Audi
Vorsprung durch Technik

Reduction in simulation time and storage requirements using LoCo for SDM

Marko Thiele[2], Thomas Hambrecht[3], Harald Schluder[3]
Stefan Mertler[1], Clemens-August Thole[1],
Heiner Müllerschön[2], Martin Liebscher[2]

[1]SIDACT GmbH
[2]SCALE GmbH
[3]AUDI AG

SCALE

SIDACT GmbH

Agenda

Audi
Vorsprung durch Technik

- Introduction
- Reduction of simulation time by optimization of domain decomposition
- Reduction of
  - Input data by data deduplication
  - Result data using FEMZIP-E

SCALE

SIDACT GmbH

Introduction

Audi
Vorsprung durch Technik

- **SCALE GmbH...**
  - is a 100% subsidiary of DYNAmore with about 20 engineers and computer scientists
  - is dedicated to provide scalable IT-solutions for CAE data and process management
  - cooperates with SIDACT as a specialist for data compression
  - offers the following products which have been developed in the past 10 years in close cooperation with AUDI

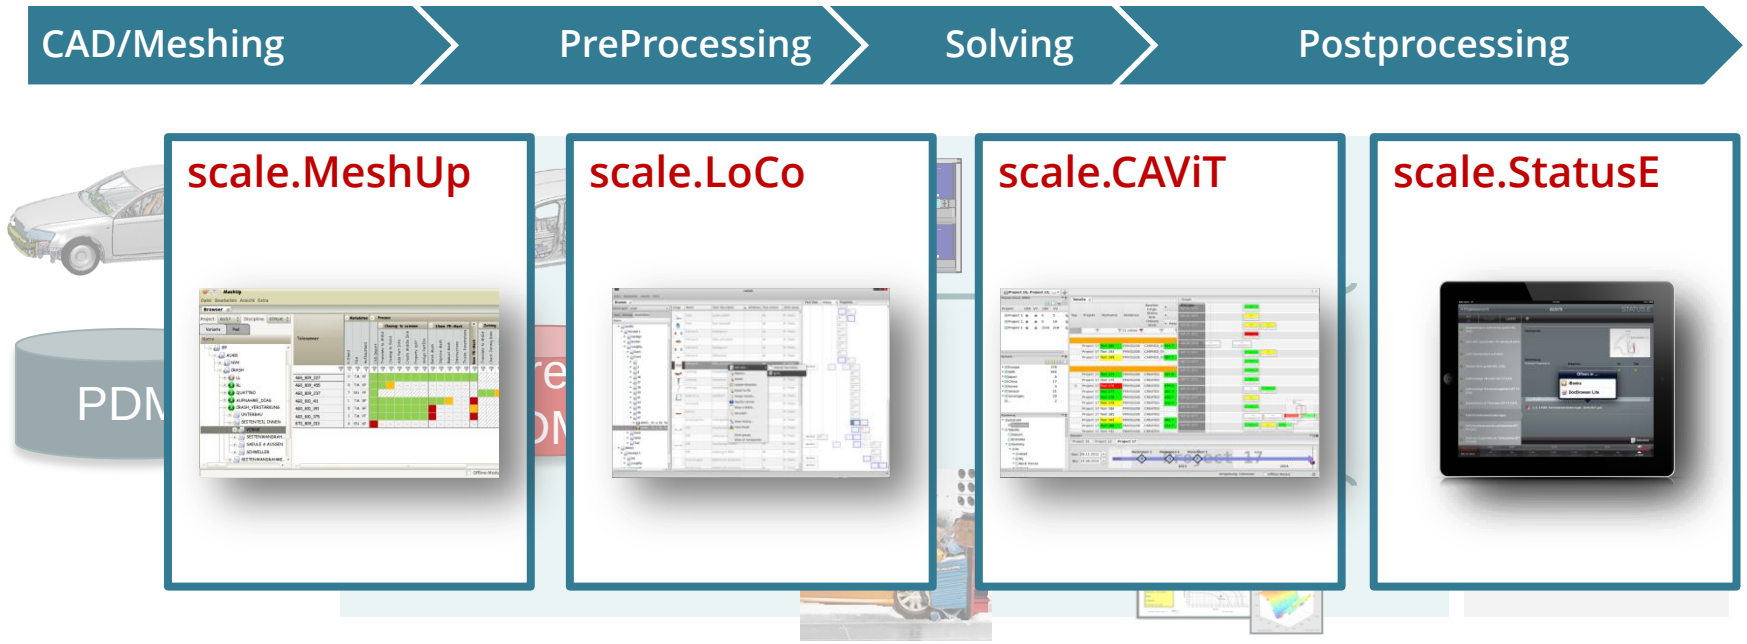

SCALE

SIDACT GmbH

Motivation

- **Increasing average model sizes**
  - average model size is still increasing
  - input data today partly exceed 1GB

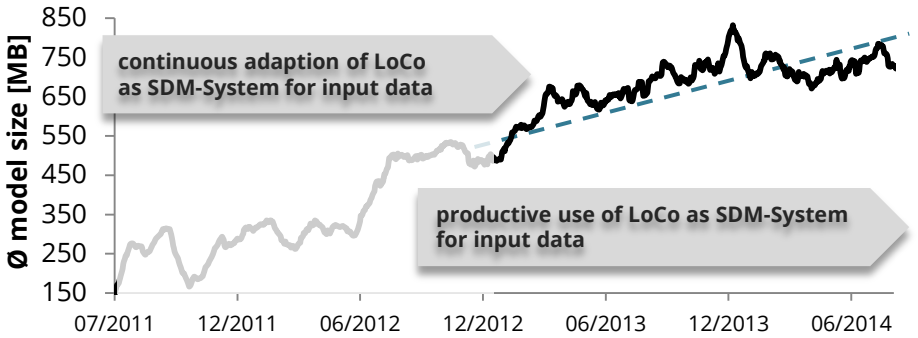

- **Increasing number of simulations**
  - more load cases
  - more vehicle models
  - more simulation disciplines
  - more ...

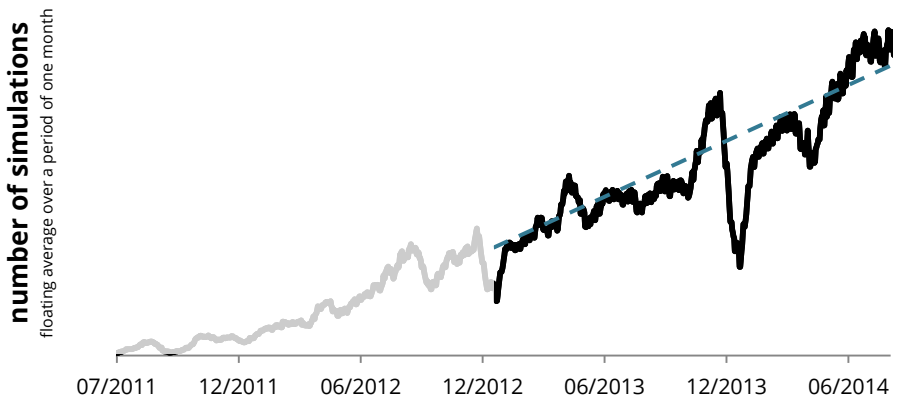

- **Increasing throughput of simulations per user**
  - Individual users are doing more simulations

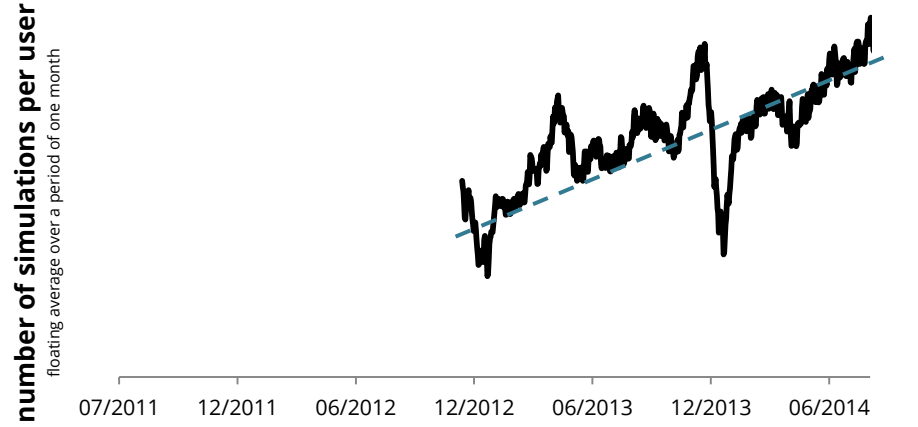

Approach

Audi
Vorsprung durch Technik

- **Typical workflow properties in SDM-System**
  - many small design changes
  - continuous evaluation of the same load cases
  - data and metadata of all simulations are stored in a central place
- **Goals**
  - harvesting information of completed simulations in order to optimize performance of future simulations
  - store only changes in model input and result data

Agenda

Audi
Vorsprung durch Technik

- Introduction
- **Reduction of simulation time by optimization of domain decomposition**
- Reduction of
  - Input data by data deduplication
  - Result data using FEMZIP-E

Optimization of Domain Decomposition

Audi
Vorsprung durch Technik

Problem

- simulation wall clock time is not scaling linear
- using more CPUs to compensate for increasing models is limited
- domain decomposition is based on estimates
- the calculation on all domains has to halt in order to wait for just one domain lacking behind

Goals

- monitoring of performance of all domains
- adjusting the domain size of new simulations according to the gathered profiling information

Optimization of Domain Decomposition

Audi
Vorsprung durch Technik

- **Investigations**
  - performance gain depends on load case and number of CPUs
  - Investigations have been performed using the same model without applying changes

- **Performance gain depending on load case**
  - 5 typical load cases have been investigated
  - 48CPUs have been used for each simulation
  - 10 simulations have been performed per iteration
  - 5 iterations have been performed for each load case
  - no changes to models between iterations

- **Performance gain depending on number of used CPUs**
  - 1 load case (front wall)
  - 16, 32, 64 and 128 numbers of CPUs
  - 5 iterations per setup
  - no changes to model between iterations
  - 10 runs for final iteration on each setup

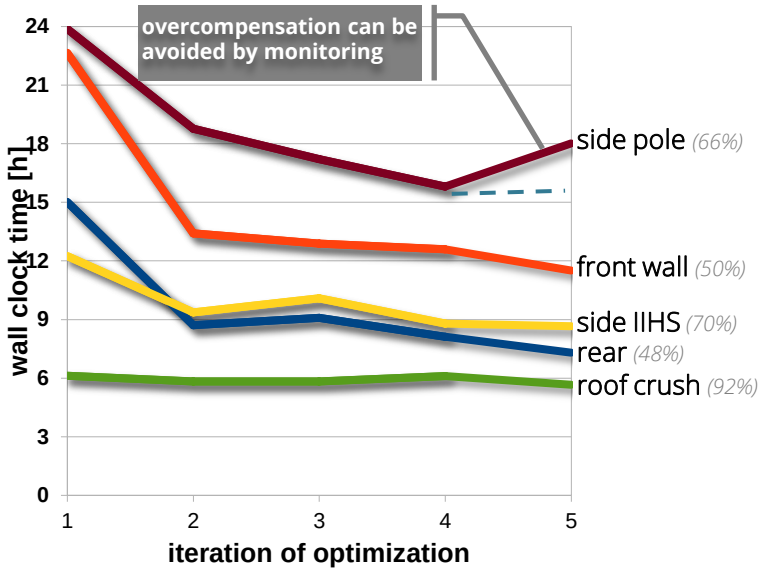

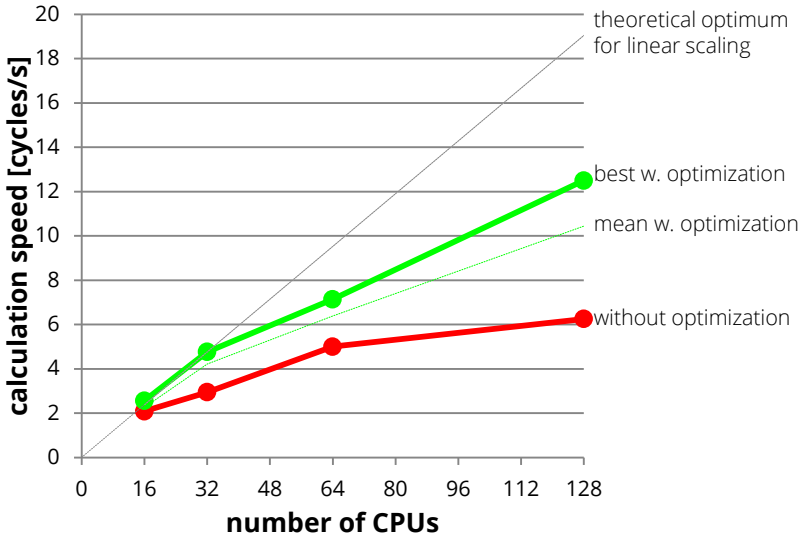

SCALE

SIDACT GmbH

Optimization of Domain Decomposition

Audi
Vorsprung durch Technik

- **Implementation in productive environment**
  - LoCo has been used as SDM-System to gather the profiling information of ongoing simulations and provide this information to new simulations in order to optimize domain decomposition
  - each simulation in LoCo may use the profiling information of one of its predecessors
  - overcompensation is avoided by monitoring the performance over multiple runs
  - all simulations started with LoCo use this approach by default, **no user interaction is required**
  - the difference between artificial tests and the productive environment is that in the productive environment there are always ongoing changes between calculations

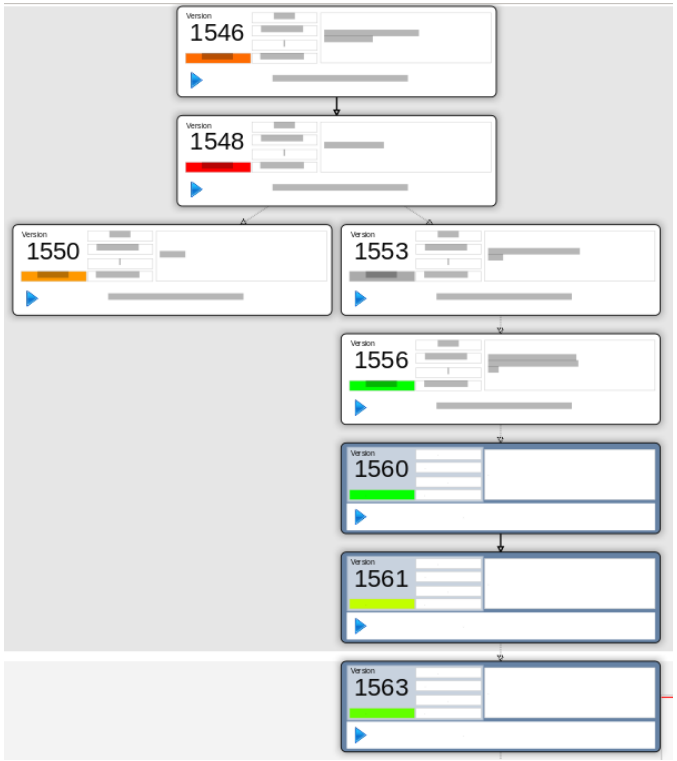

- **Results in productive environment**
  - 74 samples of original simulations of Q2/2014 have been recalculated without optimization
  - all simulations have been performed on 32CPUs
  - the overall performance gain has been **8%**

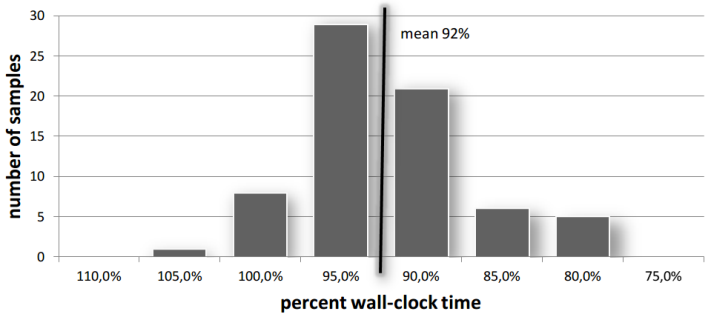

SCALE

SIDACT GmbH

Agenda

Audi
Vorsprung durch Technik

- Introduction
- Reduction of simulation time by optimization of domain decomposition
- **Reduction of**
  - Input data by data deduplication
  - Result data using FEMZIP-E

SCALE

SIDACT GmbH

Reducing Input Data

Audi
Vorsprung durch Technik

- **Motivation**
  - input data keeps increasing
  - models are best handled at the local workstation

    input data has to be transferred and stored
  - usually only small changes are applied for each simulations
- **Goal**
  - storing and transmitting only the changes that users apply
  - reducing the required data volume as much as possible
  - considering storage as well as bandwidth and transfer volume
- **Solution**
  - data deduplication
  - using the best available compression algorithms

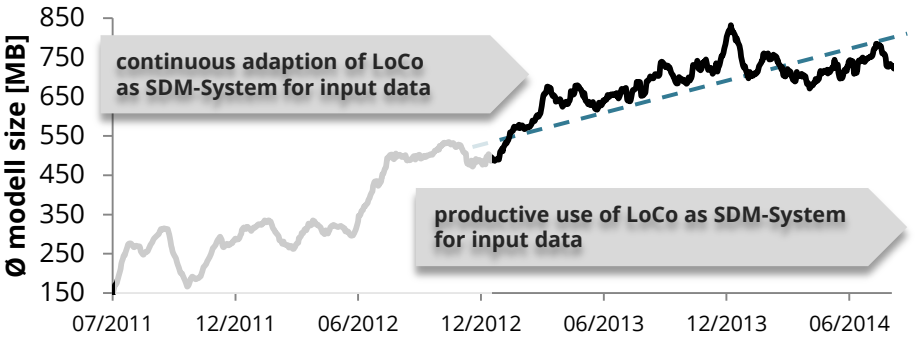

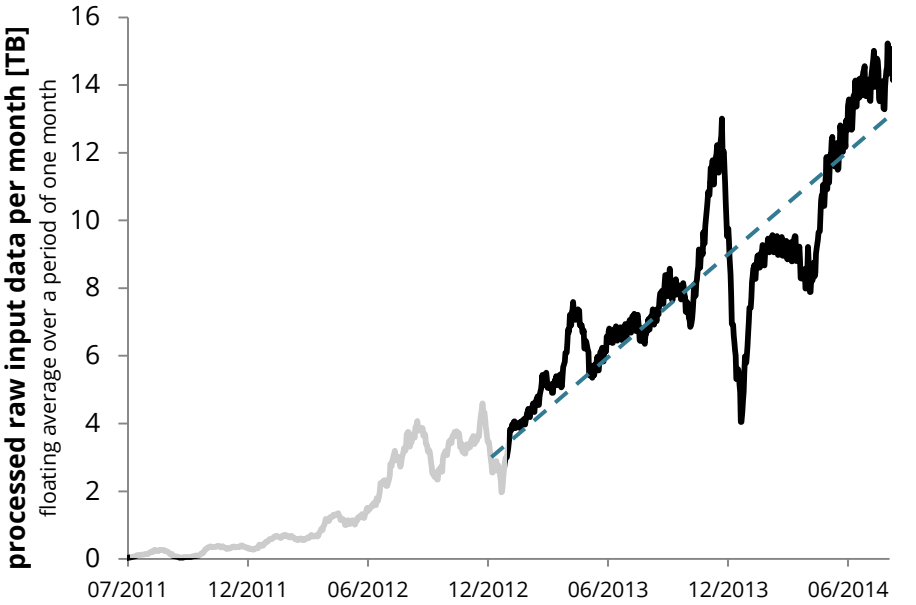

SCALE

SIDACT GmbH

Layers of Compression of Input Data

Audi
Vorsprung durch Technik

- **File level Data Deduplication**
  - each Simulation consists of multiple Files
  - changes for a simulation usually only affect a few files
  - only changed files are stored and transferred
  - savings approximately factor 20-25
  - **standard in LoCo**

- **Block level Data Deduplication**
  - changes on simulation input usually affect only a few lines
  - file is separated into blocks
  - only changed blocks are stored and transferred
  - savings approximately factor 8-10
  - **in development for LoCo** *(VAVID)*[1]

- **Standard compression algorithms**
  - simulation input files are usually ASCII
  - standard compression algorithms (e.g. zip, bzip, lzma) work best on ASCII data
  - savings approximately factor 3-4
  - **standard in LoCo**

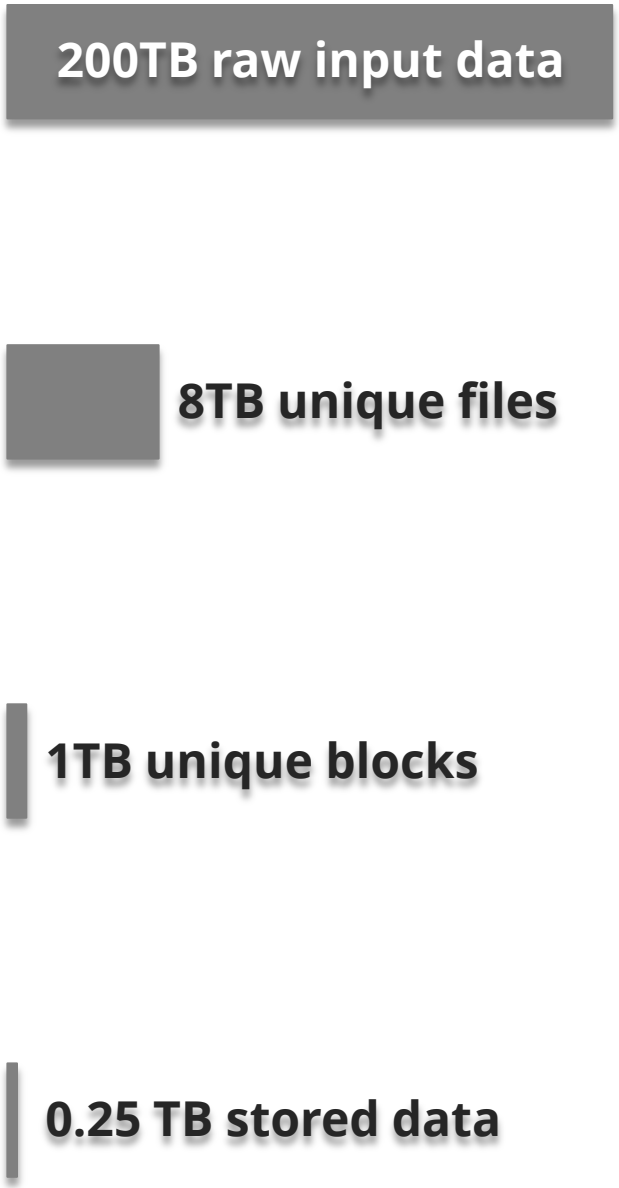

SCALE

SIDACT GmbH

[1] VAVID - BMBF Big Data research program, http://www.pt-it.pt-dlr.de/de/3138.php

Block Level Data Deduplication *(in development)*

Audi Vorsprung durch Technik

- initial file

`LoCo_speichert_nur_das_was_nötig_ist.`

Block A: `LoCo_spe`
Block B: `ichert_n`
Block C: `ur_das_w`
Block D: `as_nötig`
Block E: `_ist.`

include consists of blocks: `A` `B` `C` `D` `E`

5 + 37 = 42 characters

- changed file

`LoCo_speichert_nur_das_was_geändert_ist.`

Block F: `as_geänd`
Block G: `ert`

Include consists of blocks: `A` `B` `C` `E` `F` `G`

6 + 11 = 17 characters

Block Level Data Deduplication *(in development)*

Audi Vorsprung durch Technik

Estimates based on Test Data

	[GB]	[%]
original LoCo vault data for one project	40,5	100%
raw data *(vault decompressed)*	157,2	388%
data deduplication *(without compression)*	17,1	42%
compression gzip	5,1	12%
compression lzma	3,2	8%

Problems yet to be solved

- block index can become very large *(high demand on RAM)*
- deduplication has to be done client side *(reducing transfer volume)*
- very high performance requirements for request rates to block index *(esp. server side)*
- performance for file reconstruction is curtail
- permanently deleting individual data sets is challenging

Agenda

Audi
Vorsprung durch Technik

- Introduction
- Reduction of simulation time by optimization of domain decomposition
- **Reduction of**
  - Input data by data deduplication
  - Result data using FEMZIP-E

SCALE

SIDACT GmbH

Reducing Result Data

Audi
Vorsprung durch Technik

Motivation

- Simulation results often contain data of various different variables.
  - The more variables are used in the output the bigger becomes the final result file.
- Due to generally small changes in models the results often contain comparable data
  - Reduce data by storing just the differences between new and known results

Problems

- With the increasing number of load cases it becomes difficult to decide witch variables to consider for the output files
- Taking advantage of similarities between runs is a challenging mathematical task
  - Result data contains chaotic components
  - Unchanged parts of the models might behave similar but not identical

Solution

- Reducing the used variables in the output files to the really required ones for the specific load case and task will reduce the overall required storage.
- Applying FEMZIP-E by SIDACT in order to take advantage of similarities between runs.

Result Data *(overall measures to reduce result data)*

Audi
Vorsprung durch Technik

reduce output

- **specify necessary output in header**
  - output only relevant data from solver
  - specify relevant output for each use case
  - use features such as selective output
- **delete unnecessary output from results**
  - use tools from solver to delete unwanted data
  - use e.g. Animator 4 to cut relevant parts of the model

compress output

- **use FEMZIP**
- **tune regular FEMZIP parameter**
  - use reasonable accuracy
  - delete unnecessary functions from result files

FEMZIP-E

- **Use previous results and store only newly added data**

~ 2x (?) ~ 2x (?) ~ 3x

measures with LoCo

- **specify output individually by**
  - discipline
  - load case
  - ..
- **use scripts in LoCo to reduce data**
  - Animator session files for each load case to cut relevant parts of the model

measures with LoCo

- **specify FEMZIP parameter file individually by**
  - discipline
  - load case
  - ...

measures with LoCo

- **find correct predecessor to store new results**
- **retrieve necessary files to restore individual results**
- **delete old data by means of access time and access count**

SCALE

SIDACT GmbH

Result Data *(storage architecture, concept for deleting data)*

Audi Vorsprung durch Technik

- **storing result data with FEMZIP-E**
  - each simulation produces two Files
    - **FDB or UFDB:** Database of properties that can be used throughout multiple simulations
    - **EFZ:** Additional Information to restore individual result
  - result data of each Simulation depends on FDB and UFDB files of predecessors

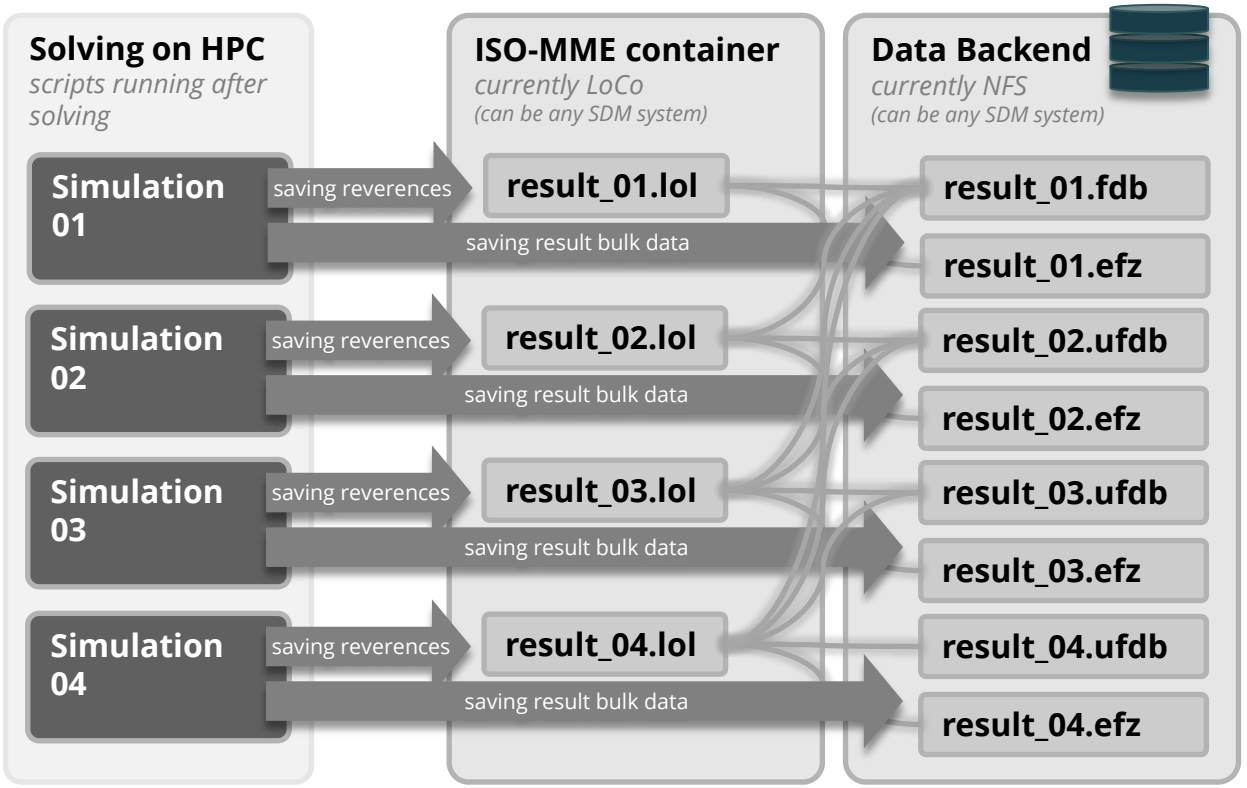

SCALE

SIDACT GmbH

Result Data *(storage architecture, concept for deleting data)*

Audi Vorsprung durch Technik

- **deleting individual Data from FEMZIP-E**
  - **~85% of the data:** EFZ files can be deleted individually
  - **~15% of the data:** FDB files and UFDB files have to be kept until all results of a chain are deleted

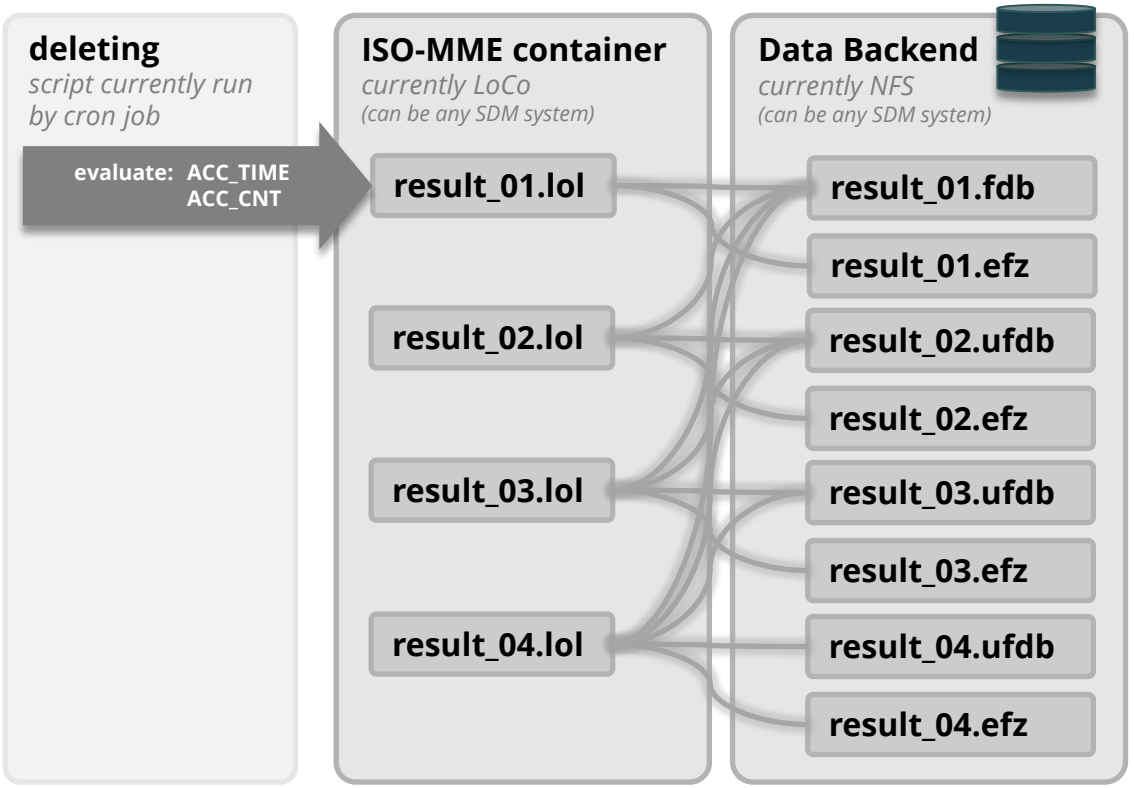

SCALE

SIDACT GmbH

Result Data *(FEMZIP-E general principles and workflow)*

Audi
Vorsprung durch Technik

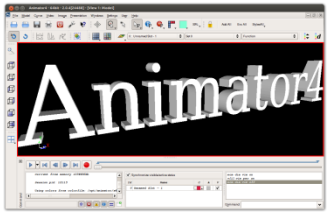

NFS Share

*.fdb *.ufdb *.efz
*.fdb *.ufdb *.efz
*.fdb *.ufdb *.efz

request missing parts of FEMZIP-E DB

LoCo-Client

open

Open Result
decompression

*.DSY.fz

FEMUnZip-E

*.fdb *.ufdb *.efz

RemoteSubmit

LoCo Server

Input Files

Assembly

Solver Input

lol-File
(referenced FEMZIP-E files)

HPC / LoCo-Clusterskripte

Job assembly with LoCo

gathering of required Files from storage

*.fdb *.ufdb

(in parallel to running simulation with low bandwidth)

Solving *(PAM-Crash)*

FEMZIP-E *compression*

old *.fdb, *.ufdb

new *.ufdb

new *.efz

new Result *.DSY

*.efz *.ufdb

SCALE

SIDACT GmbH

Result Data *(compression results)*

Audi
Vorsprung durch Technik

Test Data

- one load case *(front wall)*
- 155 Results
- compressed with respect to history/order of creation
- Results are extracted from productive environment and have been created over a period of approximately 6 month

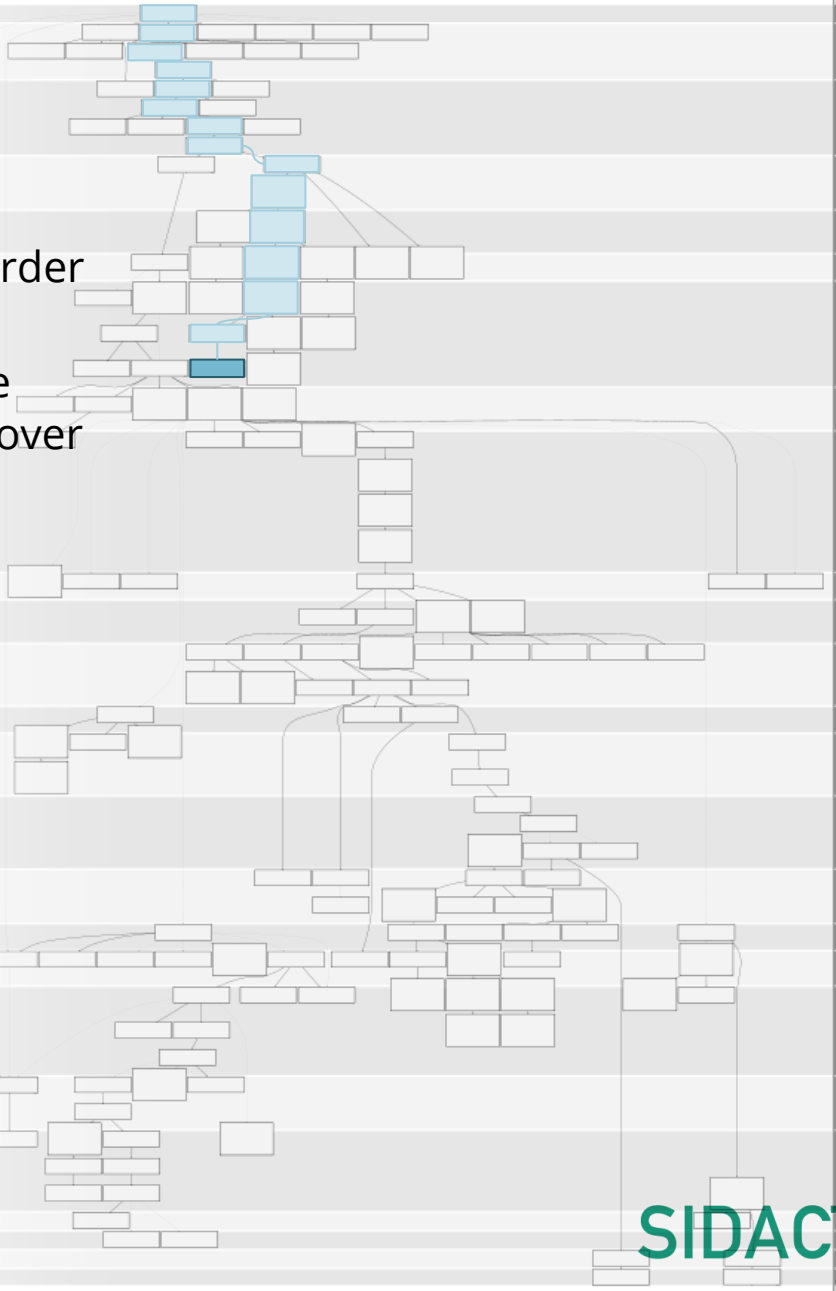

Results

FEMZIP-P		13,95 GB	
FEMZIP-E	FDB	0,85 GB	15%
	EFZ	5,00 GB	85%
	Total	5,85 GB	
FEMZIP-E + gzip		4,74 GB	

Factor P/E	**2,94***

* higher compression rates of up to factor 4 could be achieved when compressing all files at once

SCALE

SIDACT GmbH

Result Data *(concept for deleting data)*

Audi
Vorsprung durch Technik

deleting individual Data from FEMZIP-E

- **~85% of the data:** EFZ files can be deleted individually
- **~15% of the data:** FDB files and UFDB files have to be kept until all results of a chain are deleted
- parts of chains can be deleted individually
- length of chain is limited *(e.g. < 20 results)*

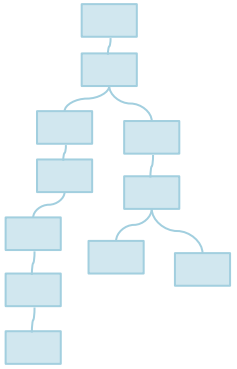

advanced scheme for deletion

- Constraints
  - MAX_KEEP *always delete, e.g. > 6 month*
  - MIN_KEEP *never delete, e.g. < 1 week*
  - MAX_STORAGE *delete until smaller than, e.g. 500GB*
- Values to be considered
  - ACCESS_TIME *last time result has been accessed*
  - ACCESS_COUNT *number of times result has been accessed*

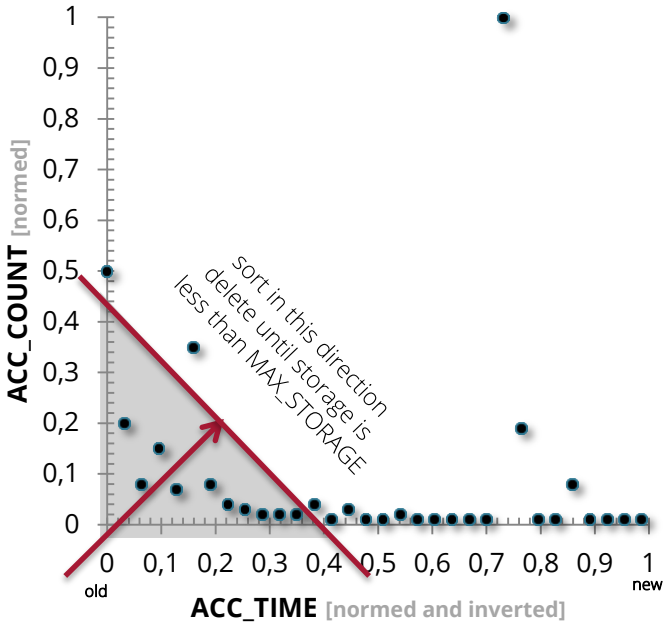

SCALE

SIDACT GmbH

Summary

Audi
Vorsprung durch Technik

- **Calculation Time**
  - can be reduced by optimizing domain decomposition
  - already implemented and proven to work in productive environment
- **Reducing Input Data**
  - Multiple levels of eliminating redundancy
    - file level data deduplication
    - block level data deduplication *(in development)*
    - standard compression algorithms
- **Reducing Output Data**
  - intelligent mechanisms to determine what really needs to be stored
  - FEMZIP-E by SIDACT *(in development)*

Audi
Vorsprung durch Technik

Vielen Dank!

SCALE

SIDACT GmbH