterns • workload configuration • performance metrics
- Experiments
  - Four different execution models
  - Four analysis
- Step towards a benchmark

Current Work

- Optimizer for workflow engines
  - Algebraic approach
    - Based on relational algebra
  - More details
    - “An Algebraic Approach for Data-Centric Scientific Workflows” VLDB 2011

1st International Workshop on **S**calable **W**orkflow **E**nactment **E**ngines and **T**echnologies

SWEET'12

EVALUATING PARAMETER SWEEP WORKFLOWS IN HIGH PERFORMANCE COMPUTING

Fernando Chirigati*, Vítor Silva, Eduardo Ogasawara, Daniel de Oliveira, Jonas Dias, Fábio Porto, Patrick Valduriez and Marta Mattoso

Thank you!

COPPE / Federal University of Rio de Janeiro, Rio de Janeiro, Brazil

LNCC, Rio de Janeiro, Brazil
INRIA & LIRMM, Montpellier, France

* Currently at Polytechnic Institute of NYU

Workload Configuration

- Instance Scale Factor (ISF)
  - 2^{ISF+8}, $ISF \geq 1$
- Activity Cost Factor (ACF)
  - Gamma distribution $\Gamma(\kappa, \theta)$, where $\kappa = 2^{ACF}$ and $\theta = 1$, $ACF \geq 1$

OrthoMCL

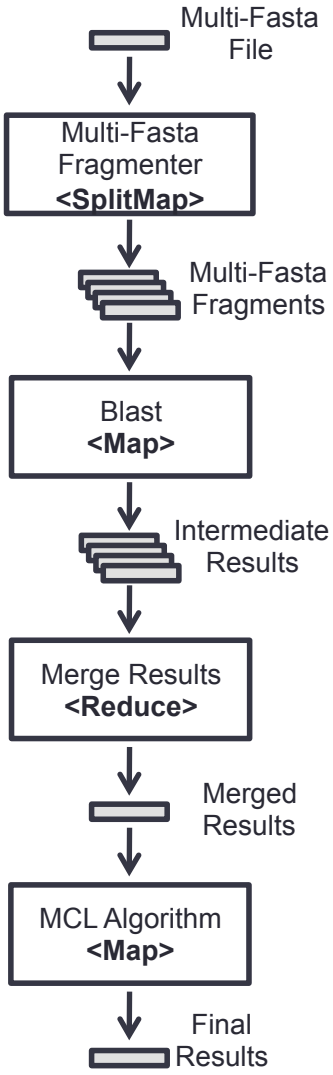

Risers Fatigue Analysis

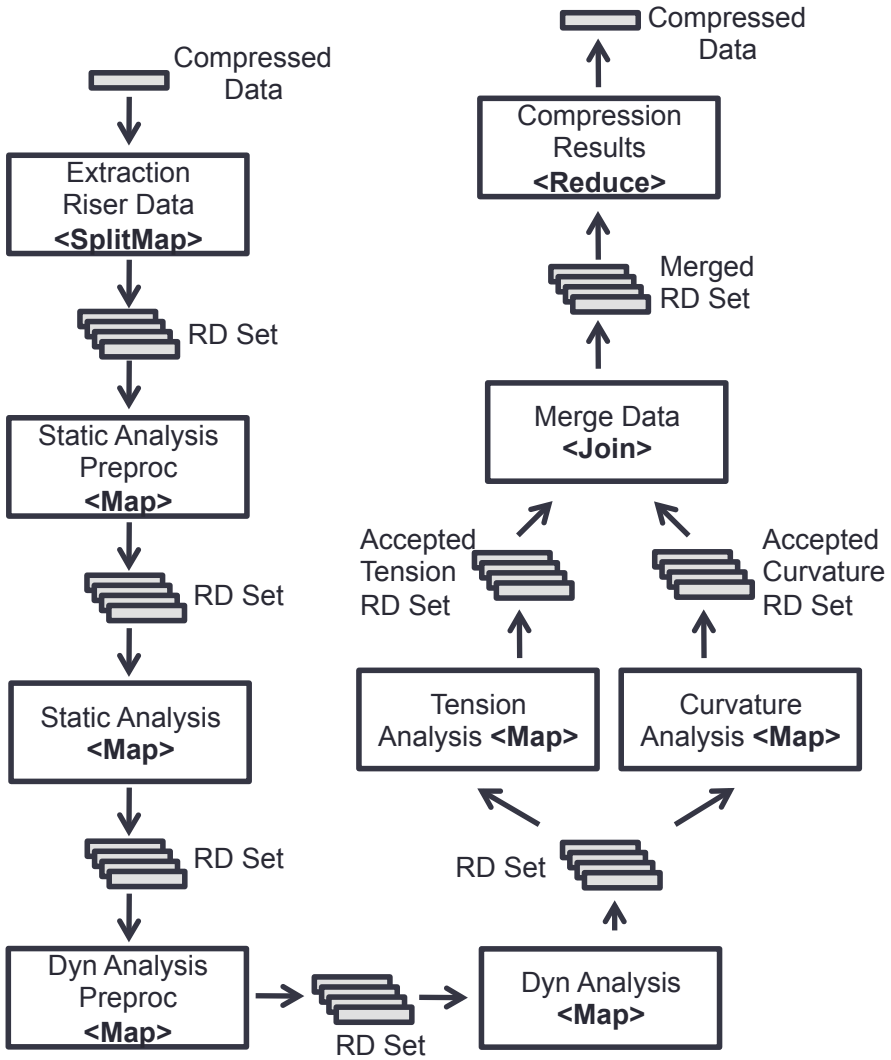

Map

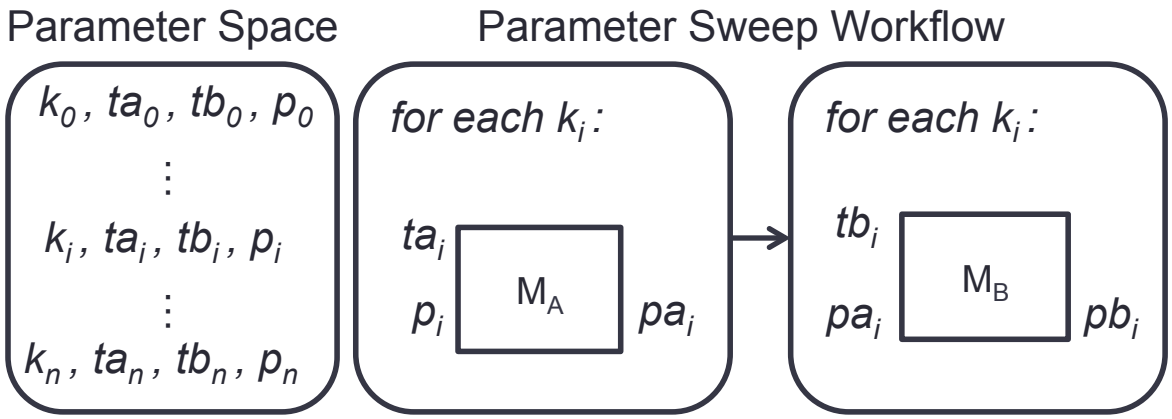
Parameter Space
Parameter Sweep Workflow
k_0, ta_0, tb_0, p_0
k_i, ta_i, tb_i, p_i
k_n, ta_n, tb_n, p_n
for each k_i :
ta_i
p_i
M_A
pa_i
for each k_i :
tb_i
pa_i
M_B
pb_i

Broadcast

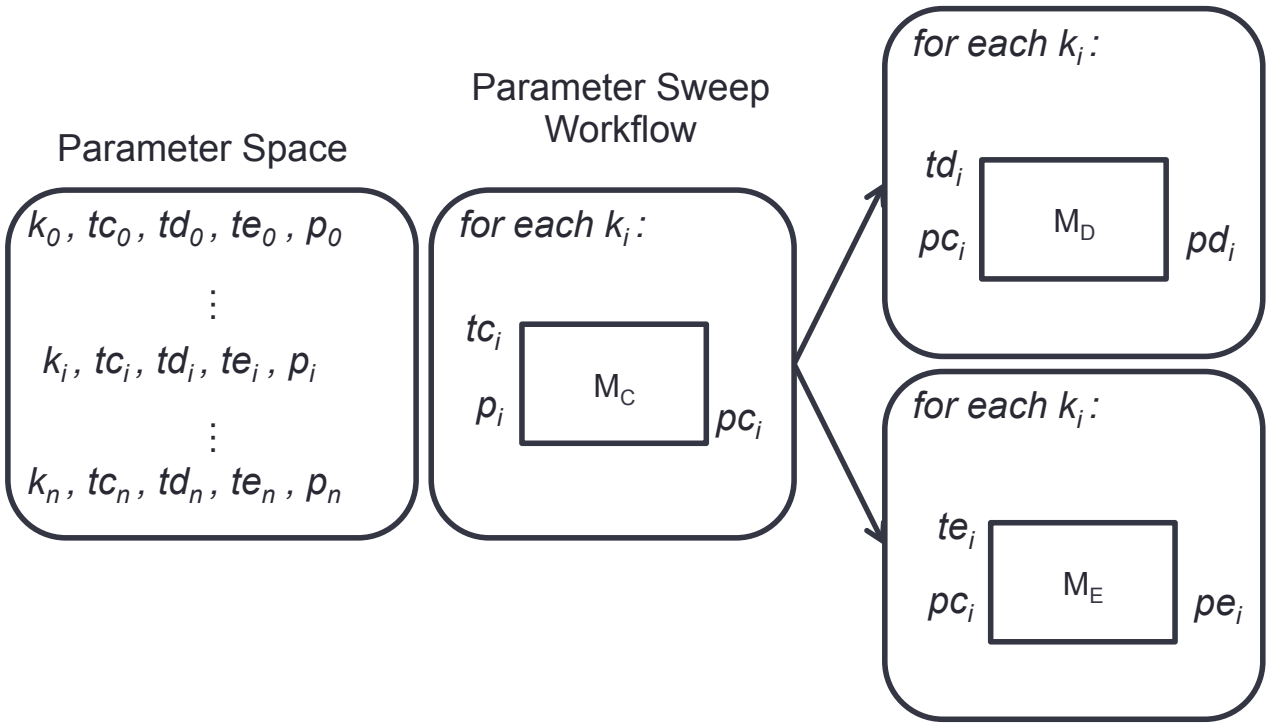
Parameter Space
$k_0, tc_0, td_0, te_0, p_0$
$k_i, tc_i, td_i, te_i, p_i$
$k_n, tc_n, td_n, te_n, p_n$
Parameter Sweep Workflow
for each k_i :
tc_i
p_i
M_C
pc_i
for each k_i :
td_i
pc_i
M_D
pd_i
for each k_i :
te_i
pc_i
M_E
pe_i

Split Map

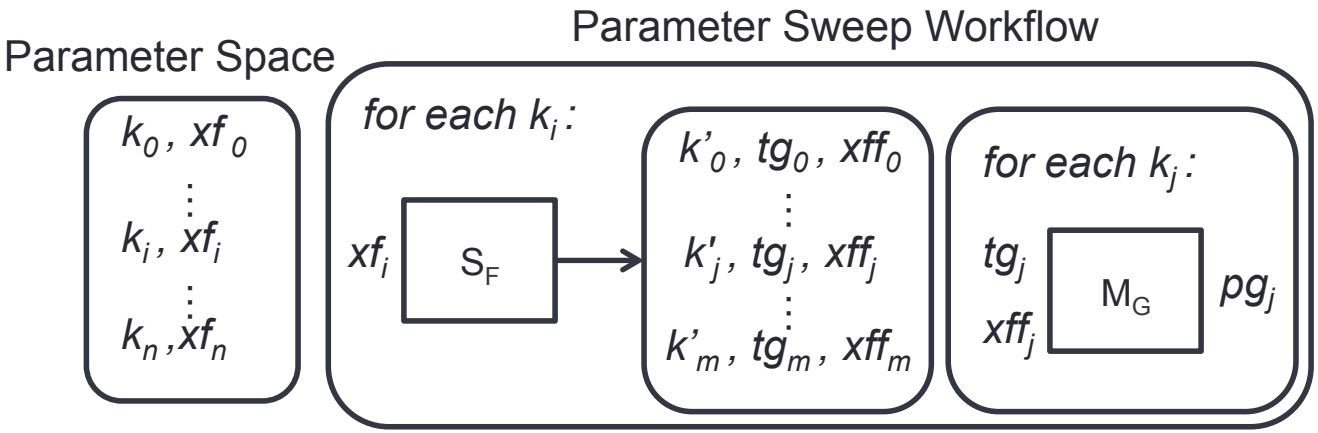
Parameter Space
Parameter Sweep Workflow
k_0, xf_0
k_i, xf_i
k_n, xf_n
for each k_i :
xf_i
S_F
k'_0, tg_0, xff_0
k'_j, tg_j, xff_j
k'_m, tg_m, xff_m
for each k_j :
tg_j
xff_j
M_G
pg_j

Join

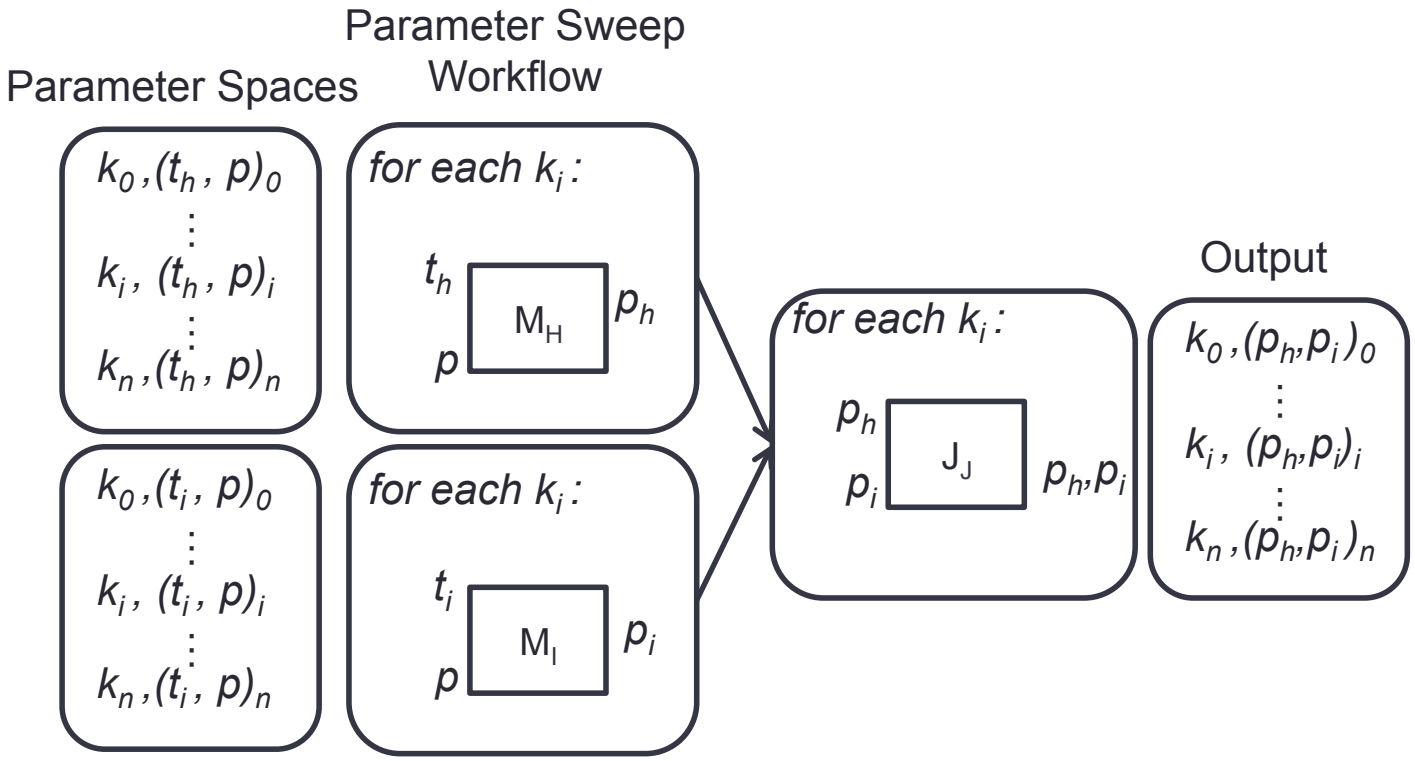
Parameter Sweep Workflow
Parameter Spaces
$k_0, (t_h, p)_0$
$k_i, (t_h, p)_i$
$k_n, (t_h, p)_n$
$k_0, (t_i, p)_0$
$k_i, (t_i, p)_i$
$k_n, (t_i, p)_n$
for each k_i :
t_h
p
M_H
p_h
for each k_i :
t_i
p
M_I
p_i
for each k_i :
p_h
p_i
J_J
p_h, p_i
Output
$k_0, (p_h, p_i)_0$
$k_i, (p_h, p_i)_i$
$k_n, (p_h, p_i)_n$

Reduce

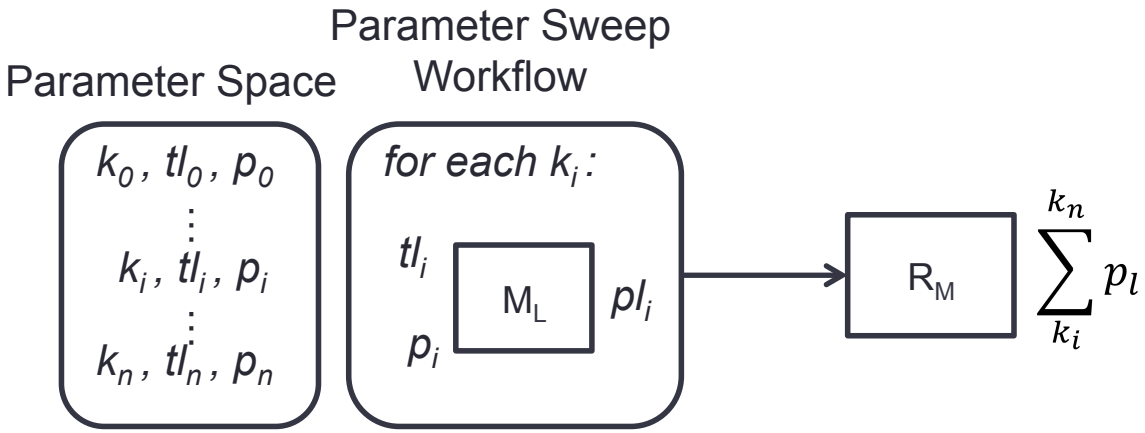
Parameter Space
Parameter Sweep Workflow
k_0, tl_0, p_0
k_i, tl_i, p_i
k_n, tl_n, p_n
for each k_i :
tl_i
p_i
M_L
pl_i
R_M
$\sum_{k_i}^{k_n} p_l$

Input Parameters

Map	k_i	t_a	t_b	p				
Broadcast	k_i	t_c	t_d	t_e	p			
Split Map	k_i	x_f						
Join	k_i	t_h	p			k_i	t_i	p
Reduce	k_i	t_l	p					

Map
ISF=2 / ACF=1

k_i	t_a	t_b	p
1	6413	1513	2420
2	4963	7011	9645
3	6670	3620	2956
⋮			
1024	4191	3083	1952

Programs

Name	Input	Output	Command line
M.jar	k, t, p_i	p_o	java –jar A.jar -K=k -T=t -P=p_i
S.jar	k, x_f	x_{ff}	java -jar S.jar -K=k -X=x_f
J.jar	j_a, j_b	j_o	java -jar J.jar -J=j_a,j_b
R.jar	r_i	v	java -jar R.jar -R=r_i