Supported by:

Federal Ministry for Economic Affairs and Energy

on the basis of a decision by the German Bundestag

High Performance Computing with GAMS

M.Bussieck (GAMS), F. Fiand (GAMS)

October 2017

2017 INFORMS Annual Meeting – Houston, Texas October 22-25

A PROJECT BY

Outline

- Motivation & Project Overview
- Parallel Interior Point Solver PIPS-IPM
- GAMS/PIPS-IPM Solver Link
  - Model Annotation
  - Distributed Model Generation
- Computational Experiments
- Summary & Outlook

Motivation

ESM REMix

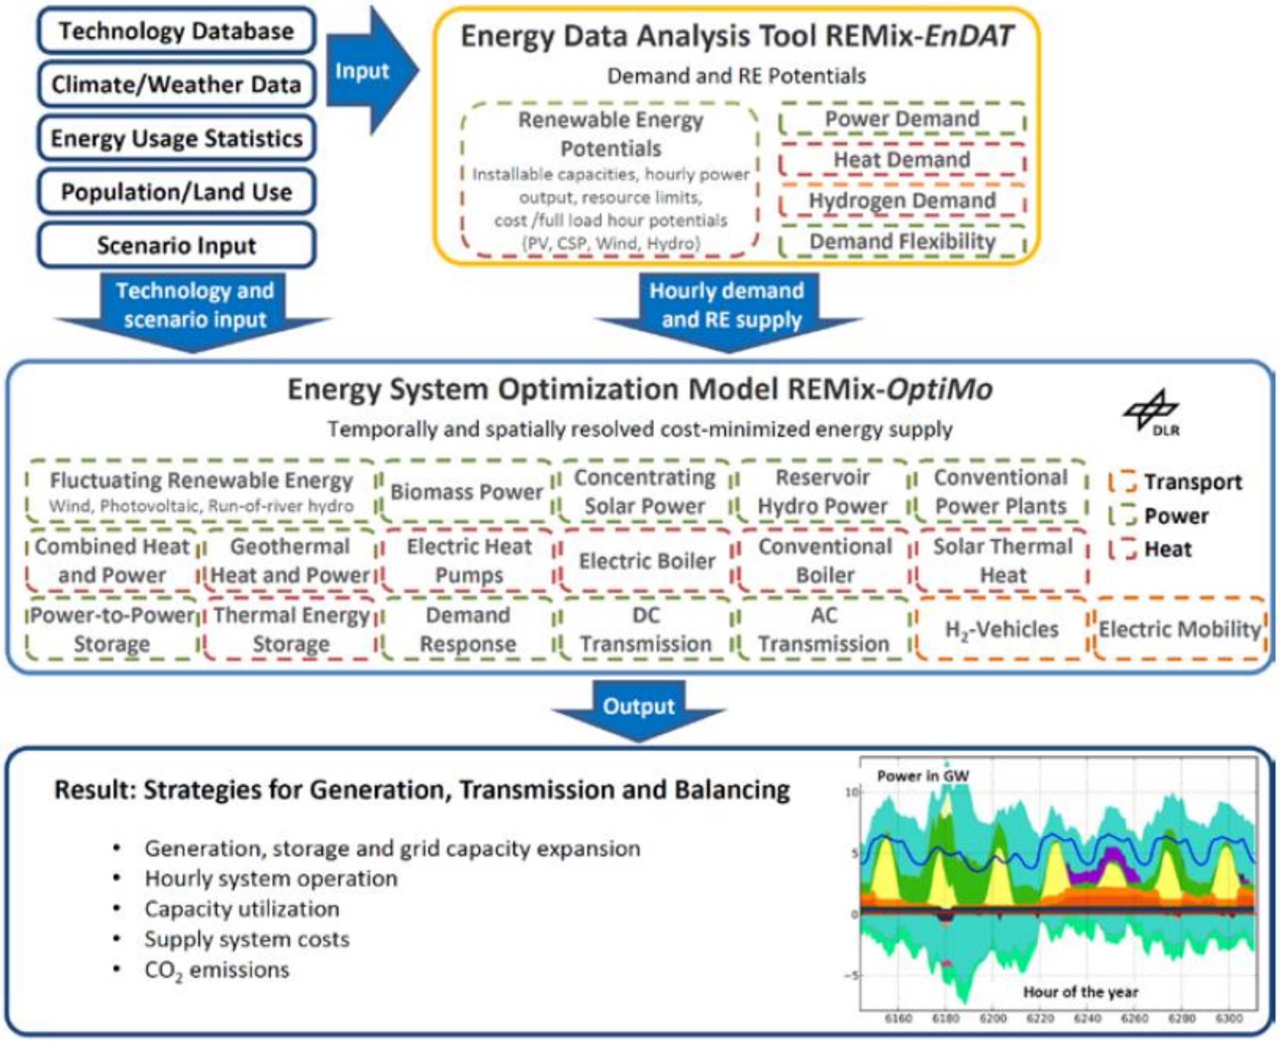

http://www.dlr.de/Portaldata/41/Resources/dokumente/institut/system/Modellbeschreibungen/DLR_Energy_System_Model_REMix_short_description_2016.pdf

Energy System Models

ESM REMix

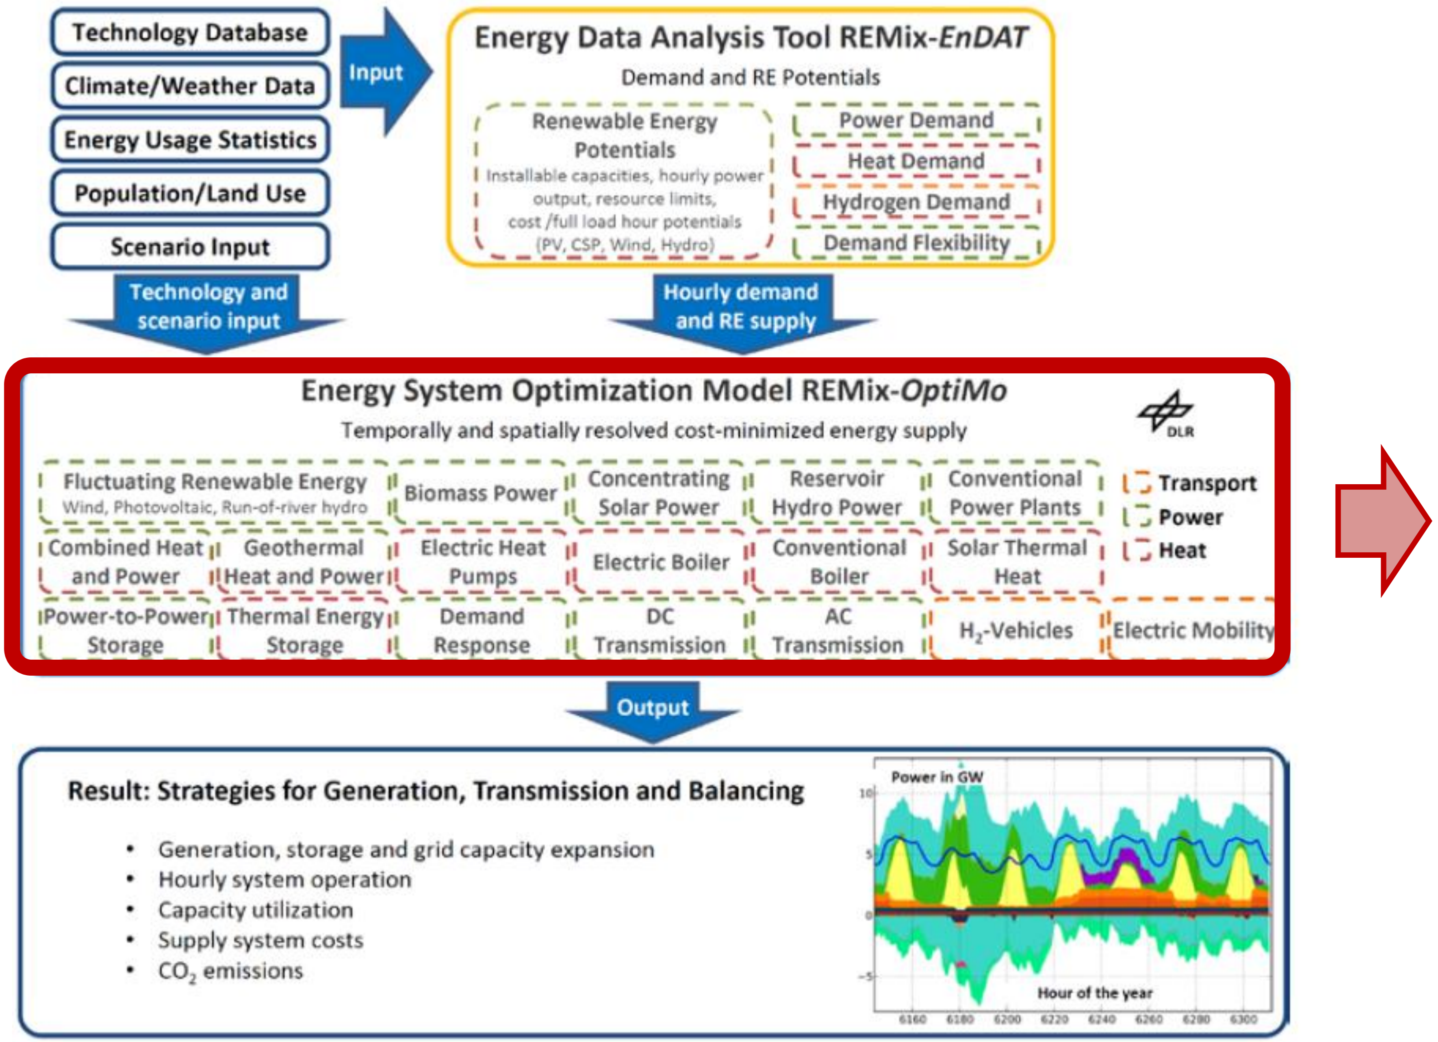

Results in large-scale LP

http://www.dlr.de/Portaldata/41/Resources/dokumente/institut/system/Modellbeschreibungen/DLR_Energy_System_Model_REMix_short_description_2016.pdf

Motivation

- Energy system models (ESM) have to increase in complexity to provide valuable quantitative insights for policy makers and industry:
  - Uncertainty
  - Large shares of renewable energies
  - Complex underlying electricity systems

- **Challenge:**
  - Increasing complexity makes solving ESM more and more difficult

→ Need for new solution approaches

Motivation

BEAM-ME

- Energy system models (ESM) have to increase in complexity to provide valuable quantitative insights for policy makers and industry:
  - Uncertainty
  - Large shares of renewable energies
  - Complex underlying el

- **Challenge:**
  - Increasing complexit
    and more difficult

→ Need for new solution approaches

Challenge appears in several areas.
ESM is just one potential field of application.

What exactly is BEAM-ME about?

Realisierung von Beschleunigungsstrategien der anwendungsorientierten Mathematik und Informatik für optimierende Energiesystemmodelle

Implementation of acceleration strategies from mathematics and computational sciences for optimizing energy system models

The BEAM-ME Project cont.

BEAM-ME

Energy System Modeling

DLR
Deutsches Zentrum für Luft- und Raumfahrt
German Aerospace Center

High Performance Computing

HLRS

JÜLICH
FORSCHUNGSZENTRUM

ZIB

TU berlin

Solver Development

GAMS

Modeling Language

Parallel Interior Point Solver PIPS-IPM

PIPS-IPM

PIPS-IPM: Parallel interior-point solver for LPs (und QPs) from stochastic energy models.

Main developer: Cosmin Petra (Argonne National Laboratory, Lawrence Livermore National Laboratory)

- PIPS-IPM is Open-Source.
- PIPS-IPM already solved problems with more than 10^9 variables.[1]
- PIPS-IPM originally supported linking variables but no linking constraints.
- PIPS-IPM extension to support linking constraints implemented by ZIB.

[1] Petra et al. 2014: *"Real-Time Stochastic Optimization of Complex Energy Systems on High-Performance Computers"*

Slide provided by D. Rehfeldt (ZIB)

All linear programs can be transformed into standard form

$$
\begin{aligned}
\min \quad & c^T x \\
\text{s.t.} \quad & Ax = b \\
& x \geq 0
\end{aligned}
$$

with $c \in \mathbb{R}^n, b \in \mathbb{R}^m, A \in \mathbb{R}^{m \times n}$.

Slide provided by D. Rehfeldt (ZIB)

Interior-Point vs. Simplex Algorithm

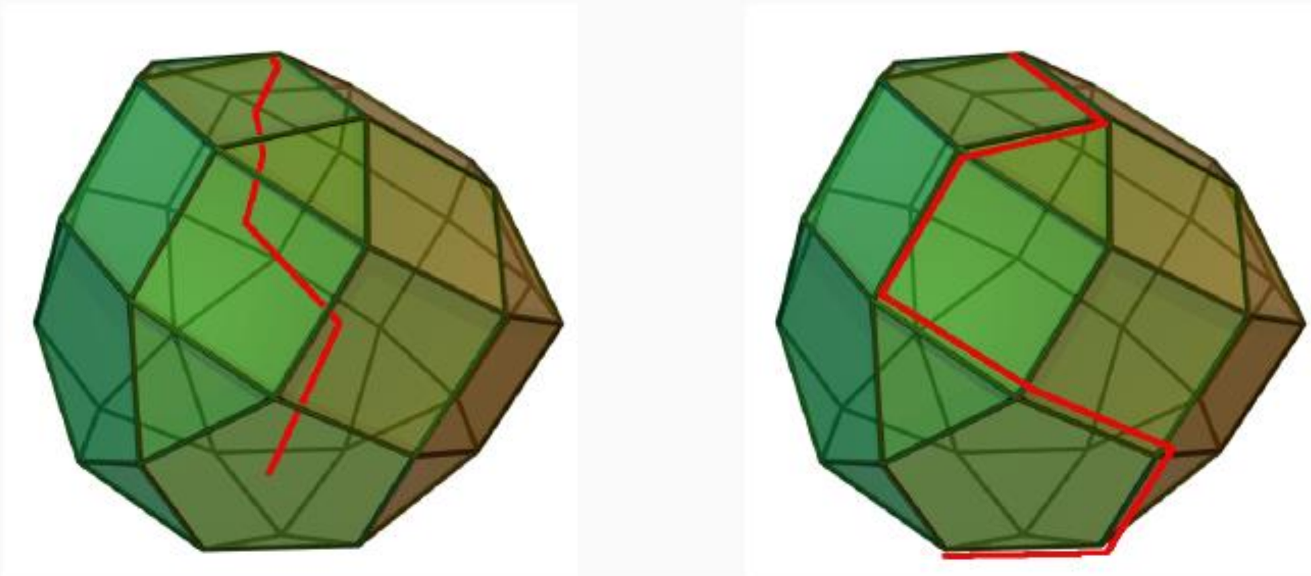

Picture: Wikipedia

- Choice of algorithm depends on problem to be solved
- But: interior-point usually faster for large-scale problems
- ...also for BEAM-ME LPs, see [Cao, Gleixner, Miltenberger, 16']

Slide provided by D. Rehfeldt (ZIB)

Primal Dual Interior Point Method cont.

Sketch: towards optimal solution along Central Path

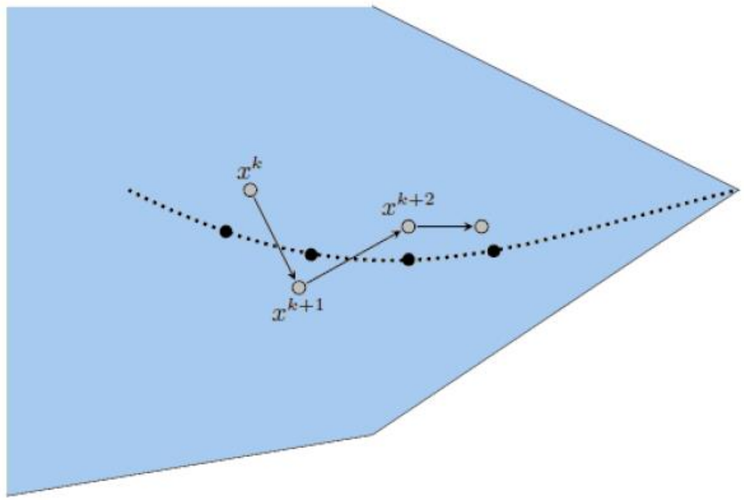

Two crucial points for practical solving:

1. Choice of direction and step length
2. **Solving system of linear equations**

Idea for BEAM-ME LPs: Exploit **block structure** to solve systems of linear equations in parallel

Slide provided by D. Rehfeldt (ZIB)

PIPS-IPM

Consider LP with block-diagonal structure, linking constraints, and linking variables (the kind of problem we want to solve):

$$
\begin{array}{cccccccc}
 & \min & c_0 & c_{s1} & & c_{sn} & & \\
 & & A^{=} & & & & = & b \\
b^{low} & \leq & A^{\leq} & & & & \leq & b^{upp} \\
 & & T_{s1}^{=} & W_{s1}^{=} & & & = & b_{s1} \\
b_{s1}^{low} & \leq & T_{s1}^{\leq} & W_{s1}^{\leq} & & & \leq & b_{s1}^{upp} \\
\vdots & & \vdots & & \ddots & & & \vdots \\
 & & T_{sn}^{=} & & & W_{sn}^{=} & = & b_{sn} \\
b_{sn}^{low} & \leq & T_{sn}^{\leq} & & & W_{sn}^{\leq} & \leq & b_{sn}^{upp} \\
 & & C^{=} & C_{s1}^{=} & \cdots & C_{sn}^{=} & = & b_C \\
b_C^{low} & \leq & C^{\leq} & C_{s1}^{\leq} & \cdots & C_{sn}^{\leq} & \leq & b_C^{upp} \\
x_0^{low} & \leq & x_0 & & & & \leq & x_0^{upp} \\
 & & & & \ddots & & & \\
x_{sn}^{low} & \leq & & & & x_{sn} & \leq & x_{sn}^{upp}
\end{array}
$$

PIPS-IPM

BEAM-ME

Consider LP with block-diagonal structure, linking constraints, and linking variables (the kind of problem we want to solve):

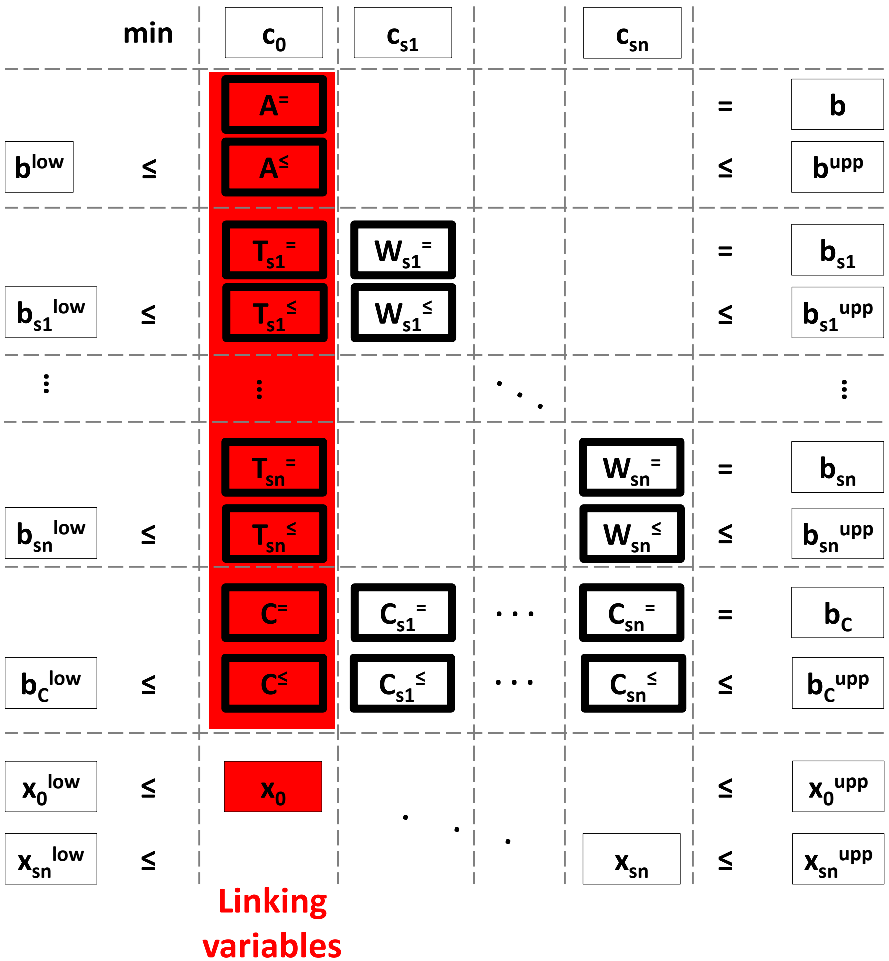

PIPS-IPM

Consider LP with block-diagonal structure, linking constraints, and linking variables (the kind of problem we want to solve):

	min	c_0	c_{s1}		c_{sn}			
		$A^=$				=	b	
b^{low}	≤	$A^\leq$				≤	b^{upp}	
		$T_{s1}^=$	$W_{s1}^=$			=	b_{s1}	
b_{s1}^{low}	≤	$T_{s1}^\leq$	$W_{s1}^\leq$			≤	b_{s1}^{upp}	
$\vdots$		$\vdots$		$\ddots$			$\vdots$	
		$T_{sn}^=$			$W_{sn}^=$	=	b_{sn}	
b_{sn}^{low}	≤	$T_{sn}^\leq$			$W_{sn}^\leq$	≤	b_{sn}^{upp}	
		$C^=$	$C_{s1}^=$	$\cdots$	$C_{sn}^=$	=	b_C	Linking
b_C^{low}	≤	$C^\leq$	$C_{s1}^\leq$	$\cdots$	$C_{sn}^\leq$	≤	b_C^{upp}	constraints
x_0^{low}	≤	x_0				≤	x_0^{upp}	
x_{sn}^{low}	≤			$\ddots$	x_{sn}	≤	x_{sn}^{upp}	

PIPS-IPM

Consider LP with block-diagonal structure, linking constraints, and linking variables (the kind of problem we want to solve):

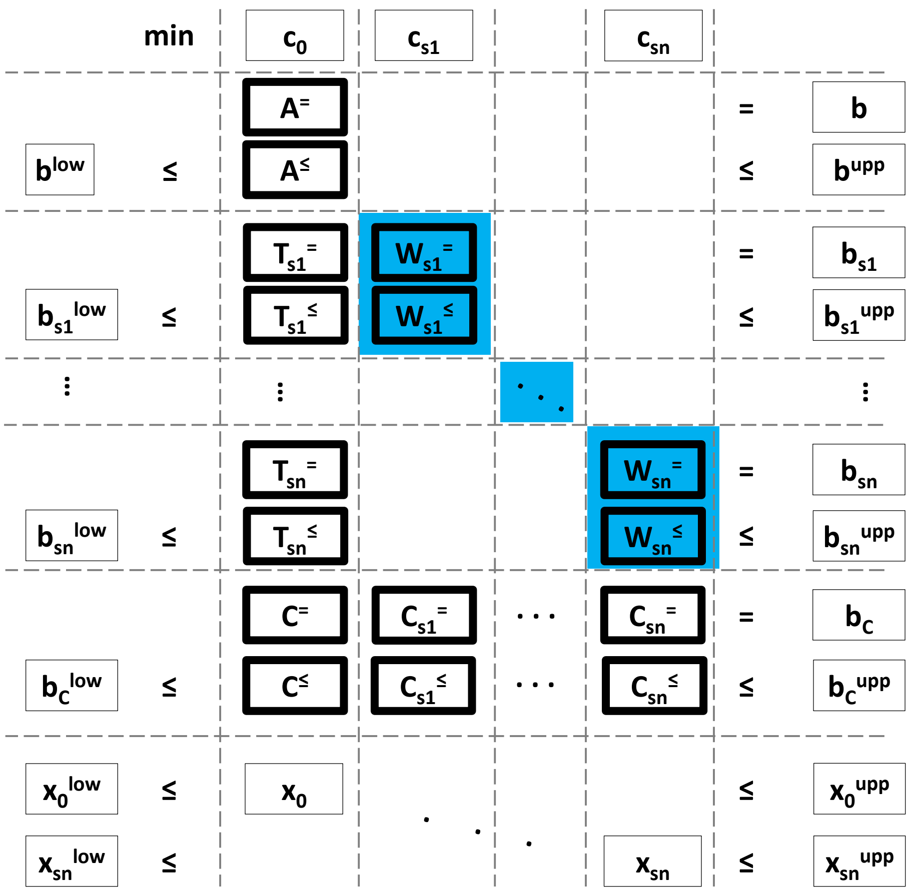

Recourse decision blocks

PIPS-IPM

Consider LP with block-diagonal structure, linking constraints, and linking variables (the kind of problem we want to solve):

	min	c_0	c_{s1}		c_{sn}		
		$A^=$				=	b
b^{low}	≤	$A^\leq$				≤	b^{upp}
		$T_{s1}^=$	$W_{s1}^=$			=	b_{s1}
b_{s1}^{low}	≤	$T_{s1}^\leq$	$W_{s1}^\leq$			≤	b_{s1}^{upp}
⋮		⋮		⋱			⋮
		$T_{sn}^=$			$W_{sn}^=$	=	b_{sn}
b_{sn}^{low}	≤	$T_{sn}^\leq$			$W_{sn}^\leq$	≤	b_{sn}^{upp}
		$C^=$	$C_{s1}^=$	…	$C_{sn}^=$	=	b_C
b_C^{low}	≤	$C^\leq$	$C_{s1}^\leq$	…	$C_{sn}^\leq$	≤	b_C^{upp}
x_0^{low}	≤	x_0				≤	x_0^{upp}
x_{sn}^{low}	≤		⋱		x_{sn}	≤	x_{sn}^{upp}

- Block diagonal structure allows parallelization of linear algebra within PIPS-IPM
- Solve *N* systems of linear equations in parallel instead of one huge system

PIPS-IPM

Consider LP with block-diagonal structure, linking constraints, and linking variables (the kind of problem we want to solve):

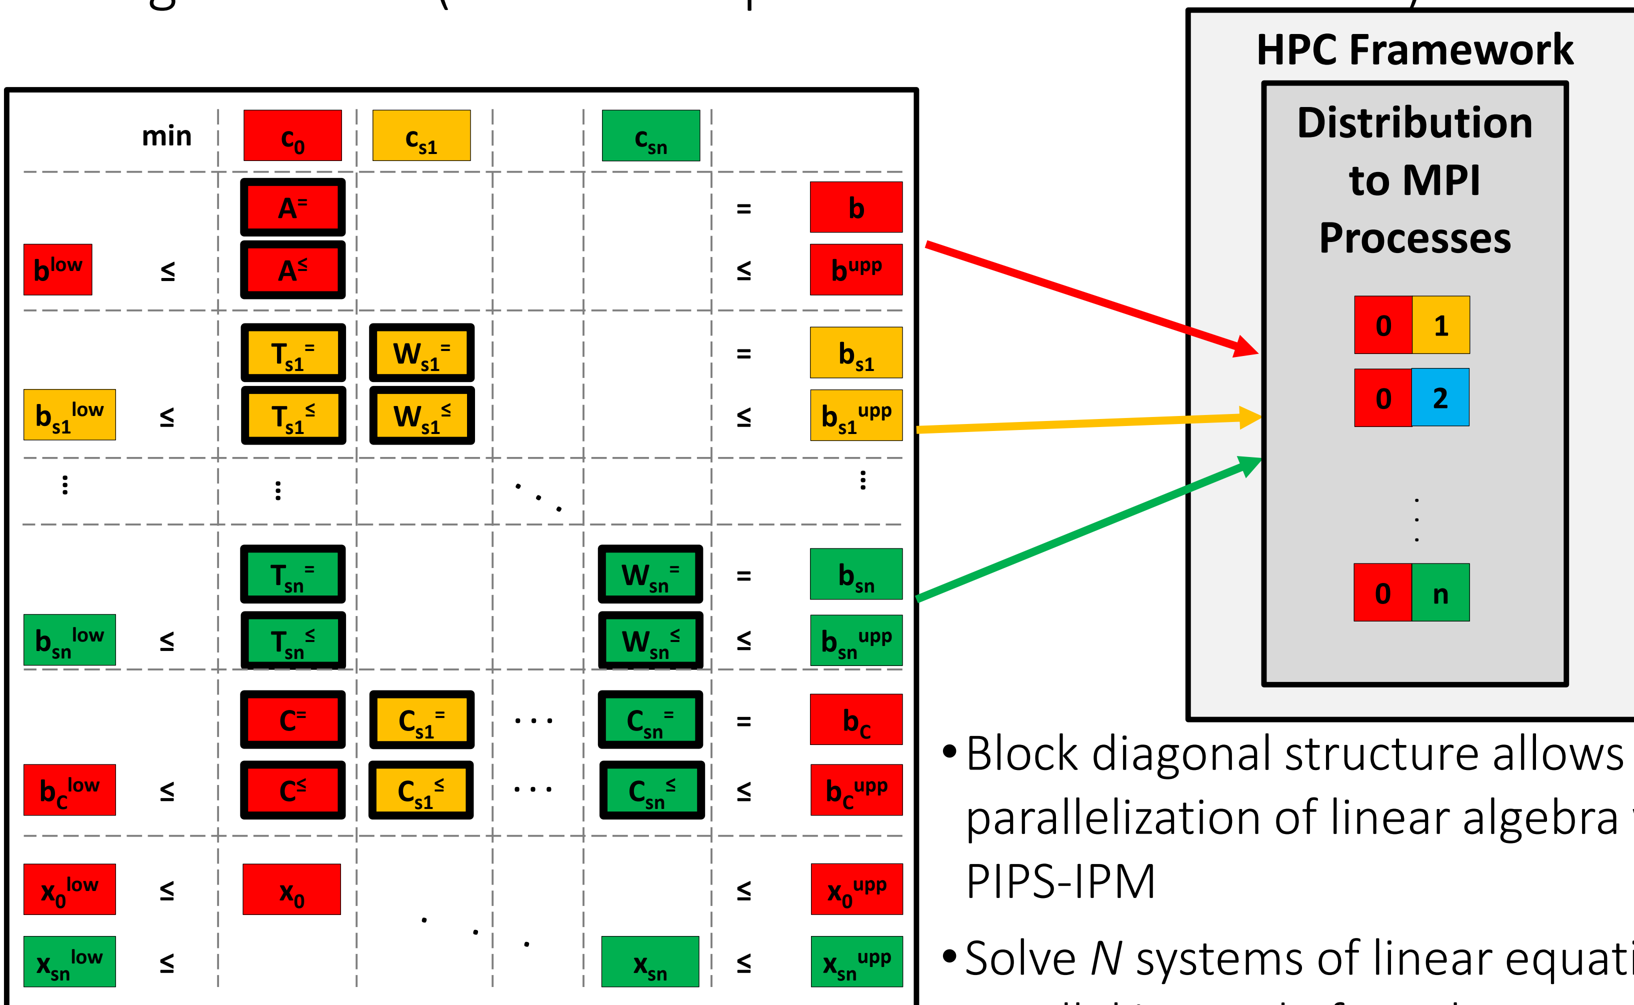

- Block diagonal structure allows parallelization of linear algebra within PIPS-IPM
- Solve *N* systems of linear equations in parallel instead of one huge system

GAMS/PIPS-IPM Solver Link

Model Annotation

GAMS/PIPS Solver Link - Overview

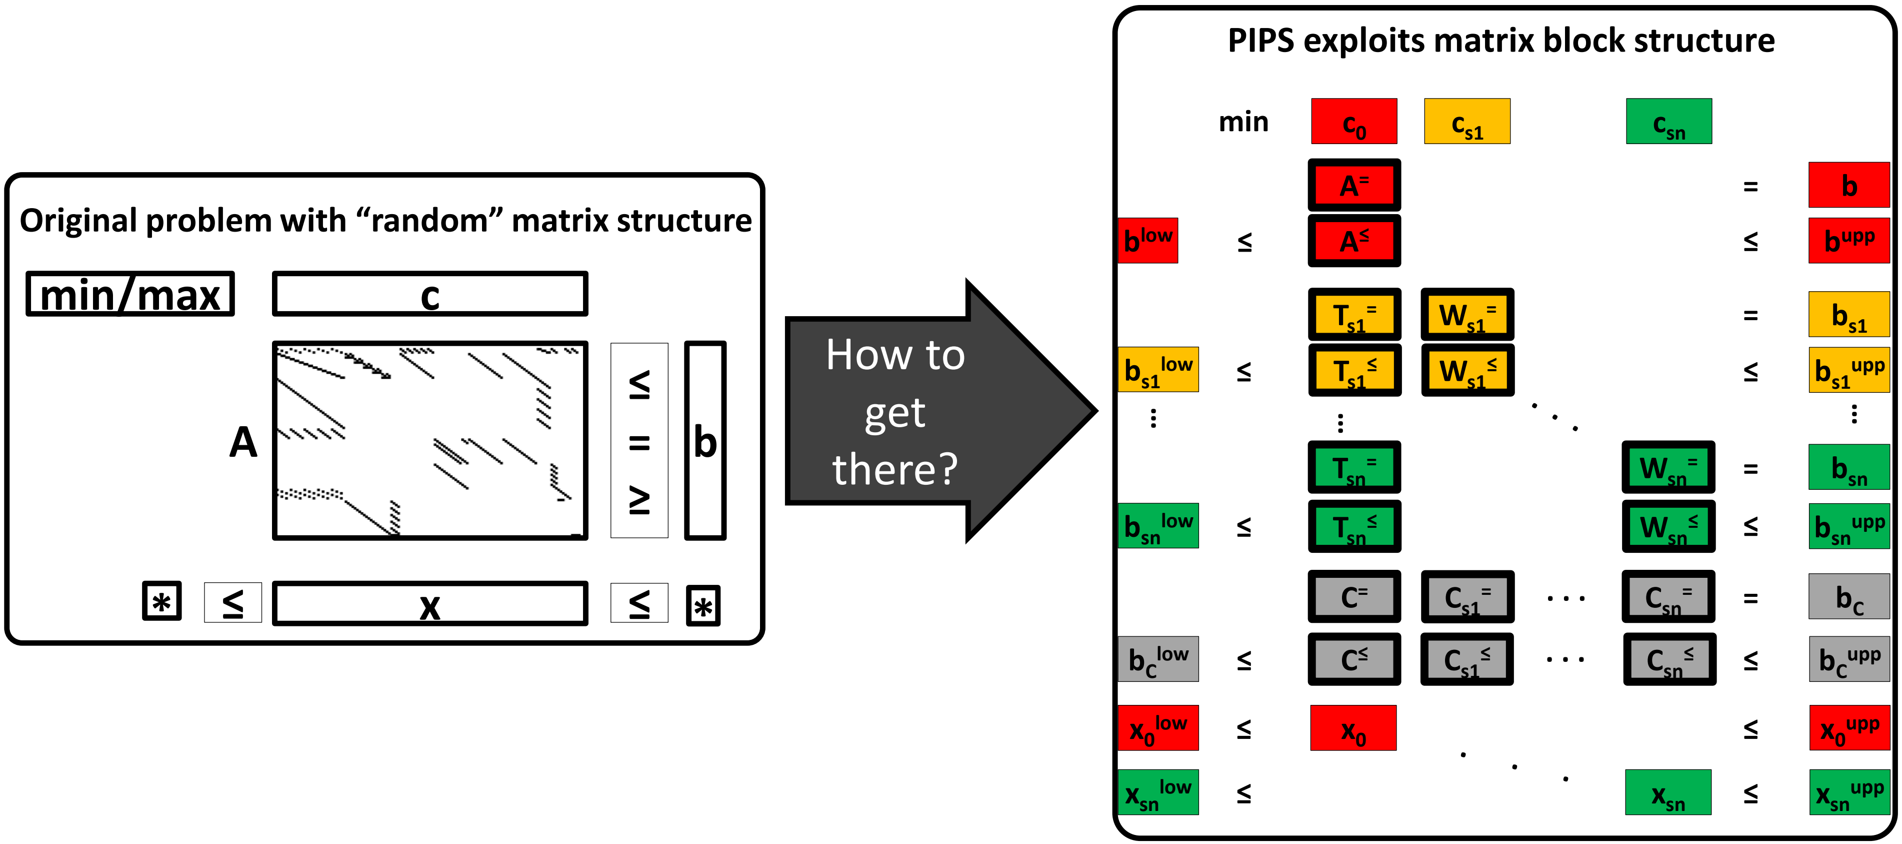

GAMS/PIPS Solver Link - Overview

Original problem with “random” matrix structure

min/max | c

A

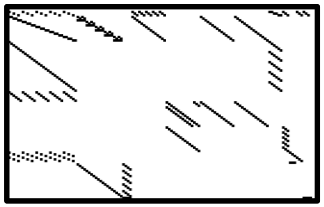

≤ / = / ≥ | b

* ≤ x ≤ *

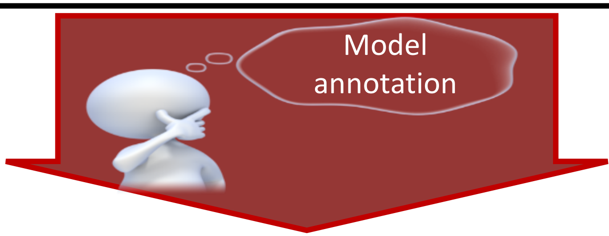

Permutation reveals block structure

min/max | c’

A’

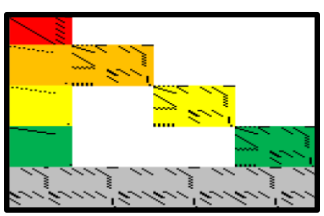

≤ / = / ≥ | b'

* ≤ x' ≤ *

PIPS exploits matrix block structure

min		c_0	c_{s1}		c_{sn}		
		$A^{=}$				=	b
b^{low}	$\le$	$A^{\le}$				$\le$	b^{upp}
		$T_{s1}^{=}$	$W_{s1}^{=}$			=	b_{s1}
b_{s1}^{low}	$\le$	$T_{s1}^{\le}$	$W_{s1}^{\le}$			$\le$	b_{s1}^{upp}
$\vdots$		$\vdots$		$\ddots$			$\vdots$
		$T_{sn}^{=}$			$W_{sn}^{=}$	=	b_{sn}
b_{sn}^{low}	$\le$	$T_{sn}^{\le}$			$W_{sn}^{\le}$	$\le$	b_{sn}^{upp}
		$C^{=}$	$C_{s1}^{=}$	$\cdots$	$C_{sn}^{=}$	=	b_C
b_C^{low}	$\le$	$C^{\le}$	$C_{s1}^{\le}$	$\cdots$	$C_{sn}^{\le}$	$\le$	b_C^{upp}
x_0^{low}	$\le$	x_0		$\ddots$		$\le$	x_0^{upp}
x_{sn}^{low}	$\le$				x_{sn}	$\le$	x_{sn}^{upp}

GAMS/PIPS Solver Link - Overview

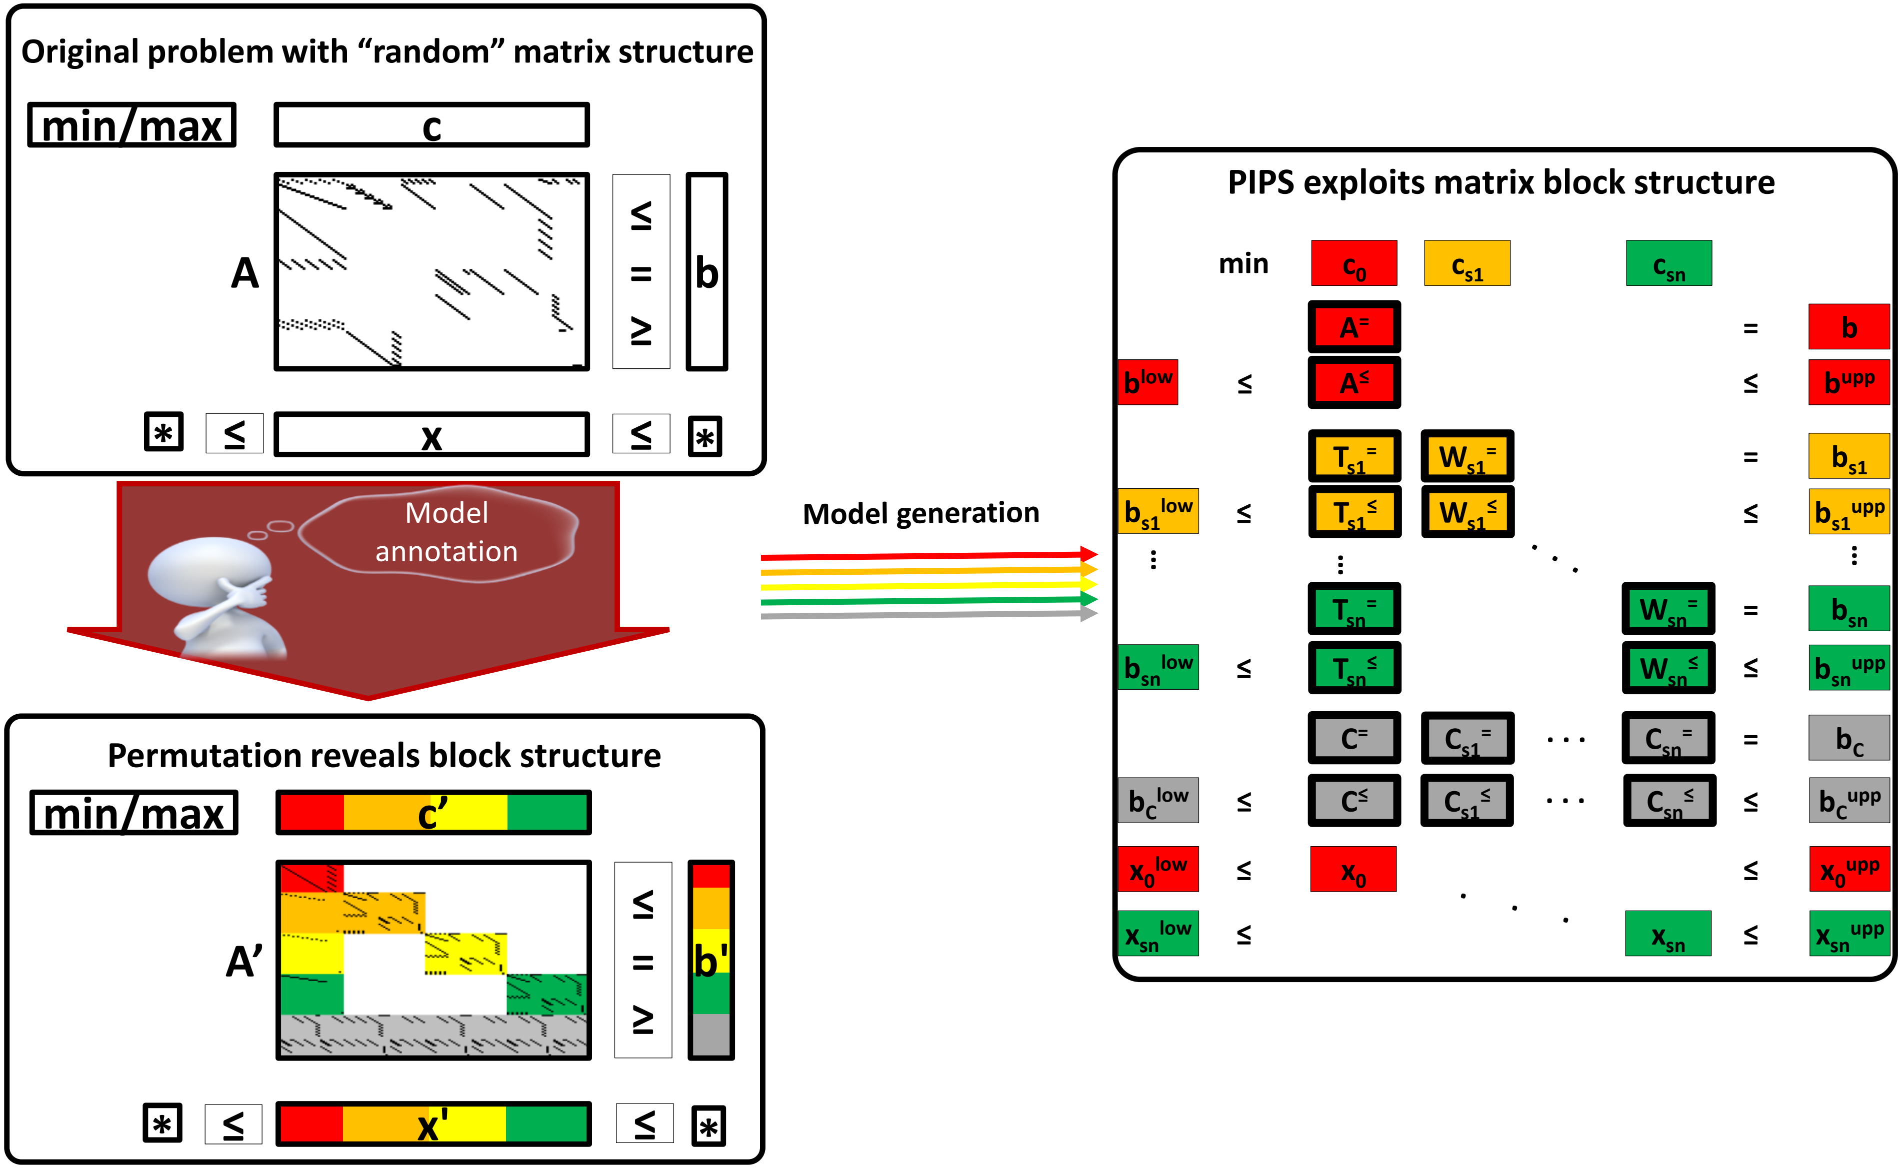

Model Annotation

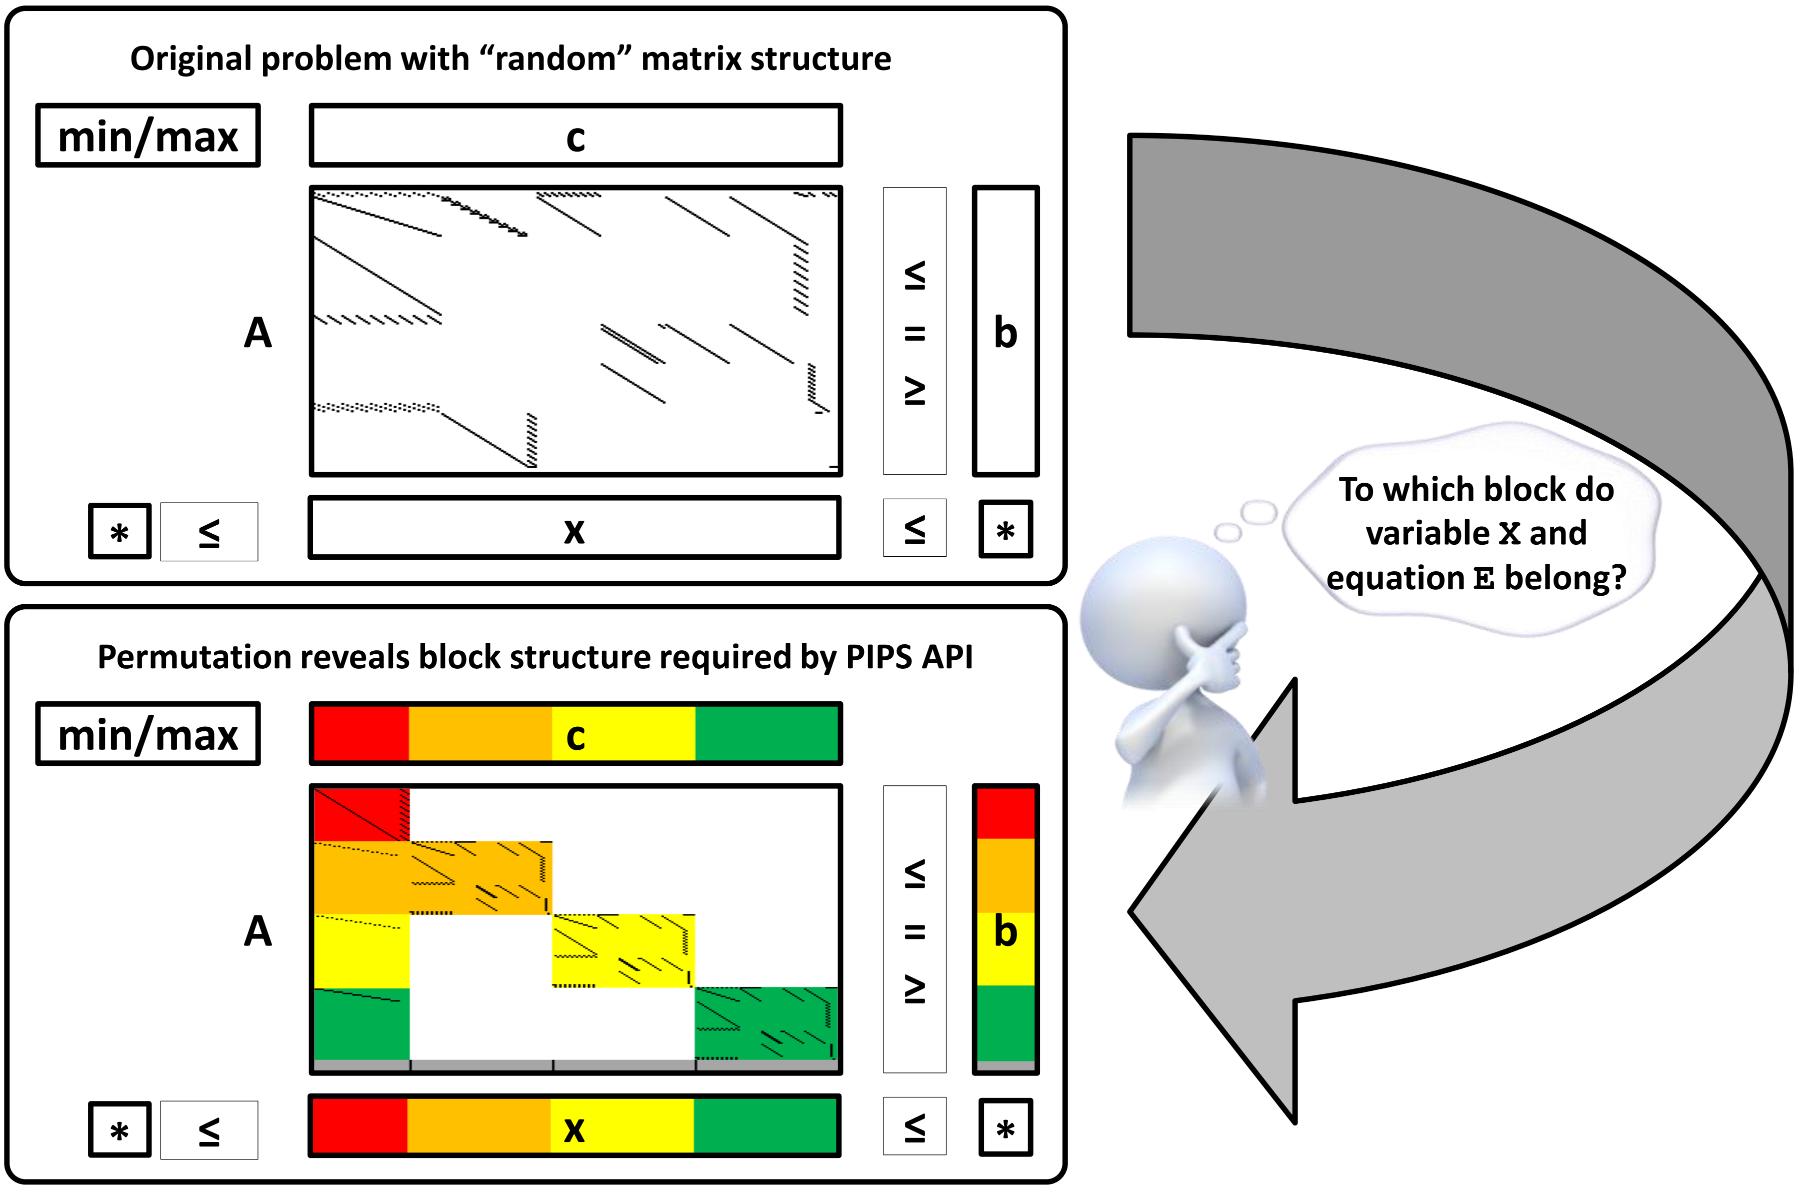

Model Annotation cont.

Model Annotation by .Stage

- The .stage attribute is available for variables/equations in GAMS

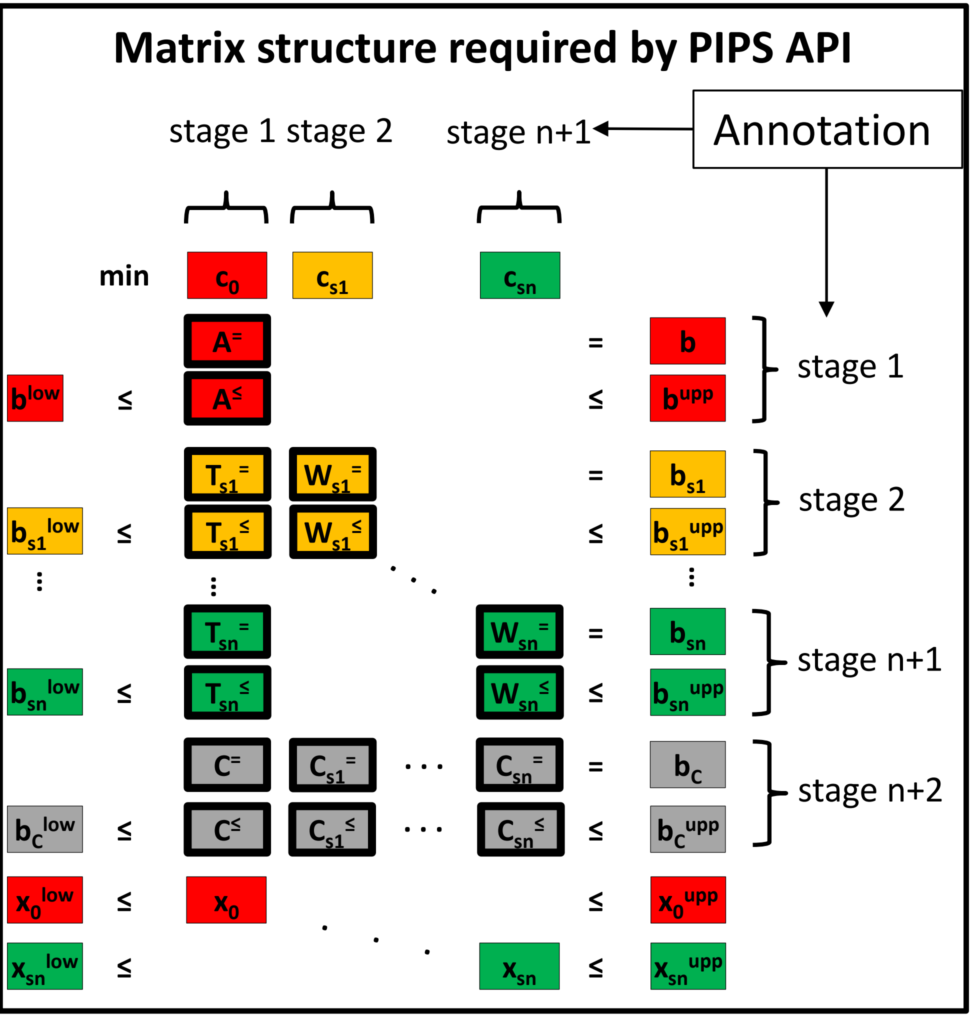

Model Annotation cont.

Model Annotation by .Stage

- The .stage attribute is available for variables/equations in GAMS

Matrix structure required by PIPS API

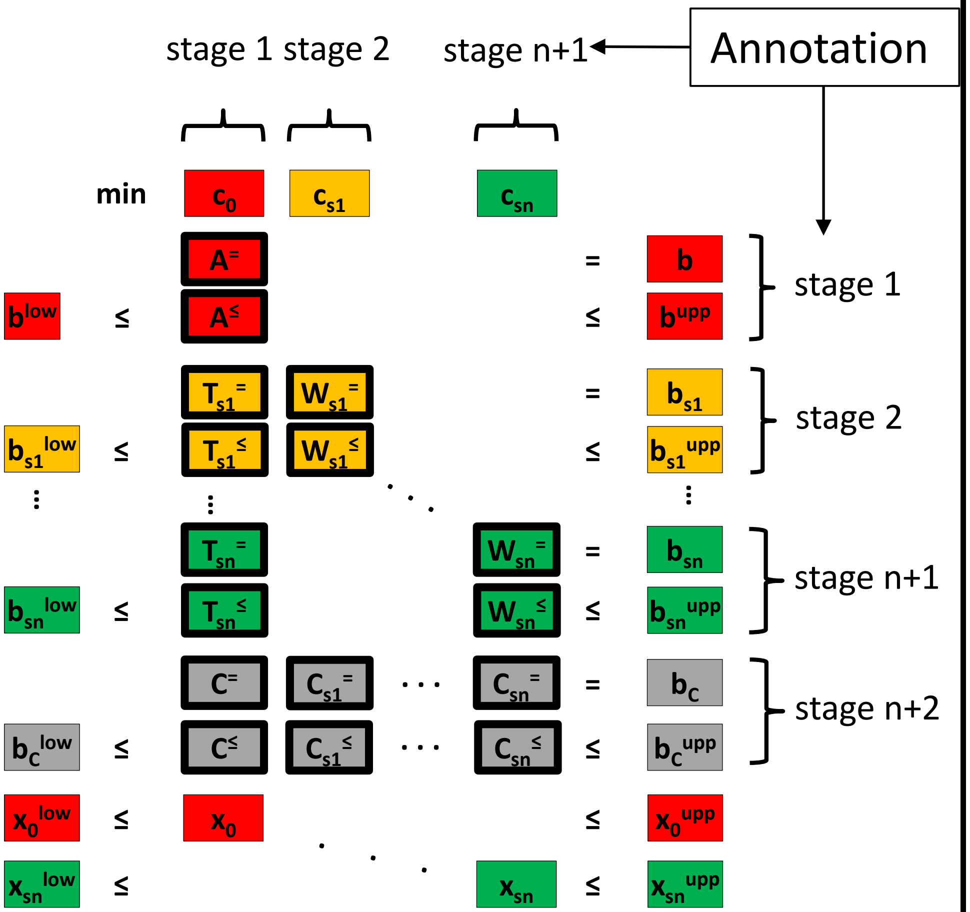

Exemplary Annotation for simple ESM (regional decomposition)

```
[...]
* Master variables and equation
FLOW.stage(t,net(rr1,rr2))        = 1;
LINK_ADD_CAP.stage(net(rr1,rr2)) = 1;
[...]
* Block variables and equations
ROBJ.stage(rr)              = ord(rr)+1;
POWER.stage(t,rp(rr,p))     = ord(rr)+1;
EMISSION_SPLIT.stage(rr,e)  = ord(rr)+1;
[...]
eq_emission_region.stage(rr,e)  = ord(rr)+1;
eq_emission_cost.stage(rr,e)    = ord(rr)+1;
[...]
* Linking Equation
eq_emission_cap.stage(e)  = n+2;
```

Model Annotation cont.

- How to annotate Model depends on how the model should be “decomposed” (by region, time,…)

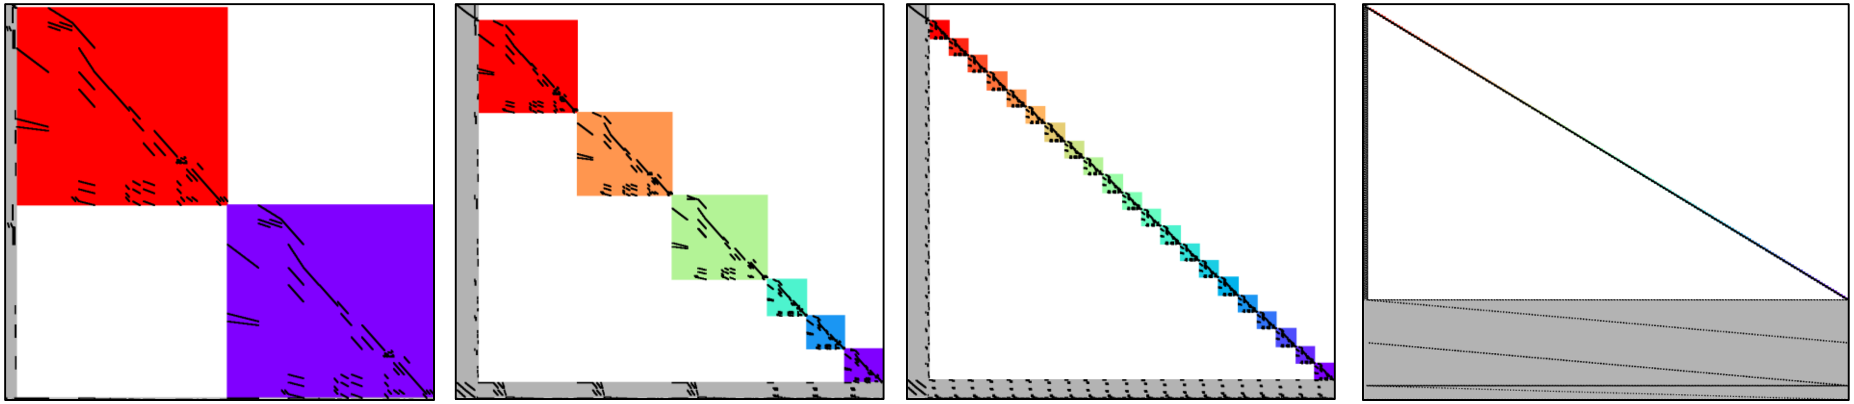

Plots show four different annotations of identical model

Model Annotation cont.

- How to annotate Model depends on how the model should be “decomposed” (by region, time,…)

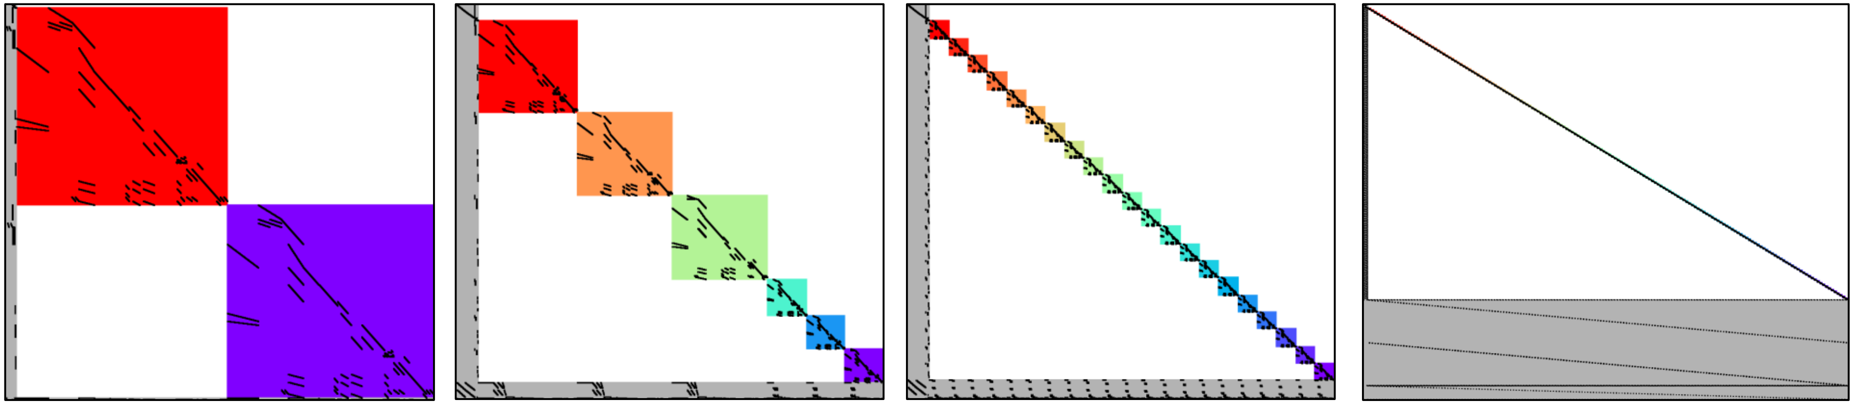

Plots show four different annotations of identical model

- Blocks of equal size are beneficial

GAMS/PIPS Solver Link

Distributed Model Generation

Distributed Model Generation

- "Usual Model": model generation time << solver time
- For **LARGE**-scale models the model generation may become significant:
  - due to time consumption
  - due to memory consumption
  - due to hard coded limitations of model size (# non-zeroes < ~2.1e9)

→ Distributed "block-wise" model setup in PIPS

→ Model annotation determines block membership of all variables and constraints

→ Distributed GAMS processes can generate the separate blocks

Distributed Model Generation cont.

Model + Annotation → Block generation

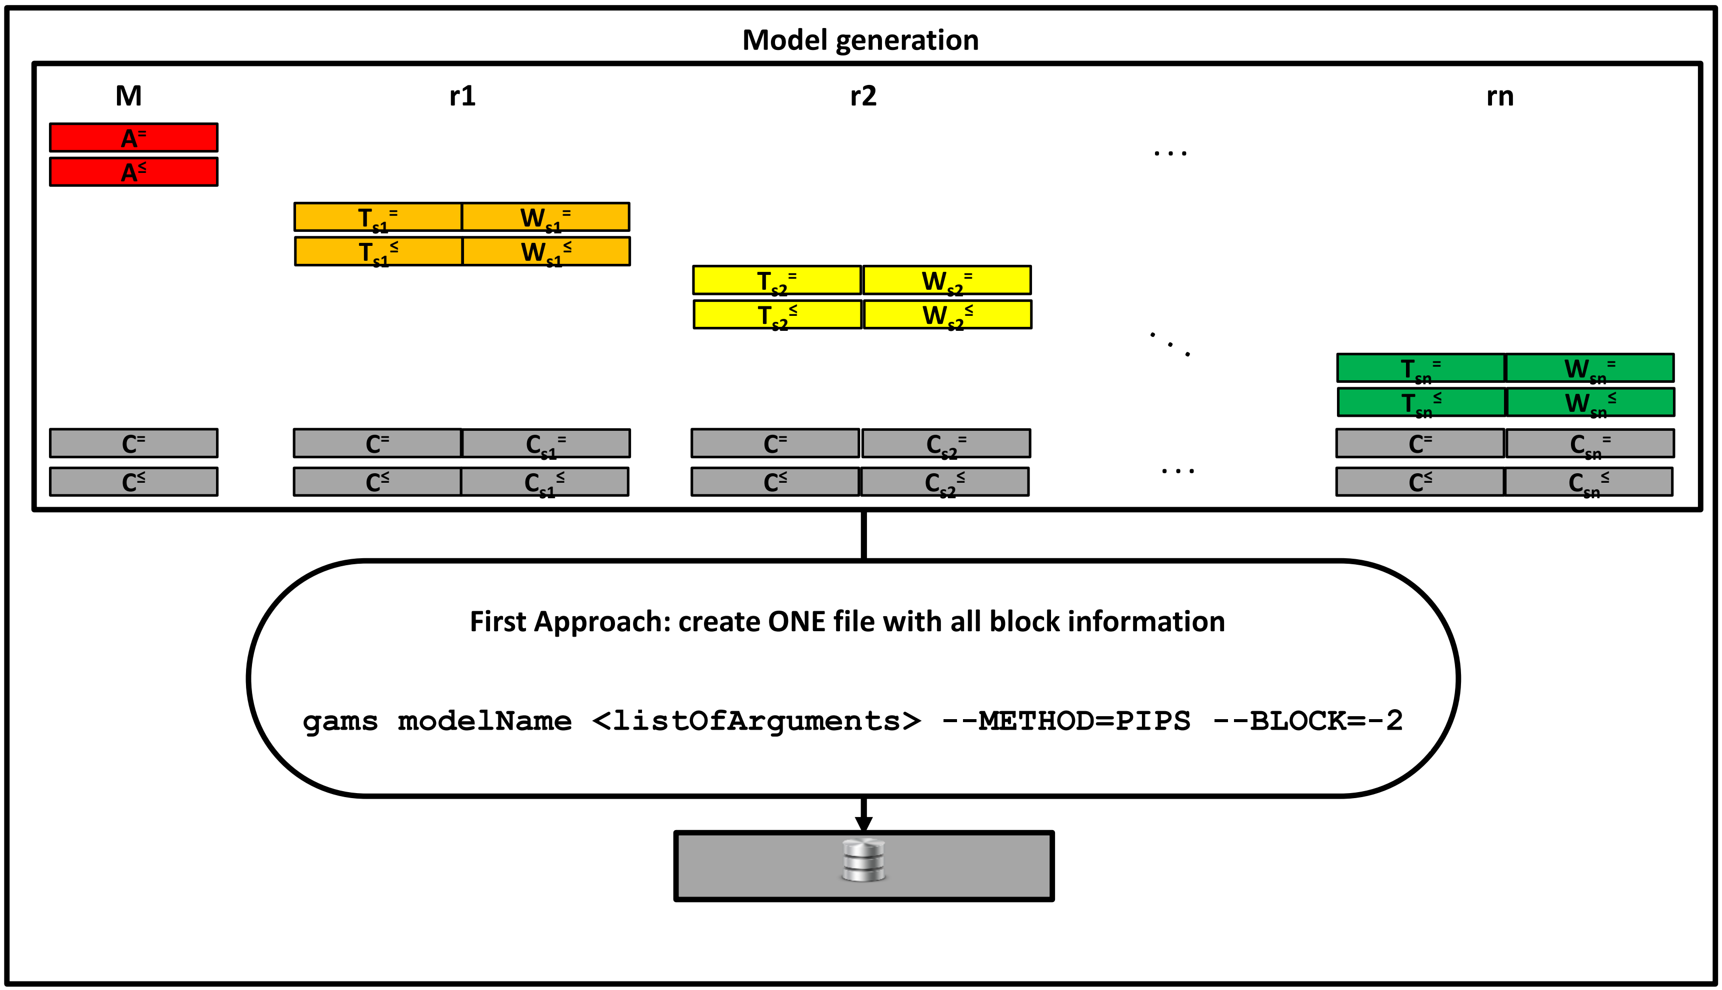

Distributed Model Generation cont.

Model + Annotation → Block generation

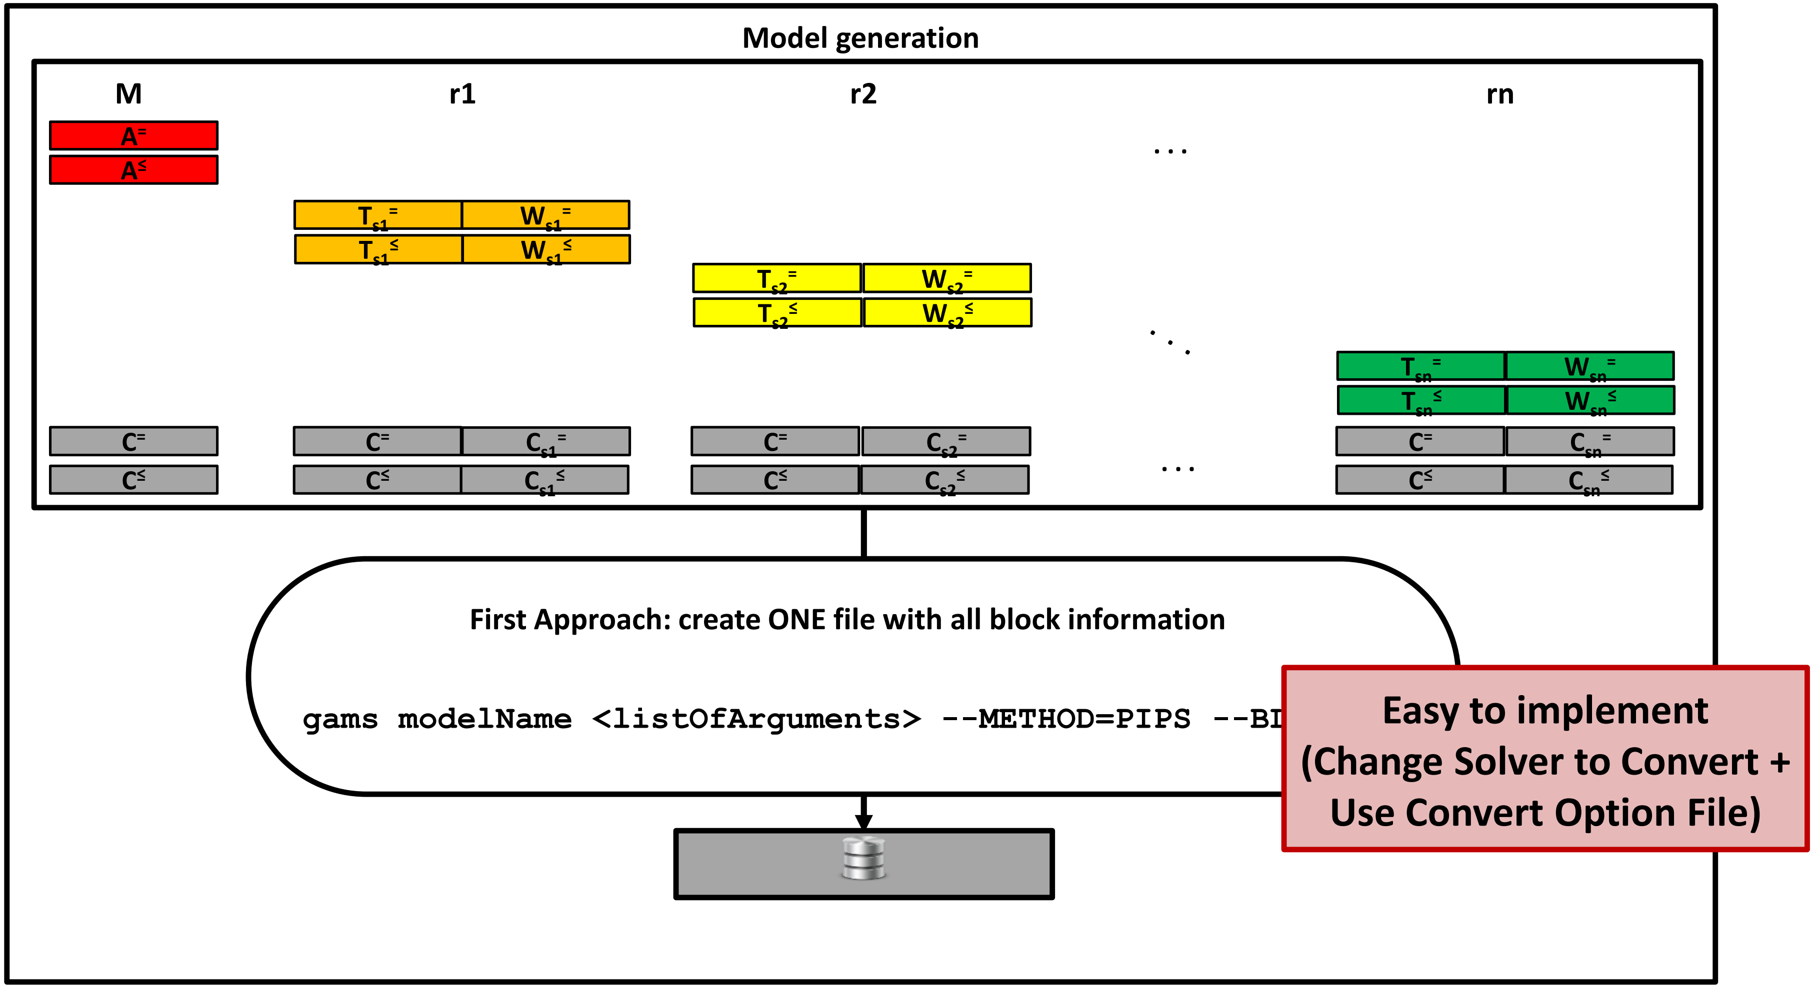

Distributed Model Generation cont.

Model + Annotation → Block generation

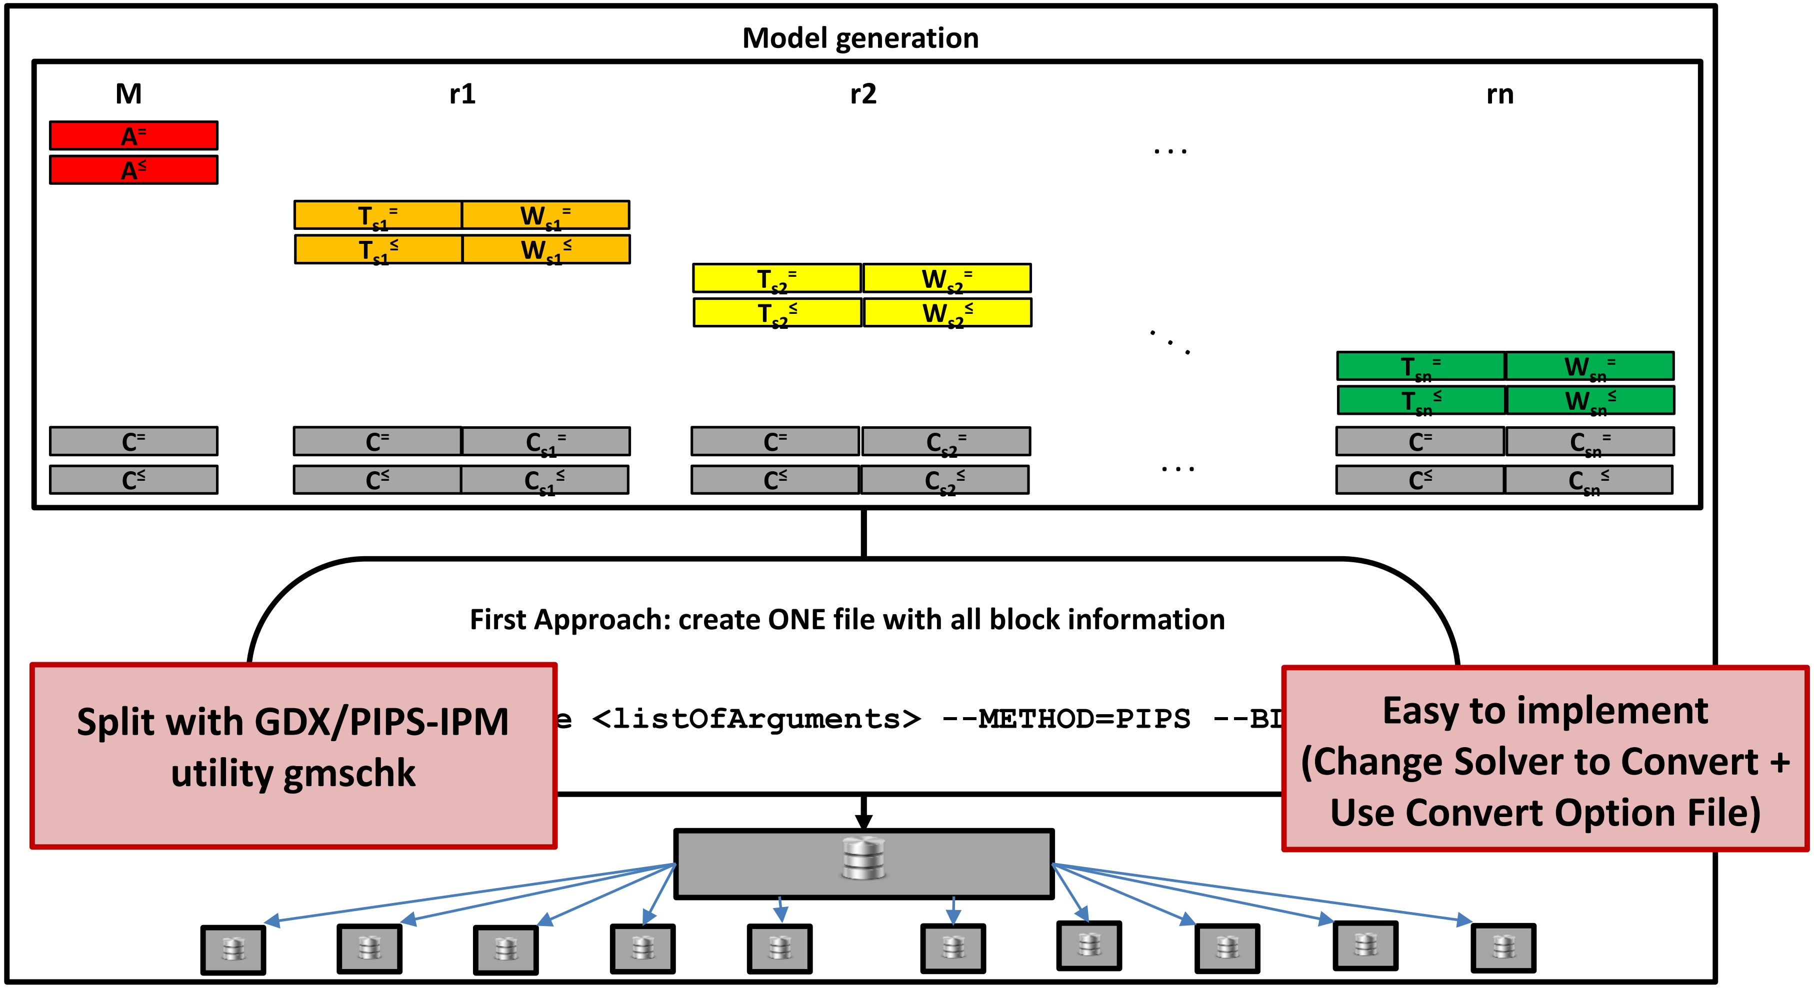

Distributed Model Generation cont.

Model + Annotation → Block generation

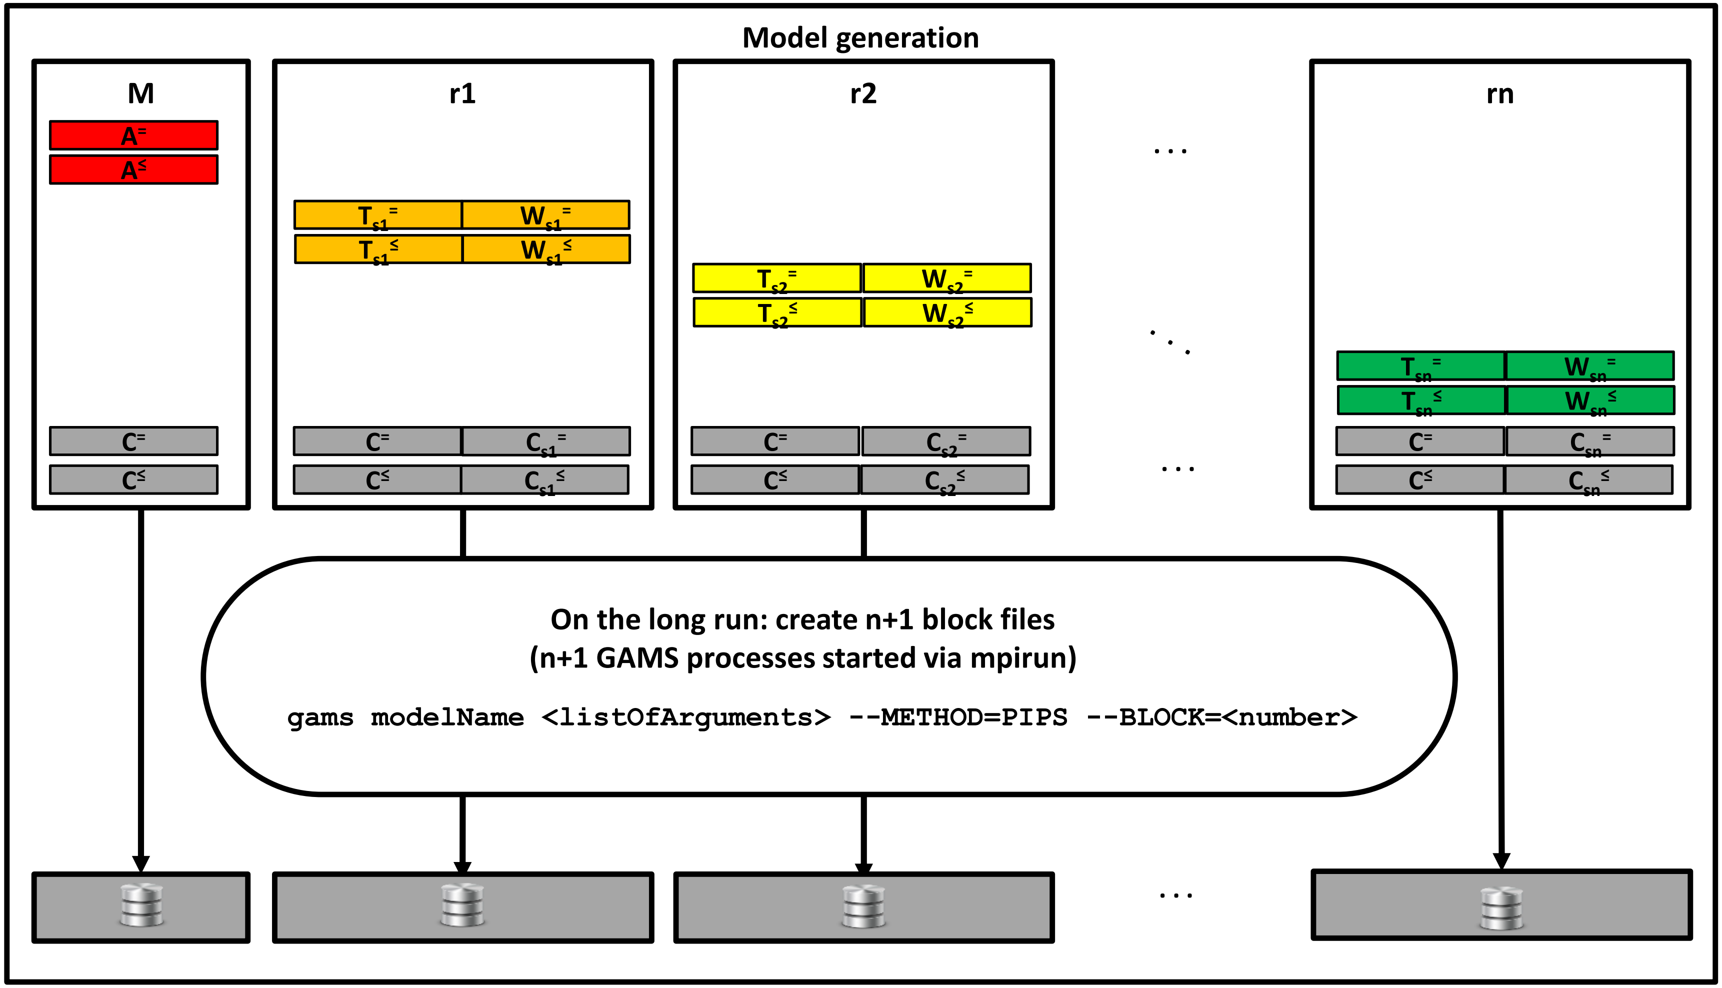

Distributed Model Generation cont.

Model + Annotation → Block generation

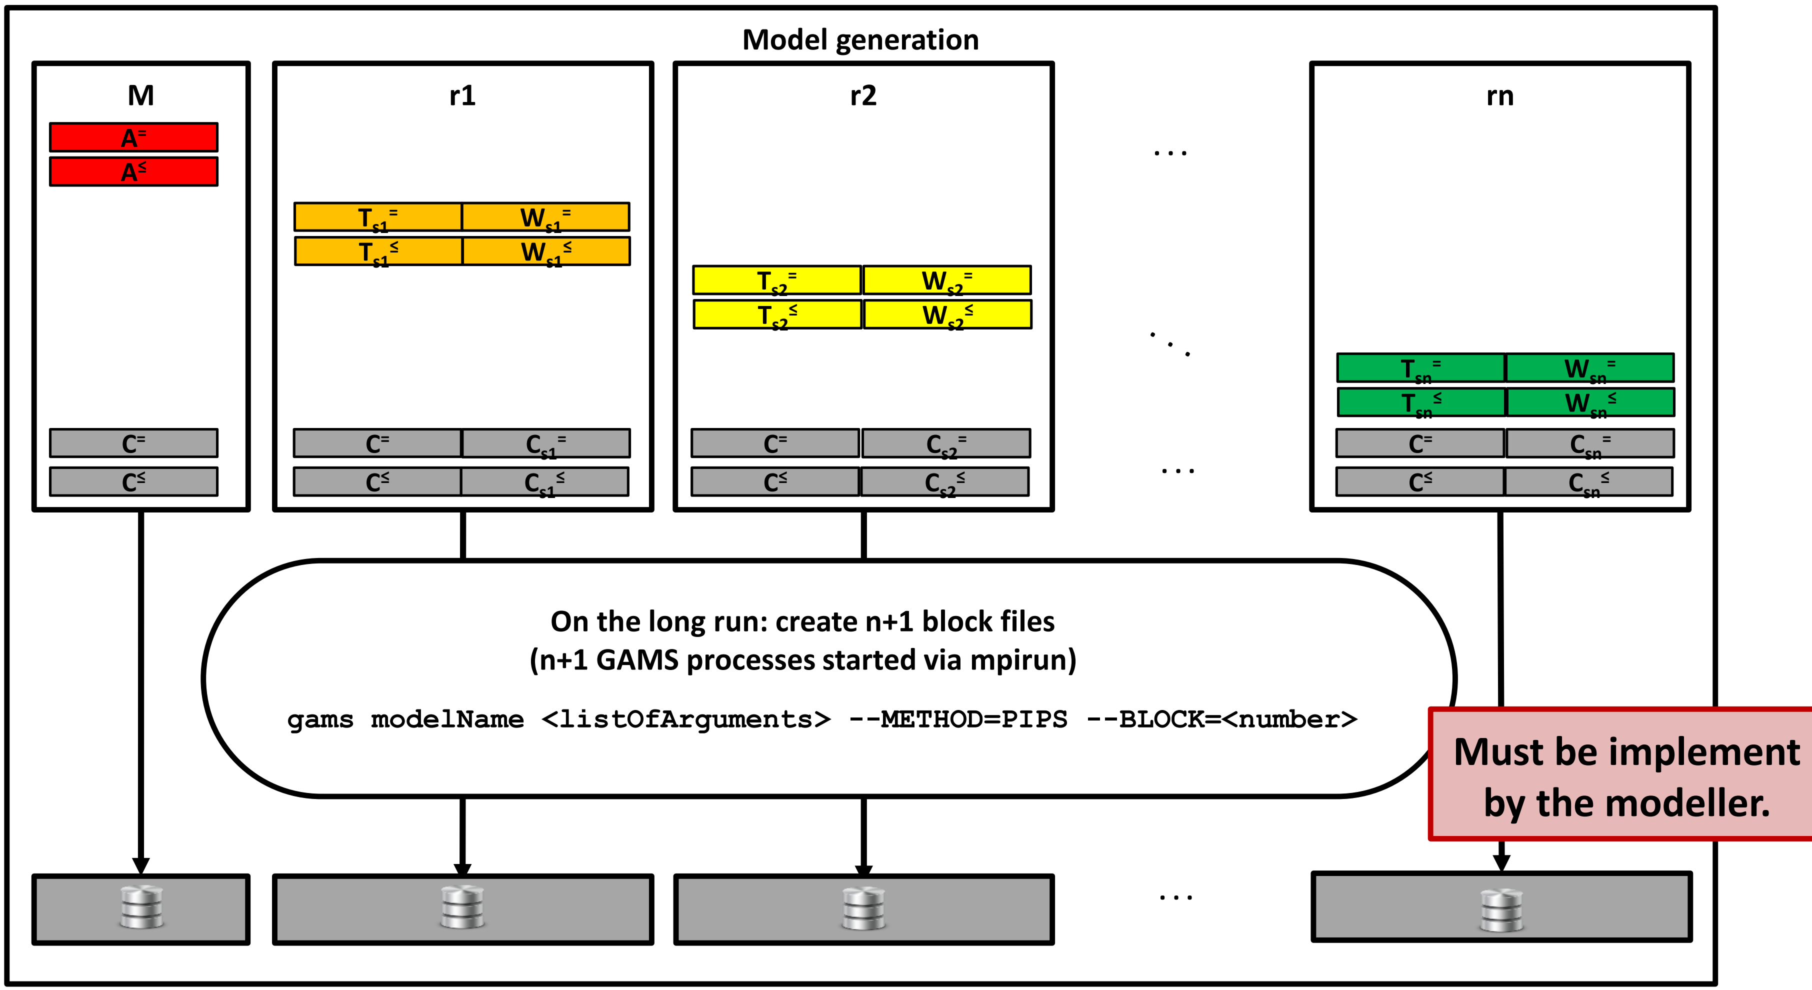

Computational Experiments

Stochastic Problem (no coupling constraints)

- #blocks vs. #MPI tasks
  - #MPI tasks <= #blocks
  - Best performance if #blocks % #MPI tasks = 0
  - Test case: 32 blocks (1 node, 1–32 MPI tasks)

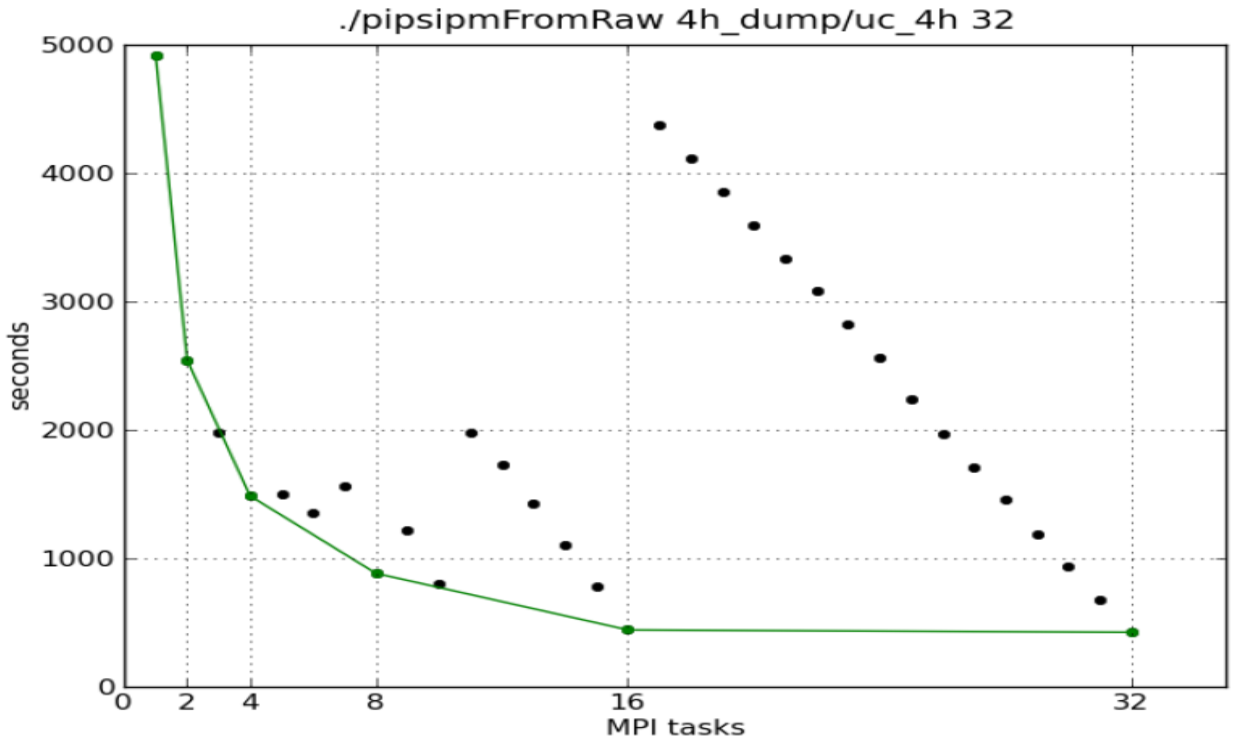

Slide provided by T. Breuer (JSC)

Profiling

Vampir Analysis of PIPS

Slide provided by T. Breuer (JSC)

Limitations of “standard” Soft- & Hardware

#t	#r	#blocks	#rows (E6)	#cols (E6)	#NZ (E6)	~Mem (GB)	time
730	10	10	0.7	0.8	2.8	2.0	00:01:22
730	10	500	35.0	38.7	142.8	95.7	01:09:36
730	10	2,500	175.3	193.5	713.9	478.8	09:32:55
730	10	4,000	280.5	309.6	1,142.2	767.1	19:22:55
730	10	7,500	526.1	580.5	2,141.2	~1,436.4	-
8,760	10	10	8.4	9.3	34.3	18.2	00:28:57
8,760	10	50	42.1	46.4	171.6	90.4	02:26:25
...							

Test runs were made on JURECA @ JSC

- 2x Intel Xeon E5-2680 v3 (Haswell), 2 x 12 cores @ 2.5GHz
- “fat” node with 1,024 GB Memory
- GAMS 24.8.5 / CPLEX 12.7.1.0
- Barrier Algorithm, Crossover disabled, 24 threads

Summary & Outlook

Summary & Outlook

BEAM-ME

- PIPS-IPM
  - Change(d) linear solver from MA27 (default) to PARDISO SC
  - Improve numerical stability
  - Implement (structure-preserving, parallel?) preprocessing
- GAMS/PIPS-IPM Link
  - Integrate model generation and solution into one user friendly process
  - Better user control of GAMS/PIPS
    - options (algorithmic, limits, tolerances)
- Annotation can be adapted for other Decomposition approaches (e.g. CPLEX Benders)
- GAMS-MPI/Embedded Code:
  - Implementation of Benders Decomposition in GAMS for ESM using the GAMS embedded code facility with Python package mpi4py to work with MPI (see talk of L. Westermann, *Tuesday, Oct 24, 10:30 - 12:00 track TB74 - room 372C*)
- Apply developed methods to several other large-scale ESM in Model Experiment: BALMOREL, DIMENSION, …

Project BEAM-ME

BEAM-ME

Supported by:

Federal Ministry for Economic Affairs and Energy

on the basis of a decision by the German Bundestag

A PROJECT BY

H L R I S

JÜLICH FORSCHUNGSZENTRUM

DLR Deutsches Zentrum für Luft- und Raumfahrt German Aerospace Center

ZIB