The 4th ISM-ZIB-IMI MODAL Workshop on Mathematical Optimization and Data Analysis

Date: 25th – 30th March 2019
Venue: The Institute of Statistical Mathematics, Tokyo, Japan

Solving Energy System Models with GAMS on HPC Platforms

Michael R. Bussieck

Economic Challenges Ahead

1. **Energy and Environmental Security**
2. Conflict and Poverty
3. Competing in a New Era of Globalization
4. Global Imbalances
5. Rise of New Powers
6. Economic Exclusion in the Middle East
7. Global Corporations, Global Impact
8. Global Health Crises
9. Global Governance Stalemate
10. Global Poverty: New Actors, New Approaches

BROOKINGS

https://www.brookings.edu/research/top-ten-global-economic-challenges-an-assessment-of-global-risks-and-priorities/

Energy and Environmental Security

"Energy and environmental security has emerged as the primary issue on the global agenda for 2007. Consensus has recently been forged on the potential for long-term economic, national security and societal damage from insecure energy supplies and environmental catastrophe, as well as the intense need for technological advances that can **provide low-polluting and secure energy sources**. Yet despite growing global momentum, there is still **little agreement on the best set of actions required to reduce global dependency on fossil fuels and greenhouse gas emissions**. Confounding the international policy challenge is the disproportionate impact of high oil prices and global warming across nations, insulating some countries from immediate concern while forcing others to press for more rapid change."

Energy System Models (ESM)

- For example: MARKAL/TIMES
  - Energy Technology Systems Analysis Program (ETSAP) of International Energy Agency (IEA) model *generators/frameworks*
    - Local, national, regional or global scale;
    - Single-region or multi-regional;
    - Partial equilibrium or general equilibrium;
    - Short term or long term (up to 2100 and beyond);
    - Perfect foresight or recursive-dynamic.

- *Essentially, all models are wrong, but some are useful.*
  (George Box)
- *The only function of economic forecasting is to make astrology look respectable.*
  (John Kenneth Galbraith/Ezra Solomon)

ESM Principles

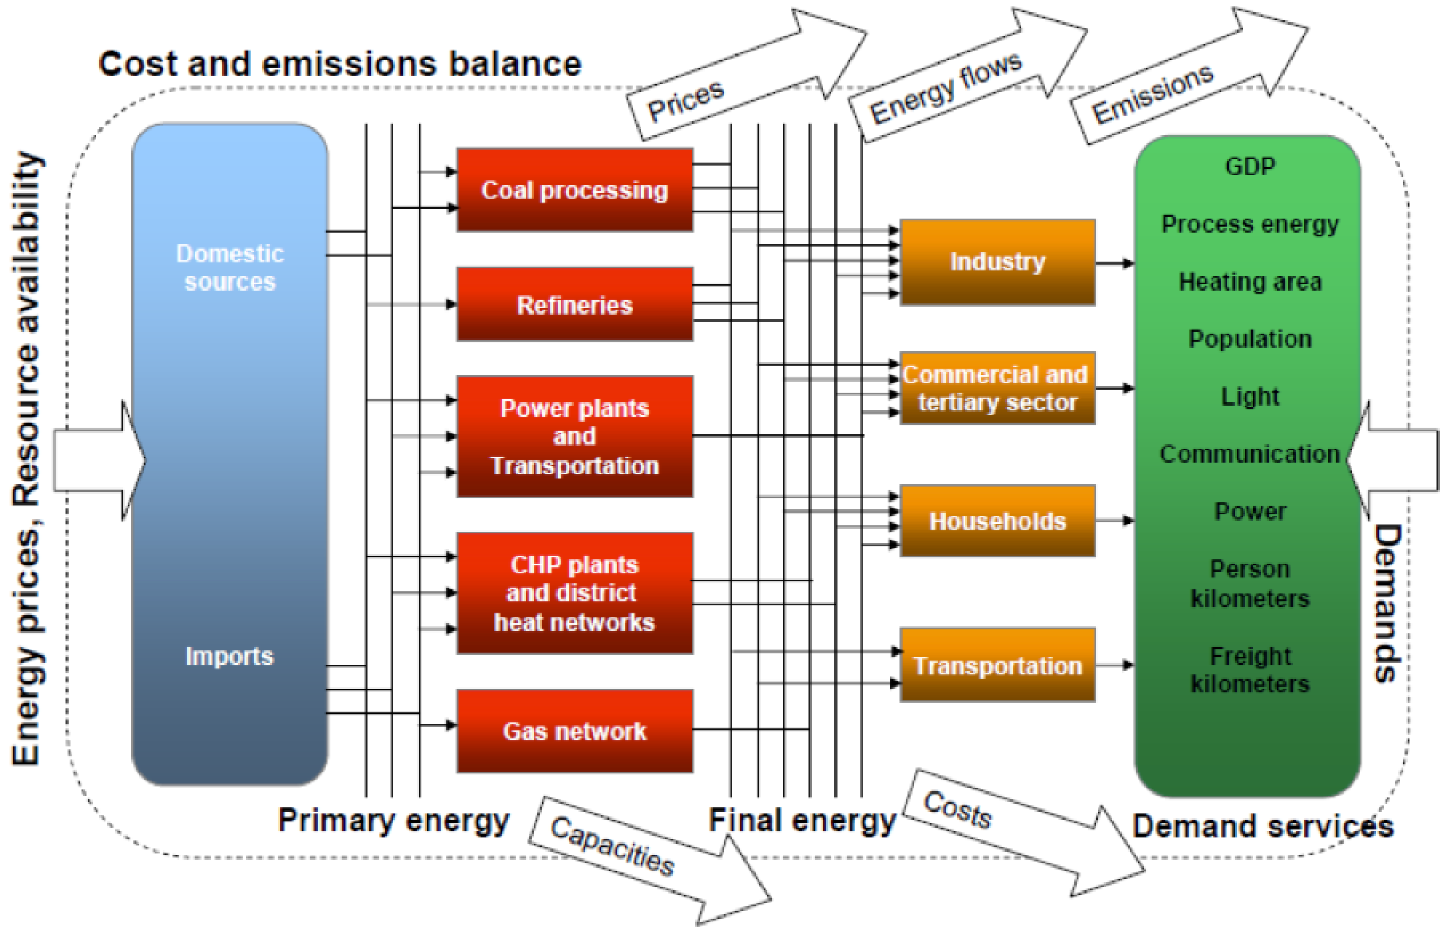

ESM Principles (2)

- Model Ingredients:
  - Technologies/Processes
    - Transform commodities into other commodities (e.g. fuel → electricity+emissions)
  - Commodities
    - Commodity is produced or consumed by a process (e.g. fuels, electricity, emission, money)
  - Time
  - Region
  - Policies
    - Minimum share of renewable energy
    - Maximum amount of (GHG) emissions
    - Minimum level of energy security
    - ...

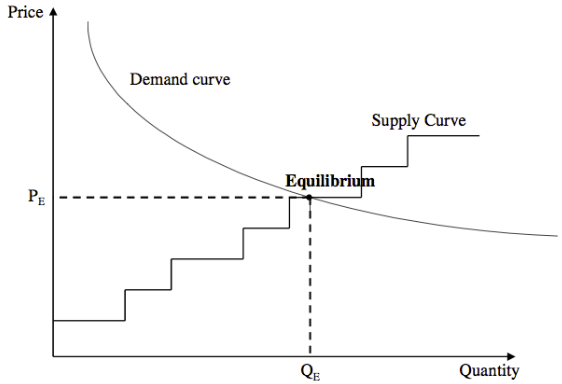

- The *mathematical*, *economic* and *engineering* relationships between these energy *producers* and *consumers* are basis of the ESM.

Building (ES) Models

- Quick (& dirty) prototypes
- Elaborate Models
  - Development & maintenance effort
    - MARKAL/TIMES ~10 person years
  - Lifetime of 15+ years
    - TIMES started ~1997, MARKAL ~1978
  - Large user base
    - TIMES is used by ~200 research teams in more than 50 countries)
- Model Development
  - Matrix generators
    - MaGen
    - OMNI, Haverly System LP “modeling language”
    - Concert, JuMP, Pyomo (concrete model)
  - Algebraic Modeling Languages (AML)
    - AIMMS, AMPL, GAMS, OPL, …

Principles of Algebraic Modeling Languages

- Matrix generators:
  - Use programming language as execution system plus use of languages' data structures
  - Often close to algorithm/solver
  - Skills: Application/Model/Algorithm + CS/IT
- AML:
  - *Describe* model algebra
    - LaTeX ←→ AML
  - Simple language and data model
  - Abstraction from algorithm/solver
    - Separation of UI, data, model, and algorithm
  - Skills: Application/Model
  - →Domain experts build models

Model Types

- Models build by domain experts (engineers/ economists):

Paul de Vos - A dog (Detail),1638 Oil

- Models build by mathematicians/ computer scientists

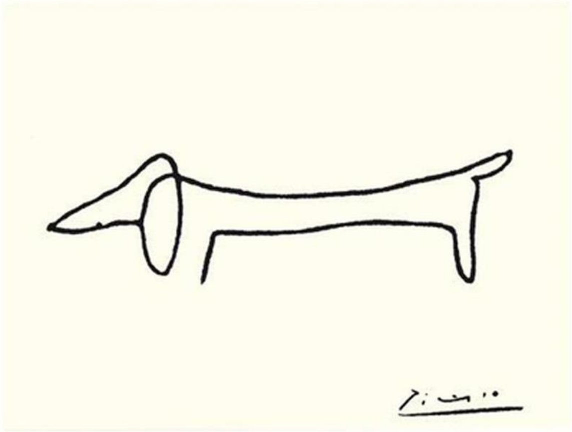

- **Models** of domain experts are widely known:
  - DICE model of Nobel Laureate (2018) William Nordhaus
- **Algorithms** of math/CS experts are widely known

Challenges from (ES) Models

- Many ESM implemented in AML (GAMS in particular)
  - https://www.energyplan.eu/othertools/
  - https://wiki.openmod-initiative.org/wiki/Open_Models
- ESM implement many well know problem classes:
  - Unit commitment (MIP)
  - Economic dispatch (LP)
  - Optimal Power Flow (NLP)
  - …
- Models challenge algorithms
  - Citius, altius, fortius (faster, higher, stronger)
    - Level of granularity
    - Time horizon
    - Global regions

Pushing the Envelope

- Expert (algorithm) knowledge helps
  - “We have a MIP problem that solving with B&B takes 17 hours to run in my notebook & it takes 27 hours ‼ to run in my big server computer. ”
  - “We are sure that will need lots of additional computational resources, we made an agreement with GOOGLE (thru a google partner in Chile) to run our MIP models in google platforms and explore parallel processing”
  - Analysis:
    - “big server”: Slow VM as preferred by IT nowadays
    - Root LP (simplex takes 6.5h)
    - Lots of cut generation without improving bound
    - Solution found by heuristic in node 0 with 0.12% gap
    - Setting of a few solver options: time down to 2h

  - Large ESM Model REMix (LP) becomes computationally intractable
  - Analysis:
    - LP solved “quickly” by barrier
    - Crossover takes forever
    - No need for a basic solution: disable crossover

→ Brain beats machine

Limitations of “standard” Soft- & Hardware

#t	#r	#scen	#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 with model ESM REMix 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

BEAM-ME: An Interdisciplinary Approach

Energy System Modeling → **Modeling Language** → **Solver Development** → **High Performance Computing**

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

ESM REMix from DLR (German Aerospace Center)

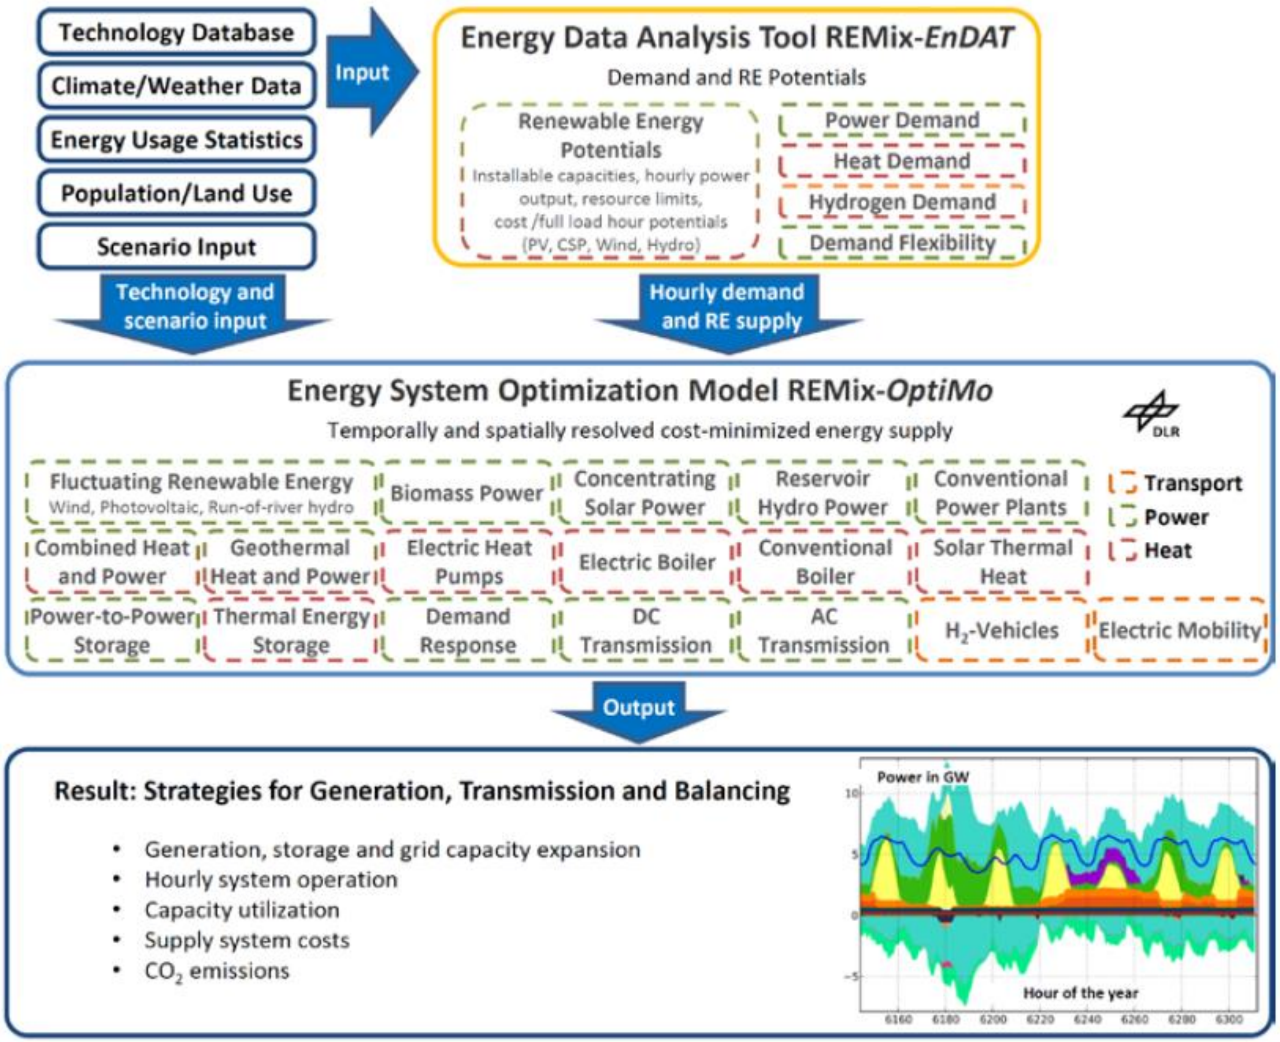

REMix Model Investment

• Start year:	2006
• #PhD Thesis:	13 (6 in progress)
• Person years (devel/use):	10-20
• Maintenance:	1 PY/a
• #Users:	11
• #Developers:	4
• #IT/UI Maintenance	0.25 PY/a

Explore Algorithms

- Block structure of ESM (time, region, technology, …)
- Speed-up of *traditional* methods:
  - Benders/Lagrange Decomposition
  - Rolling horizon
  - …
- Interior Point Algorithms
  - Barrier/Interior Point works way better than Simplex
  - PIPS (Parallel Solvers for Optimization Problems)
    - PIPS-IPM:
      - Interior Point Algorithm based on OOQP
      - Exploit block structure of the underlying LP to run massive parallel on HPC platforms

PIPS-IPM

- Parallel interior-point solver for LPs (und QPs) from *stochastic* energy models
- Exploit block structure when solving the Central Path equation system
- Main developers: Cosmin Petra, Miles Lubin
- Good reasons to work with PIPS-IPM:
  - PIPS-IPM is open source
  - PIPS-IPM ran already successfully on HPC architectures
- Extensions of PIPS-IPM (by ZIB, Daniel Rehfeld)
  - Linking constraints
  - Presolve
  - …

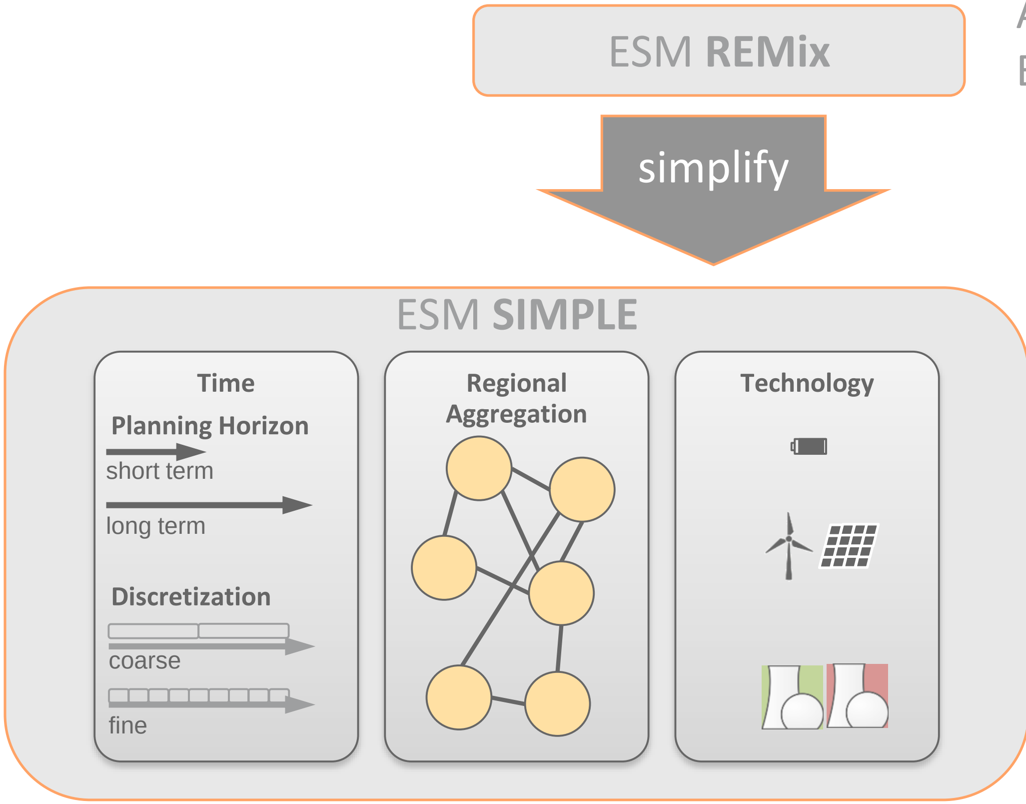

A large-scale highly complex ESM developed by DLR

A simplified generic ESM that maintains the relevant model structure.

A “sandbox model” for algorithmic experiments.

Slide from Sep 2016, BEAM-ME Meeting in Stuttgart

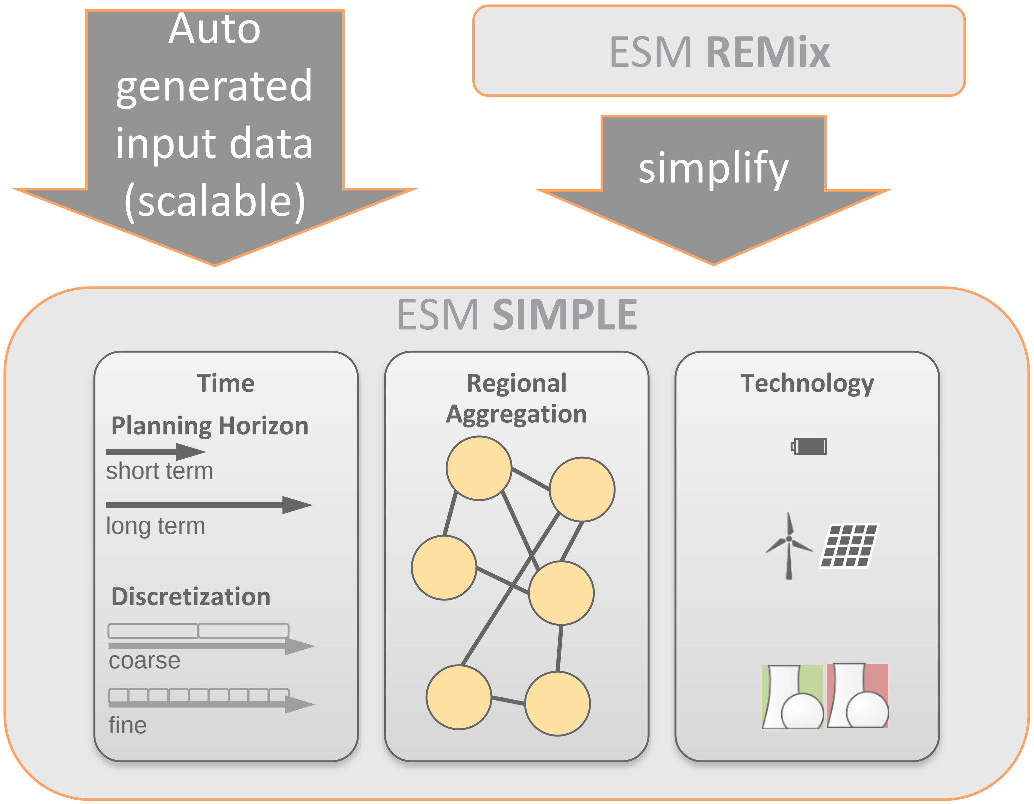

Fast and easy generation of input data (based on REMix instance)

Slide from Sep 2016, BEAM-ME Meeting in Stuttgart

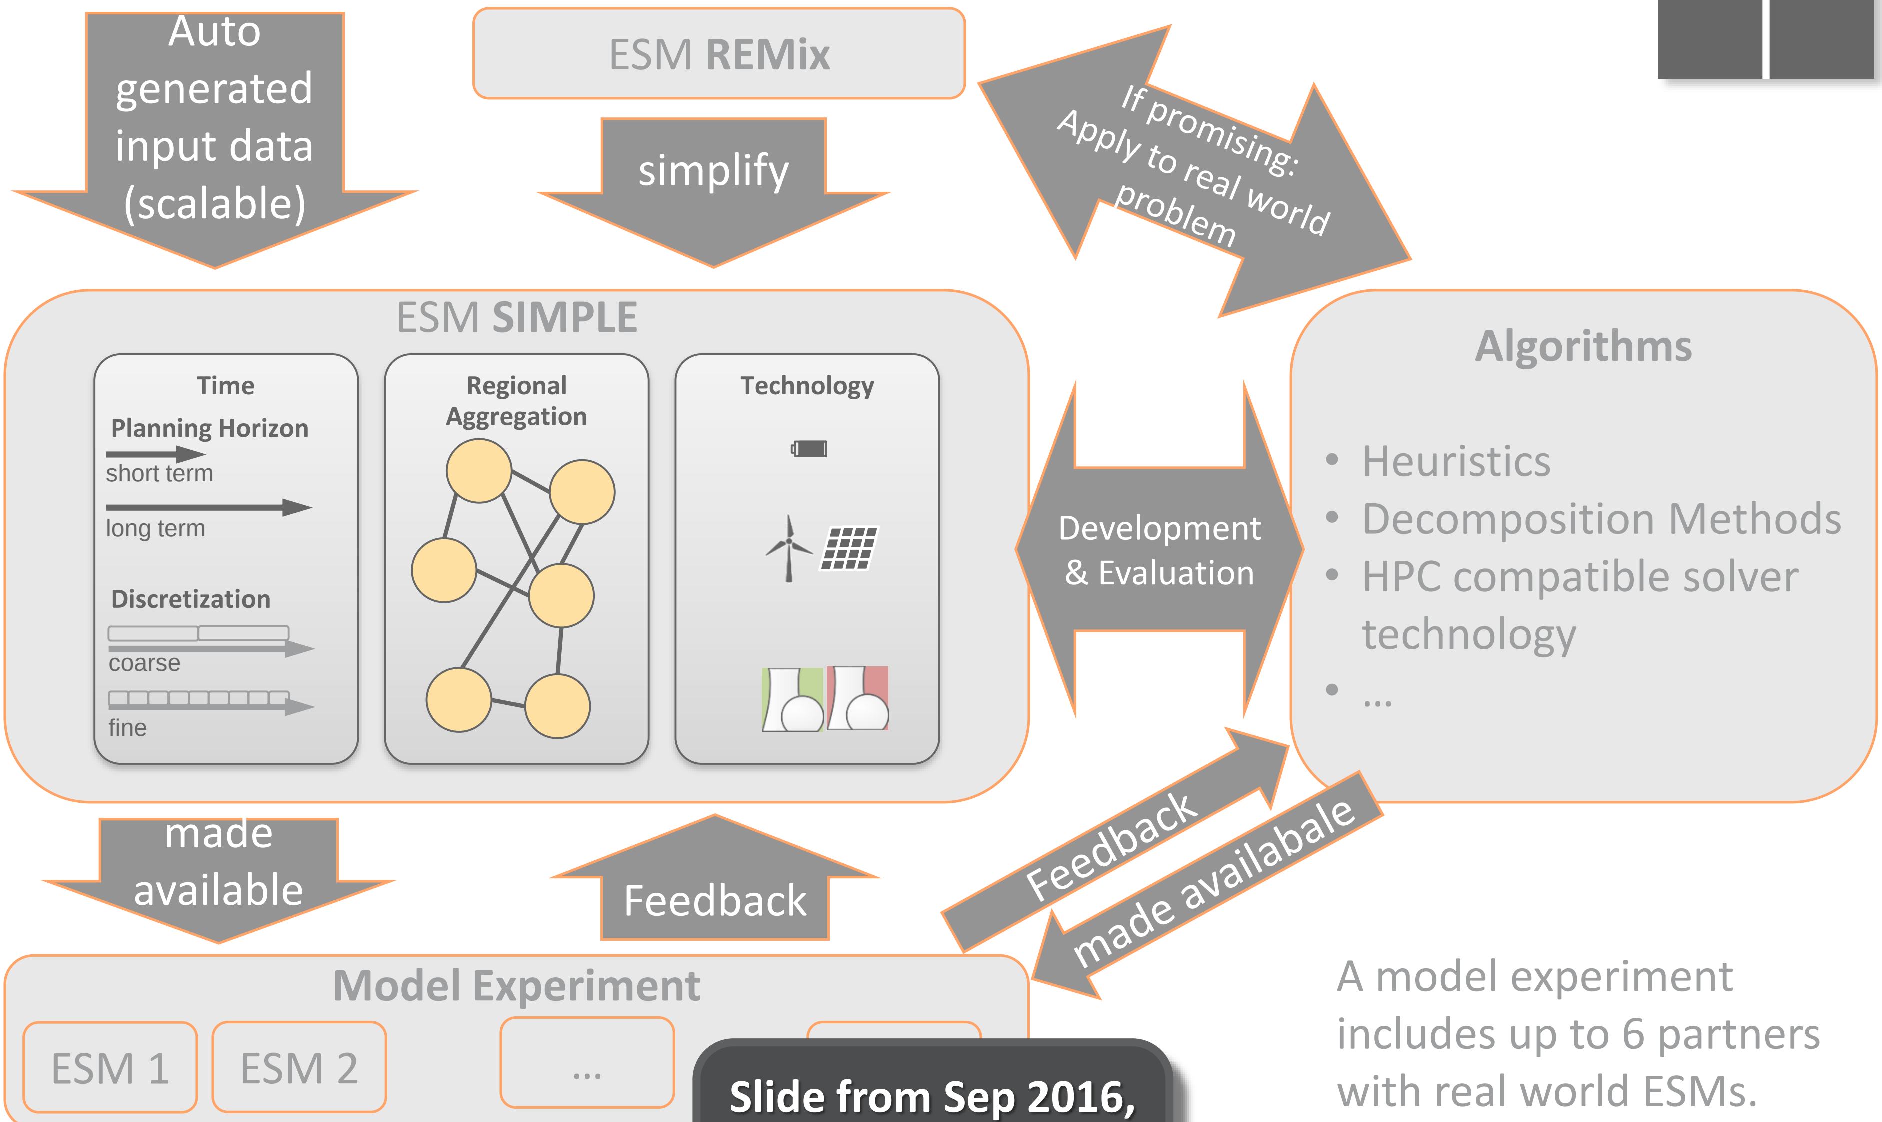

A model experiment includes up to 6 partners with real world ESMs.

GAMS

SIMPLE Model Structure

	Power flow between regions	Expand link capacity	Power generation per plant	Uncovered demand (slack)	Storage level	Storage inflow	Storage outflow	Expand plant capacity	Expand storage capacity	Regional emission allowance	Emission costs	Regional costs	Global costs
Objective function (min cost)		x										x	x
Regional objectives			x	x				x	x		x	x	
Power balance	x		x	x		x	x						
Plant capacity			x					x					
Storage balance					x	x	x						
Storage capacity					x				x				
Regional emission cap			x							x			
Emission costs	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]					x		
Global emission cap	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]				x			
Link capacity	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]	[illegible]							

Slide from Sep 2016, BEAM-ME Meeting in Stuttgart

BEAM-ME

GAMS/PIPS-IPM[1,2]

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

$$
\begin{array}{llllllll}
\min & \sum_{i=0}^{N} c_i^T x_i & & & & & & \\
\text{s.t.} & T_0 x_0 & & & & & & = b \\
 & T_1 x_0 & + W_1 x_1 & & & & & = h_1 \\
 & T_2 x_0 & + & & W_2 x_2 & & & = h_2 \\
 & \vdots & & & & \ddots & & \vdots \\
 & T_N x_0 & + & & & & W_N x_N & = h_N \\
 & F_0 x_0 & + F_1 x_1 & + & F_2 x_2 & \cdots & F_N x_N & = g
\end{array}
$$

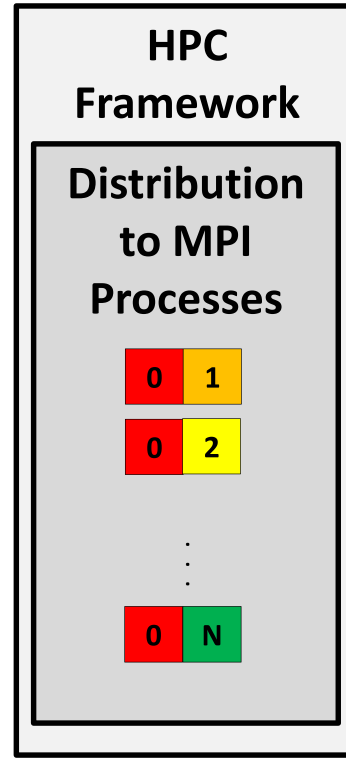

- Block diagonal structure allows parallelization of linear algebra within PIPS-IPM
- Solve N systems of linear equations in parallel instead of one huge system

[1] Petra et al. 2014: *"Real-Time Stochastic Optimization of Complex Energy Systems on High-Performance Computers"*
[2] Breuer et al. 2017: *"Optimizing Large-Scale Linear Energy System Problems with Block Diagonal Structure by Using Parallel Interior-Point Methods."*

GAMS/PIPS-IPM cont.

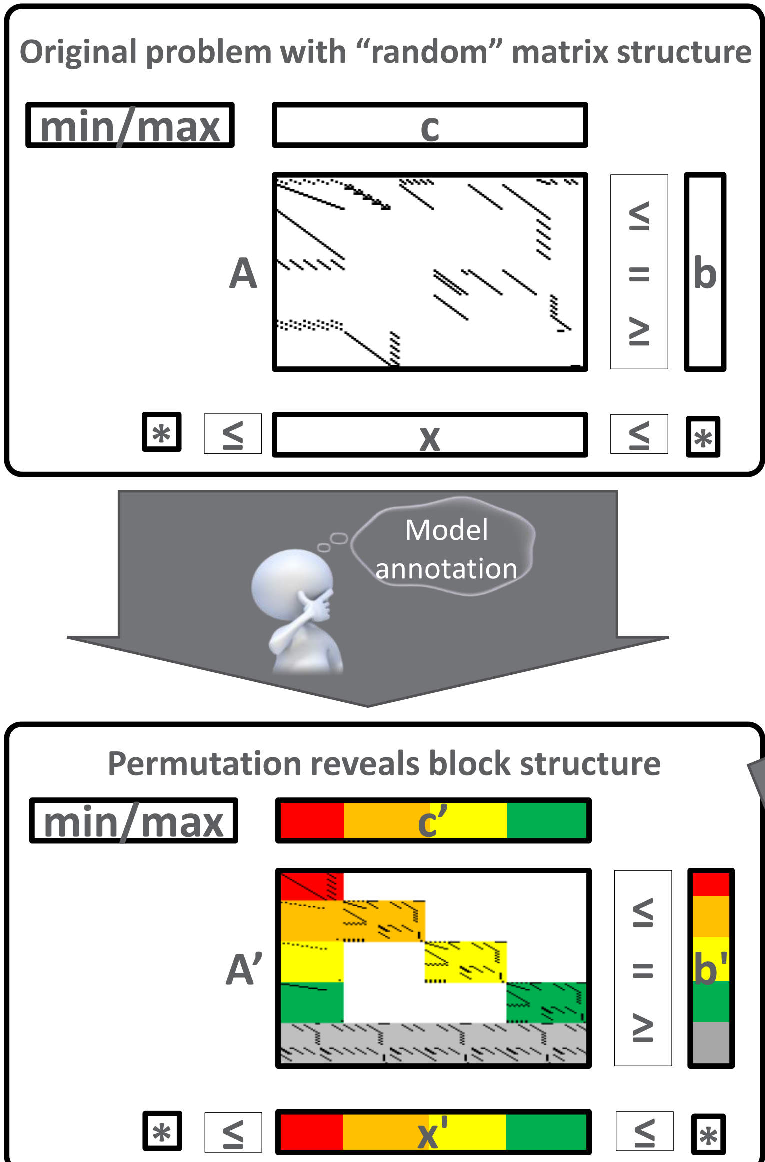

$$
\begin{array}{lllllllll}
\min & \sum_{i=0}^{N} c_i^T x_i & & & & & & & \\
\text{s.t.} & T_0 x_0 & & & & & & & = b \\
& T_1 x_0 & + & W_1 x_1 & & & & & = h_1 \\
& T_2 x_0 & + & & & W_2 x_2 & & & = h_2 \\
& \vdots & & & & & \ddots & & \vdots \\
& T_N x_0 & + & & & & & W_N x_N & = h_N \\
& F_0 x_0 & + & F_1 x_1 & + & F_2 x_2 & \cdots & F_N x_N & = g
\end{array}
$$

Model Annotation cont.

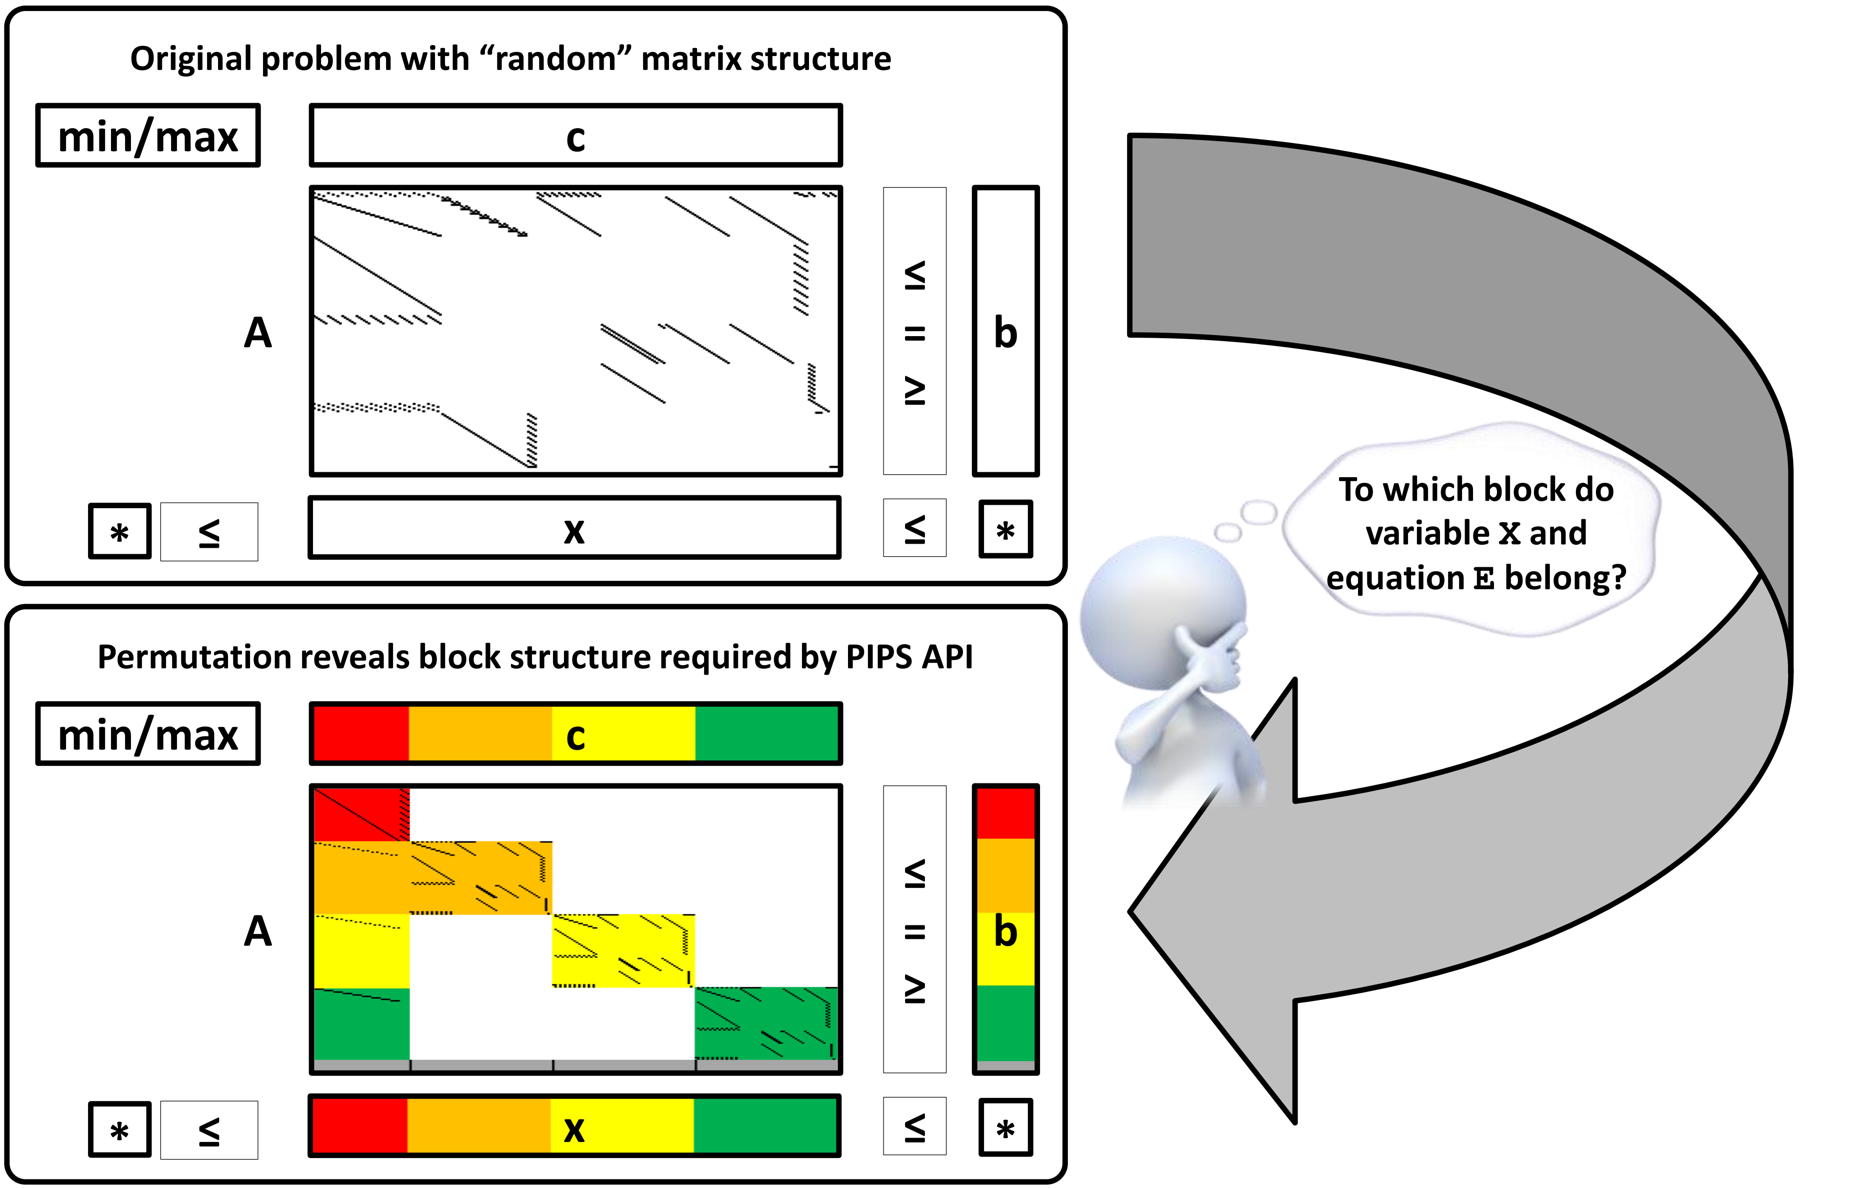

GAMS/PIPS-IPM cont.

Model Annotation by .stage attribute

Matrix structure required by PIPS API

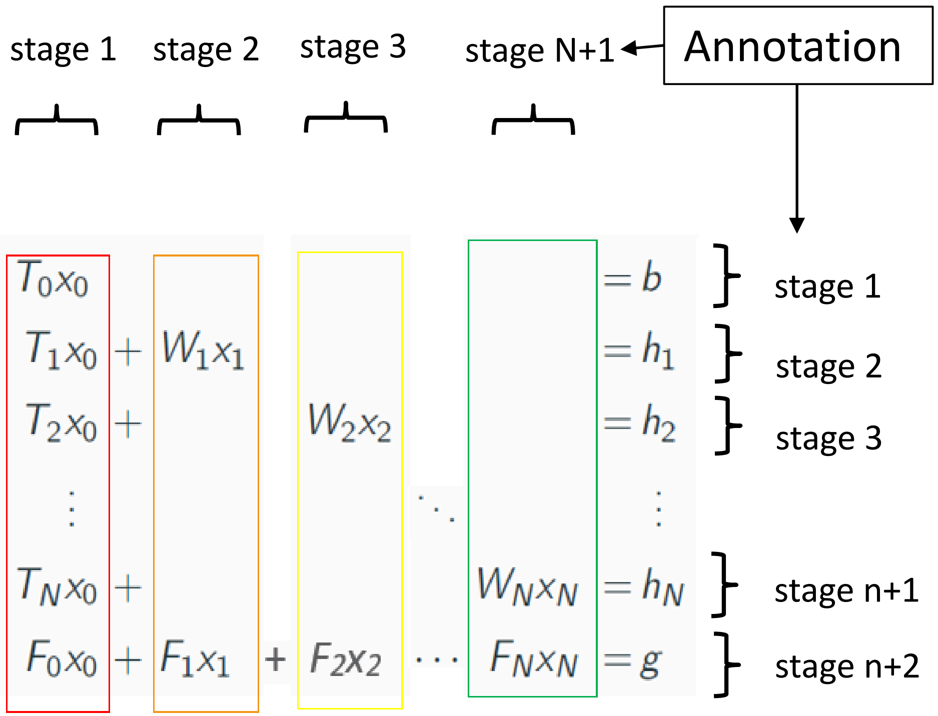

Exemplary Annotation for simple energy system model (regional decomposition)

```
[...]
* Master variables and equation
FLOW.stage(t,net(rr1,rr2))        = 1;
LINK_ADD_CAP.stage(net(rr1,rr2)) = 1;
[...]
* Block variables and equations
POWER.stage(t,rp(rr,p))      = ord(rr)+1;
EMISSION_SPLIT.stage(rr,e)  = ord(rr)+1;
[...]
eq_power_balance.stage(t,rr)     = 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)  = card(rr)+2;
```

Other Ways to Specify Block Structure

- StructJuMP (PIPS)
  - C. Petra et. al.
  - https://github.com/StructJuMP/StructJuMP.jl
- Structured Modelling Language (OOPS)
  - A. Grothey et. al.
  - https://www.maths.ed.ac.uk/ERGO/sml/

```
block MCNFArcs{a in ARCS}: {
  set ARCSDIFF = ARCS diff {a};  # local ARCS still present

  block Net{k in COMM}: {
    var Flow{ARCSDIFF} >= 0;
    # flow into node - flow out of node equals demand
    subject to FlowBalance{i in NODES}:
      sum{j in ARCSDIFF:arc_target[j]==i} Flow[j]
      - sum{j in ARCSDIFF:arc_source[j]==i} Flow[j] = b[k,i];
  }
```

GAMS/PIPS-IPM cont.

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

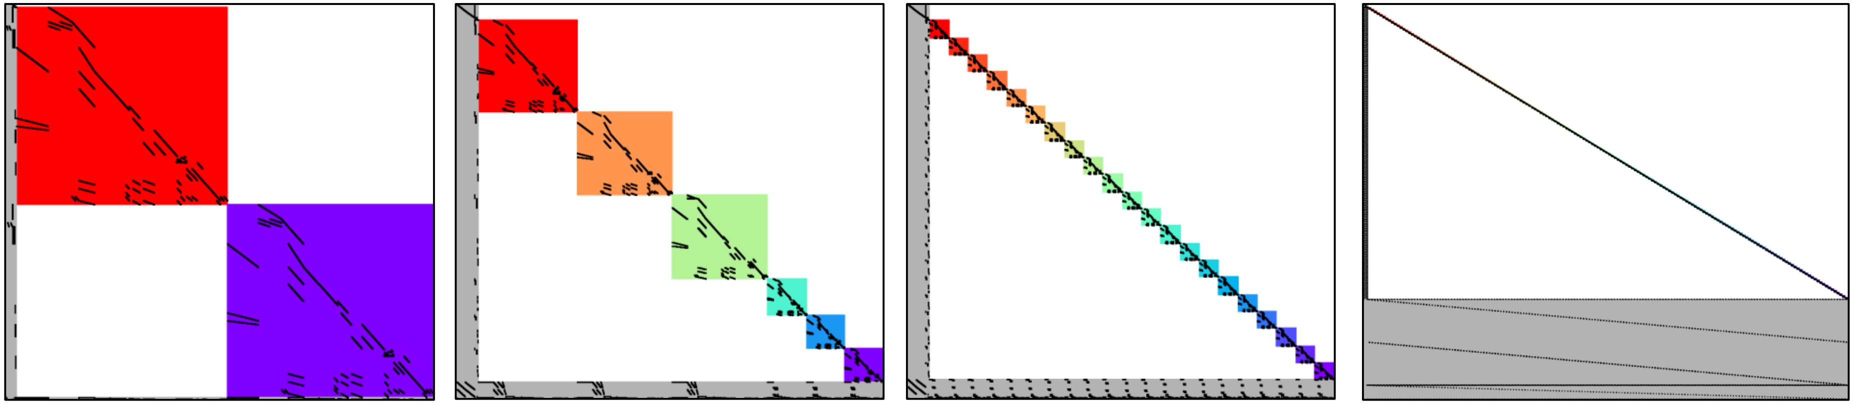

Plots show four different annotations of identical model

- Blocks of equal size are beneficial
- Few linking variables/constraints

GAMS/PIPS-IPM - Model Generation

BEAM-ME

- Process of *model generation* is separated from running the solver
- Allows different platforms for model generation and running PIPS
  - GAMS sometimes tied to WIN (e.g. because data comes from Excel, …)
  - PIPS & GDX run on ARM, GAMS does not
- Model Instances are stored in Gams Data eXchange (GDX) files:

```
//model definition
//model annotation
option lp=convertd;
$echo jacobian <filename[noVEnames]>.gdx > convertD.opt
myModel.optfile=1;
solve myModel min z use lp;
```

GAMS/PIPS-IPM - Model Generation

BEAM-ME

Monolithic Model Generation + Sequential Slicing

Create n+1 (gdx files containing) annotated model blocks | **Run PIPS-IPM**

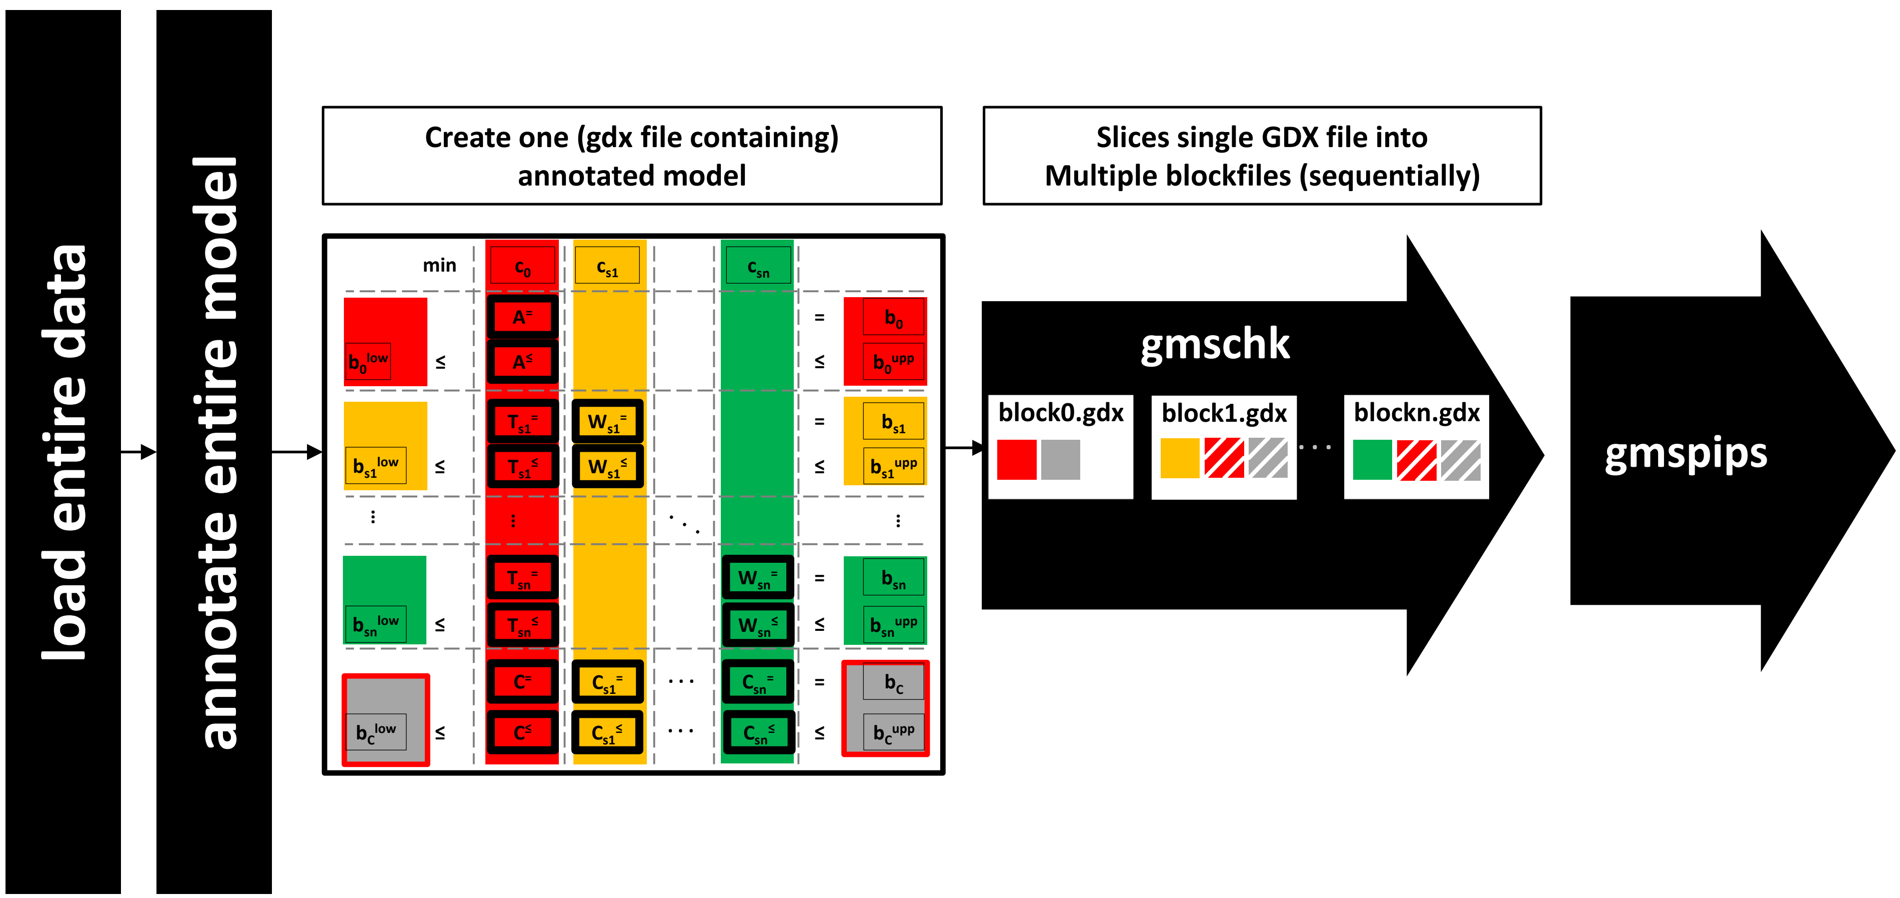

GAMS/PIPS-IPM - Model Generation

Parallel Model Generation

Simultaneously create n+1 (gdx files containing) annotated model blocks

Run PIPS-IPM

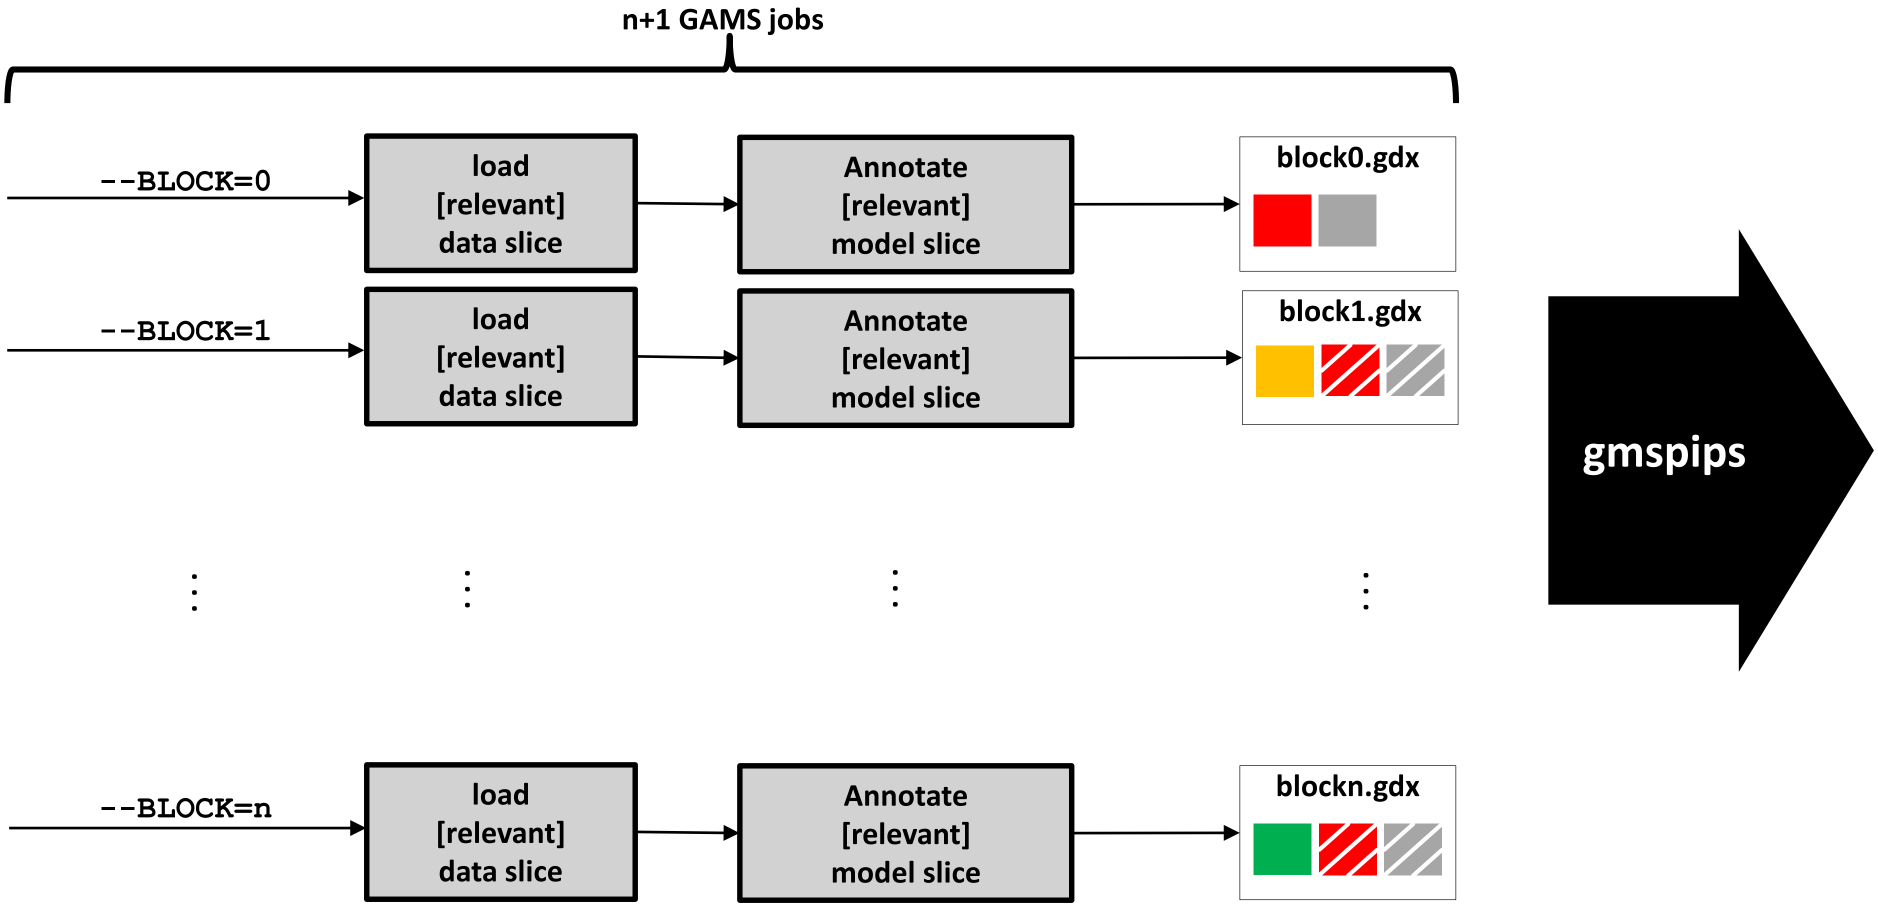

GAMS/PIPS-IPM - Model Generation

Computational results (simple4pips.gms)

# regions	#time steps	# blocks	#rows	#columns	#NZ	mode	Time [hh:mm:ss]	Memory [GB]
250	8760	366	21,725,667	23,917,146	86,449,311	mono + slicing	08:36	16.50
						parallel[1]	00:26	366 x 0.08
2000	35040	1461	654,329,207	720,310,215	2,603,662,965	mono + slicing	07:46:41	540.00
						parallel[2]	03:04	1461 x 1.29

[1]: 366 blocks in parallel (8 nodes @JURECA)
[2]: 1461 blocks in parallel (31 nodes @JURECA)

Computational Result(s)

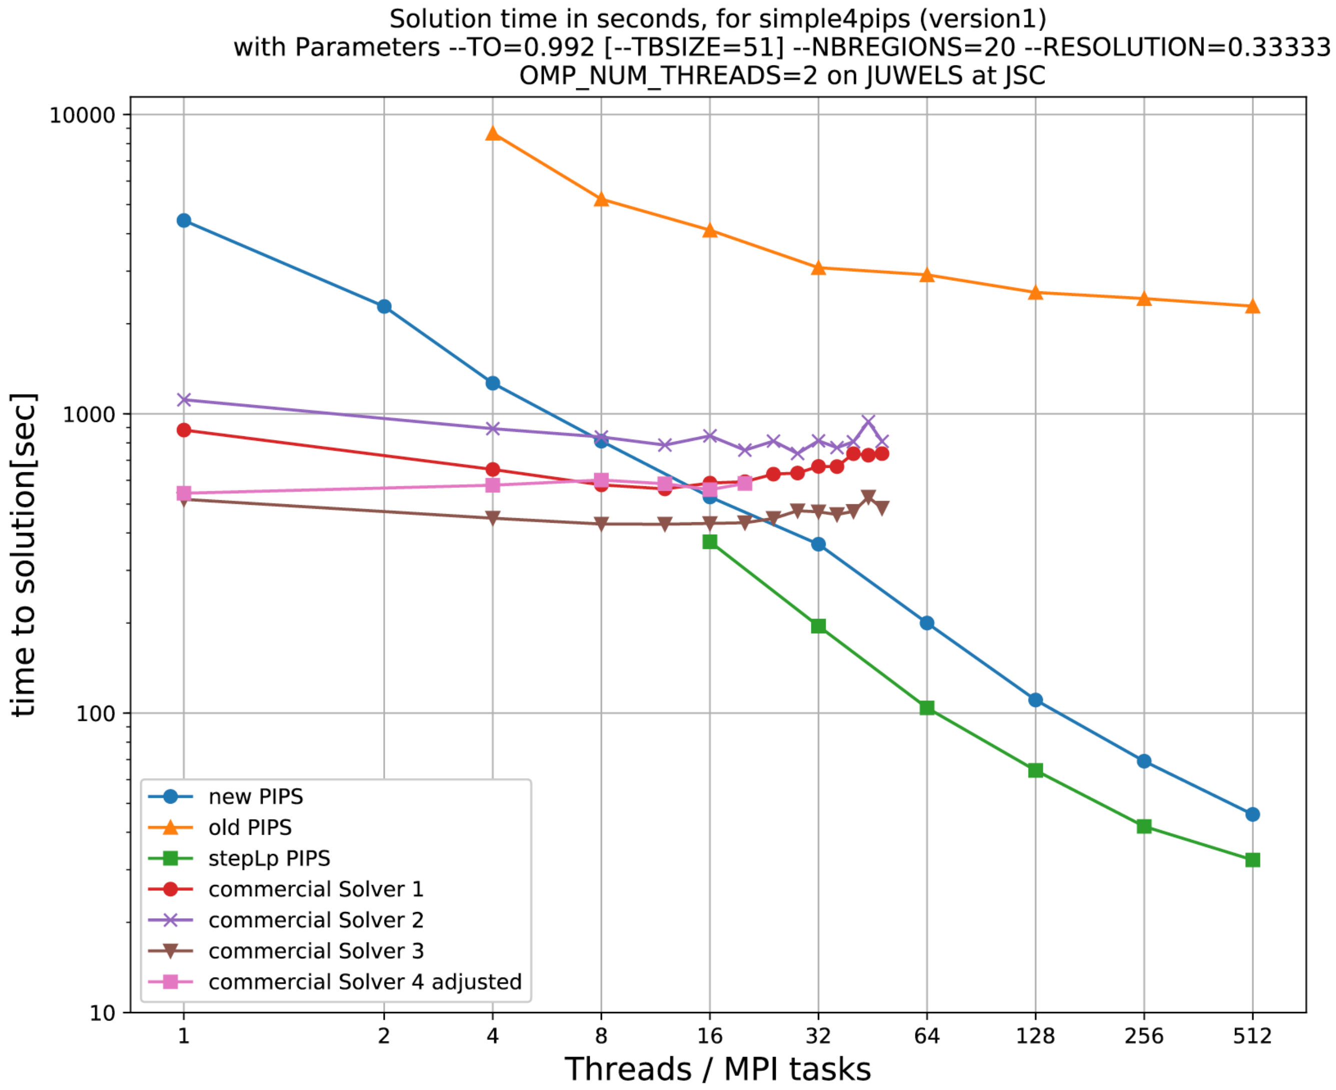

Computational Result(s) cont.

- TU Dresden, Modell ELMOD
- Model stats:
  - 254,304,961 rows
  - 224,677,686 columns
  - 712,452,541 non-zeroes
  - 438 blocks
- GAMS sequential generation time: 1,010 secs
- Commercial codes (16 threads, 2TB, barrier, no crossover):
  - CPLEX: 127,799.61 secs (~36 h)
  - Gurobi: 132,944.93 secs (~37 h)
- "*Hab gerad bei PIPS noch einen Memory Bug, kurz, bevor er konvergiert. Aber wenn der gefixt ist, sollte die Lösung nicht viel mehr als 10 Minuten dauern (allerdings auch mit 438 MPI Prozessen).*"

Outlook

- Many PIPS-IPM improvements (by ZIB) under way
- Additional HPC Solver Link to OOPS is “currently” under development
- Parallelization can be extended to 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)
  - Generation of separate model blocks as required by solver
    - Fully implemented by user: possible (significant refactorization of code)
    - Annotation provided by user → block sharp generation by GAMS: work in progress

Outlook

- Make sure we don’t work on a special case of a special case
- Model Experiment Partners
  - **Balmorel**, DTU, Denmark, http://www.balmorel.com/
  - **DIMENSION**, EWI Research and Scenarios, https://www.ewi.uni-koeln.de/en/models/dimension/
  - **PERSEUS**, Karlsruhe Institut für Technologie, https://www.energyplan.eu/othertools/global/perseus/
  - **WILMAR/JMM**, Universität Duisburg Essen, https://www.energyplan.eu/othertools/global/wilmar-planning-tool/
  - **TIMES**, Paul Scherrer Insitut (PSI), Switzerland, https://iea-etsap.org/index.php/etsap-tools/model-generators/times
  - **ELMOD**, Universität Dresden, https://tu-dresden.de/bu/wirtschaft/ee2/forschung/modelle/elmod
- Life After Research Project Ends
  - Siemens AG
  - Amazon’s EC2 Cloud