13/09/2017

FULL-SCALE SIMULATION OF AN INDUSTRIAL STEAM CRACKING FURNACE USING OPENFOAM ON TIER-1

Pieter A. Reyniers, Laurien A. Vandewalle, Kevin M. Van Geem, Guy B. Marin

Laboratory for Chemical Technology (LCT) – Ghent University
laurien.vandewalle@ugent.be
pieter.reyniers@ugent.be

GHENT UNIVERSITY

LCT DRIVING CHEMICAL TECHNOLOGY

MULTISCALE MODELING APPROACH

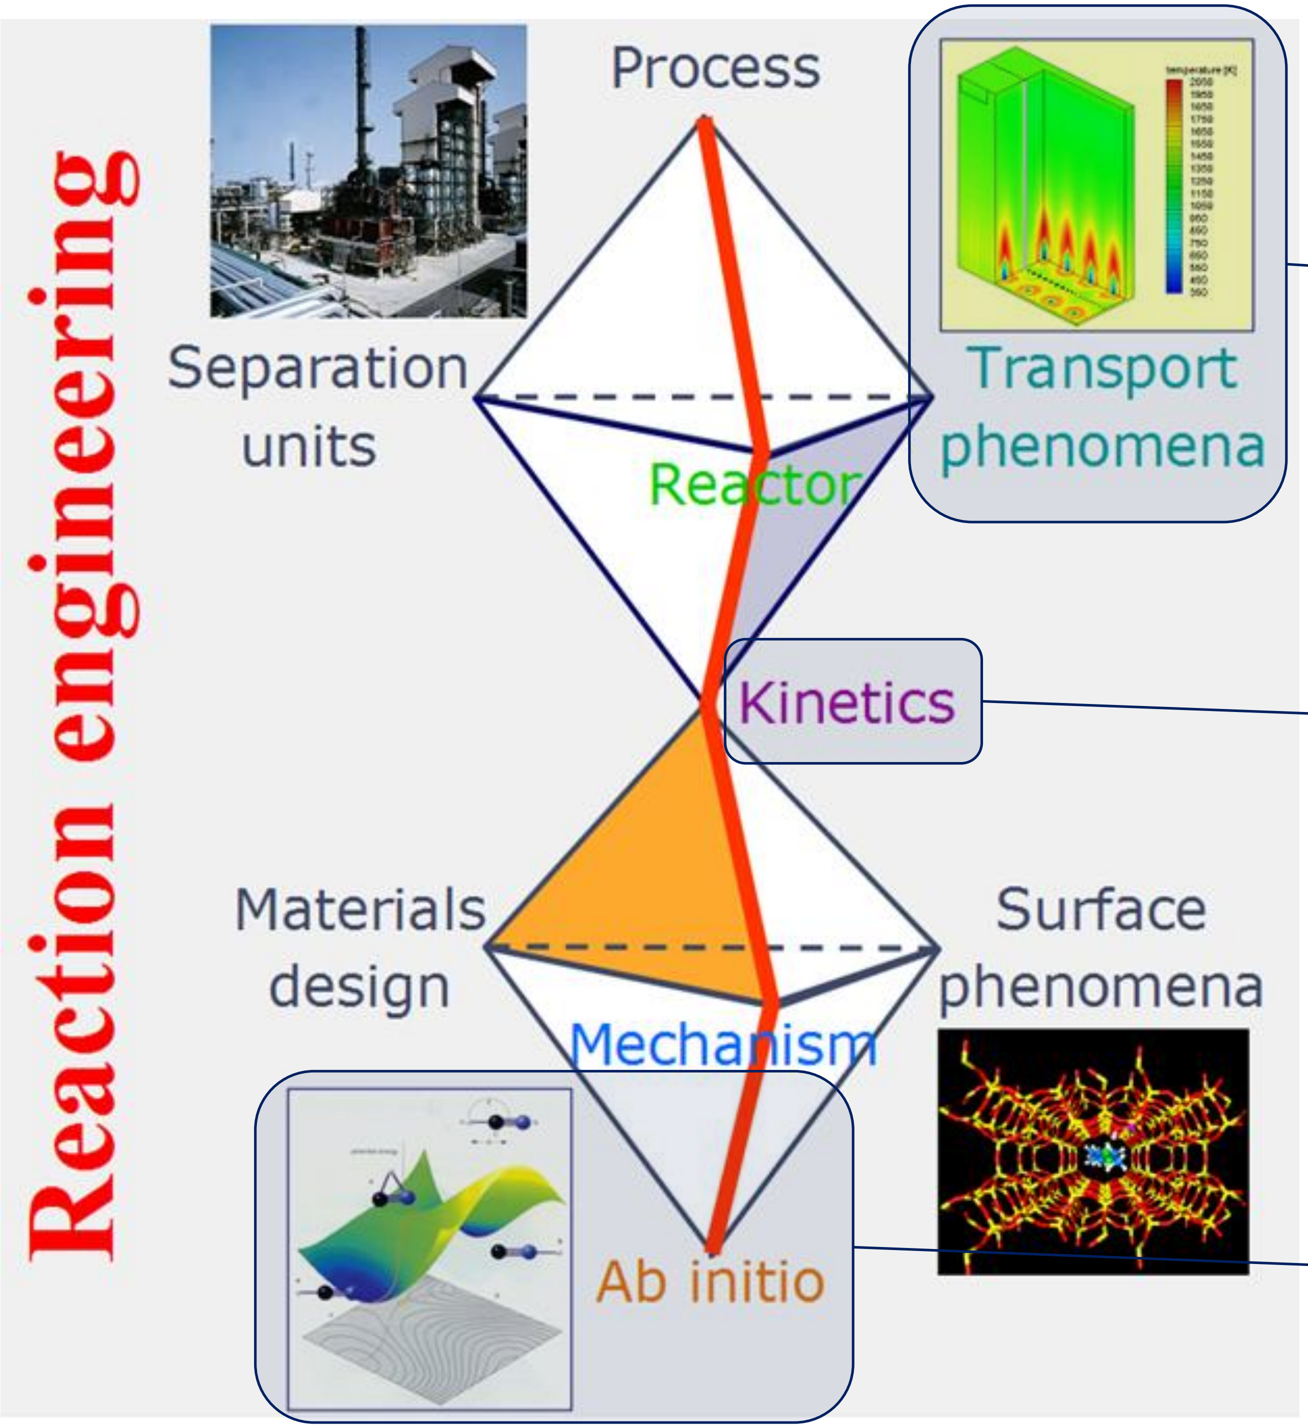

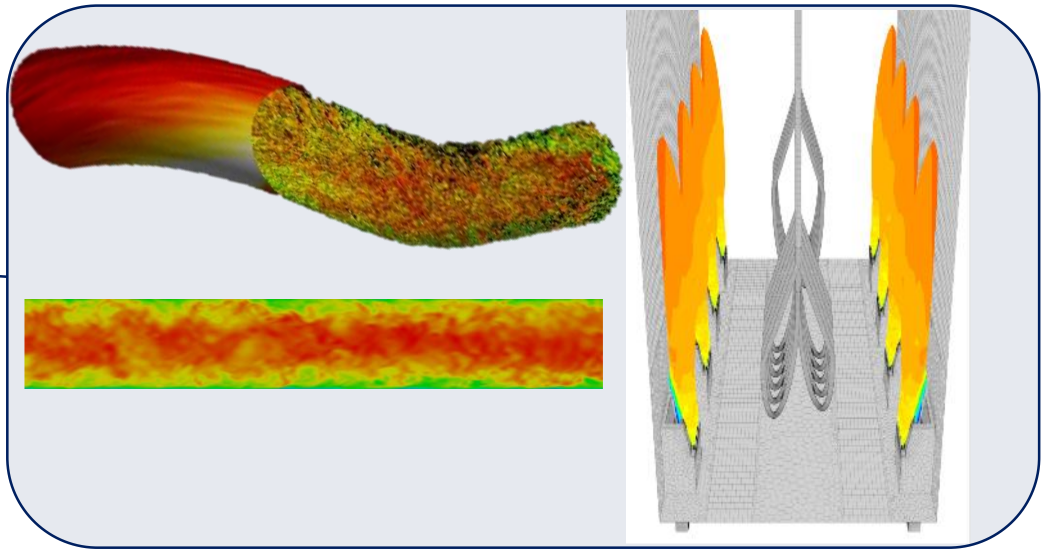

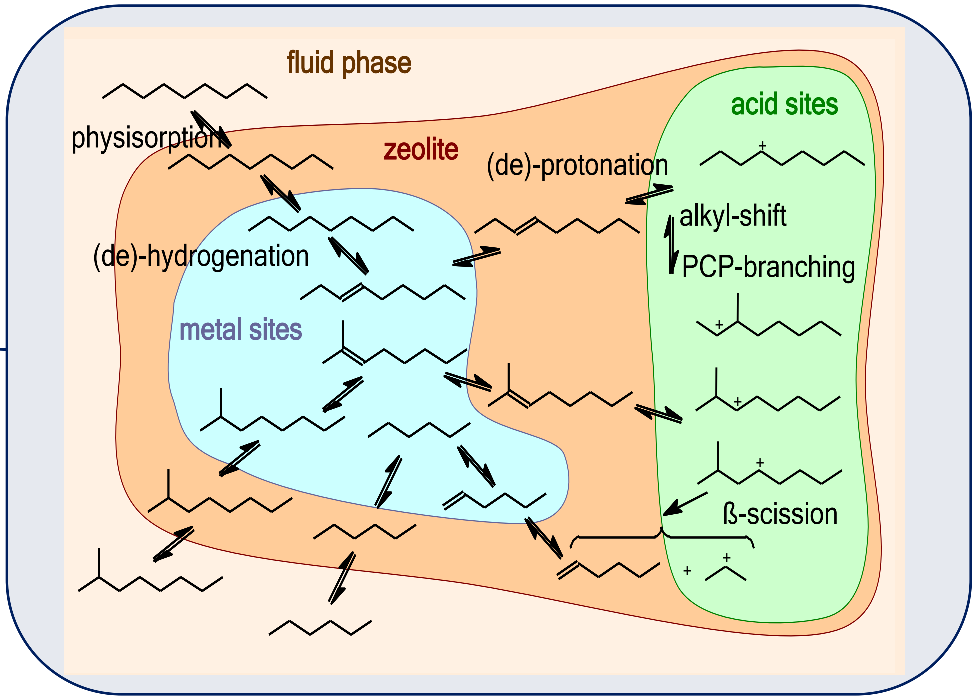

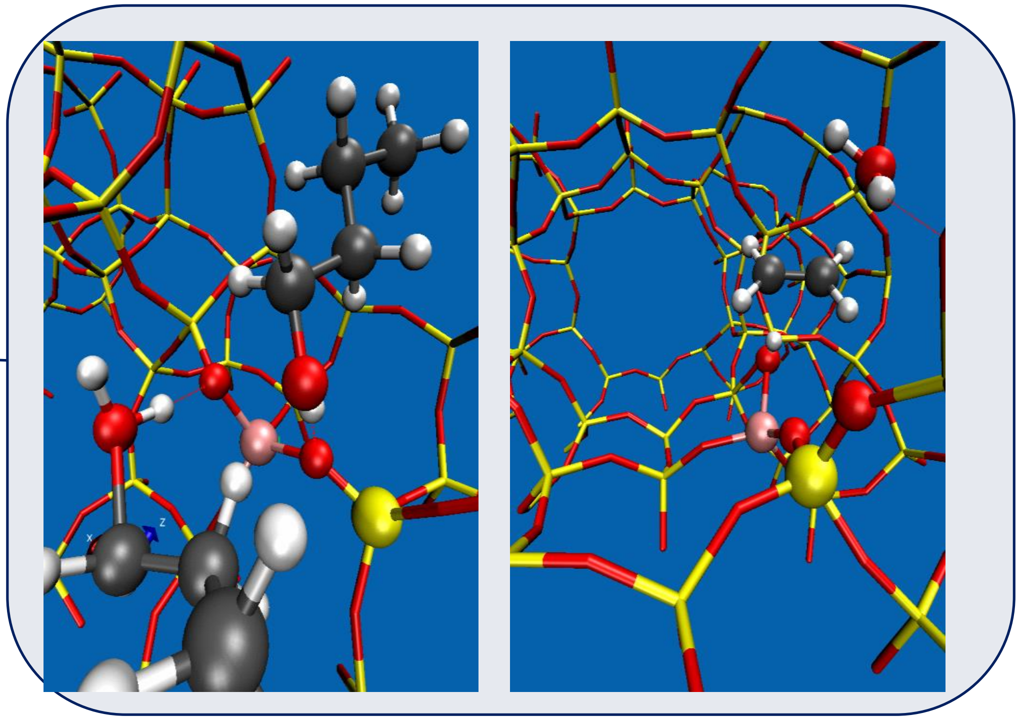

GHENT UNIVERSITY

LCT DRIVING CHEMICAL TECHNOLOGY

STEAM CRACKING

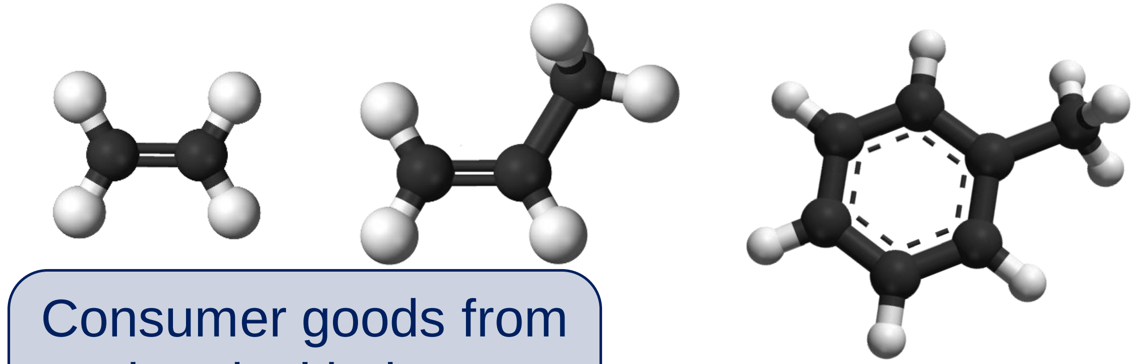

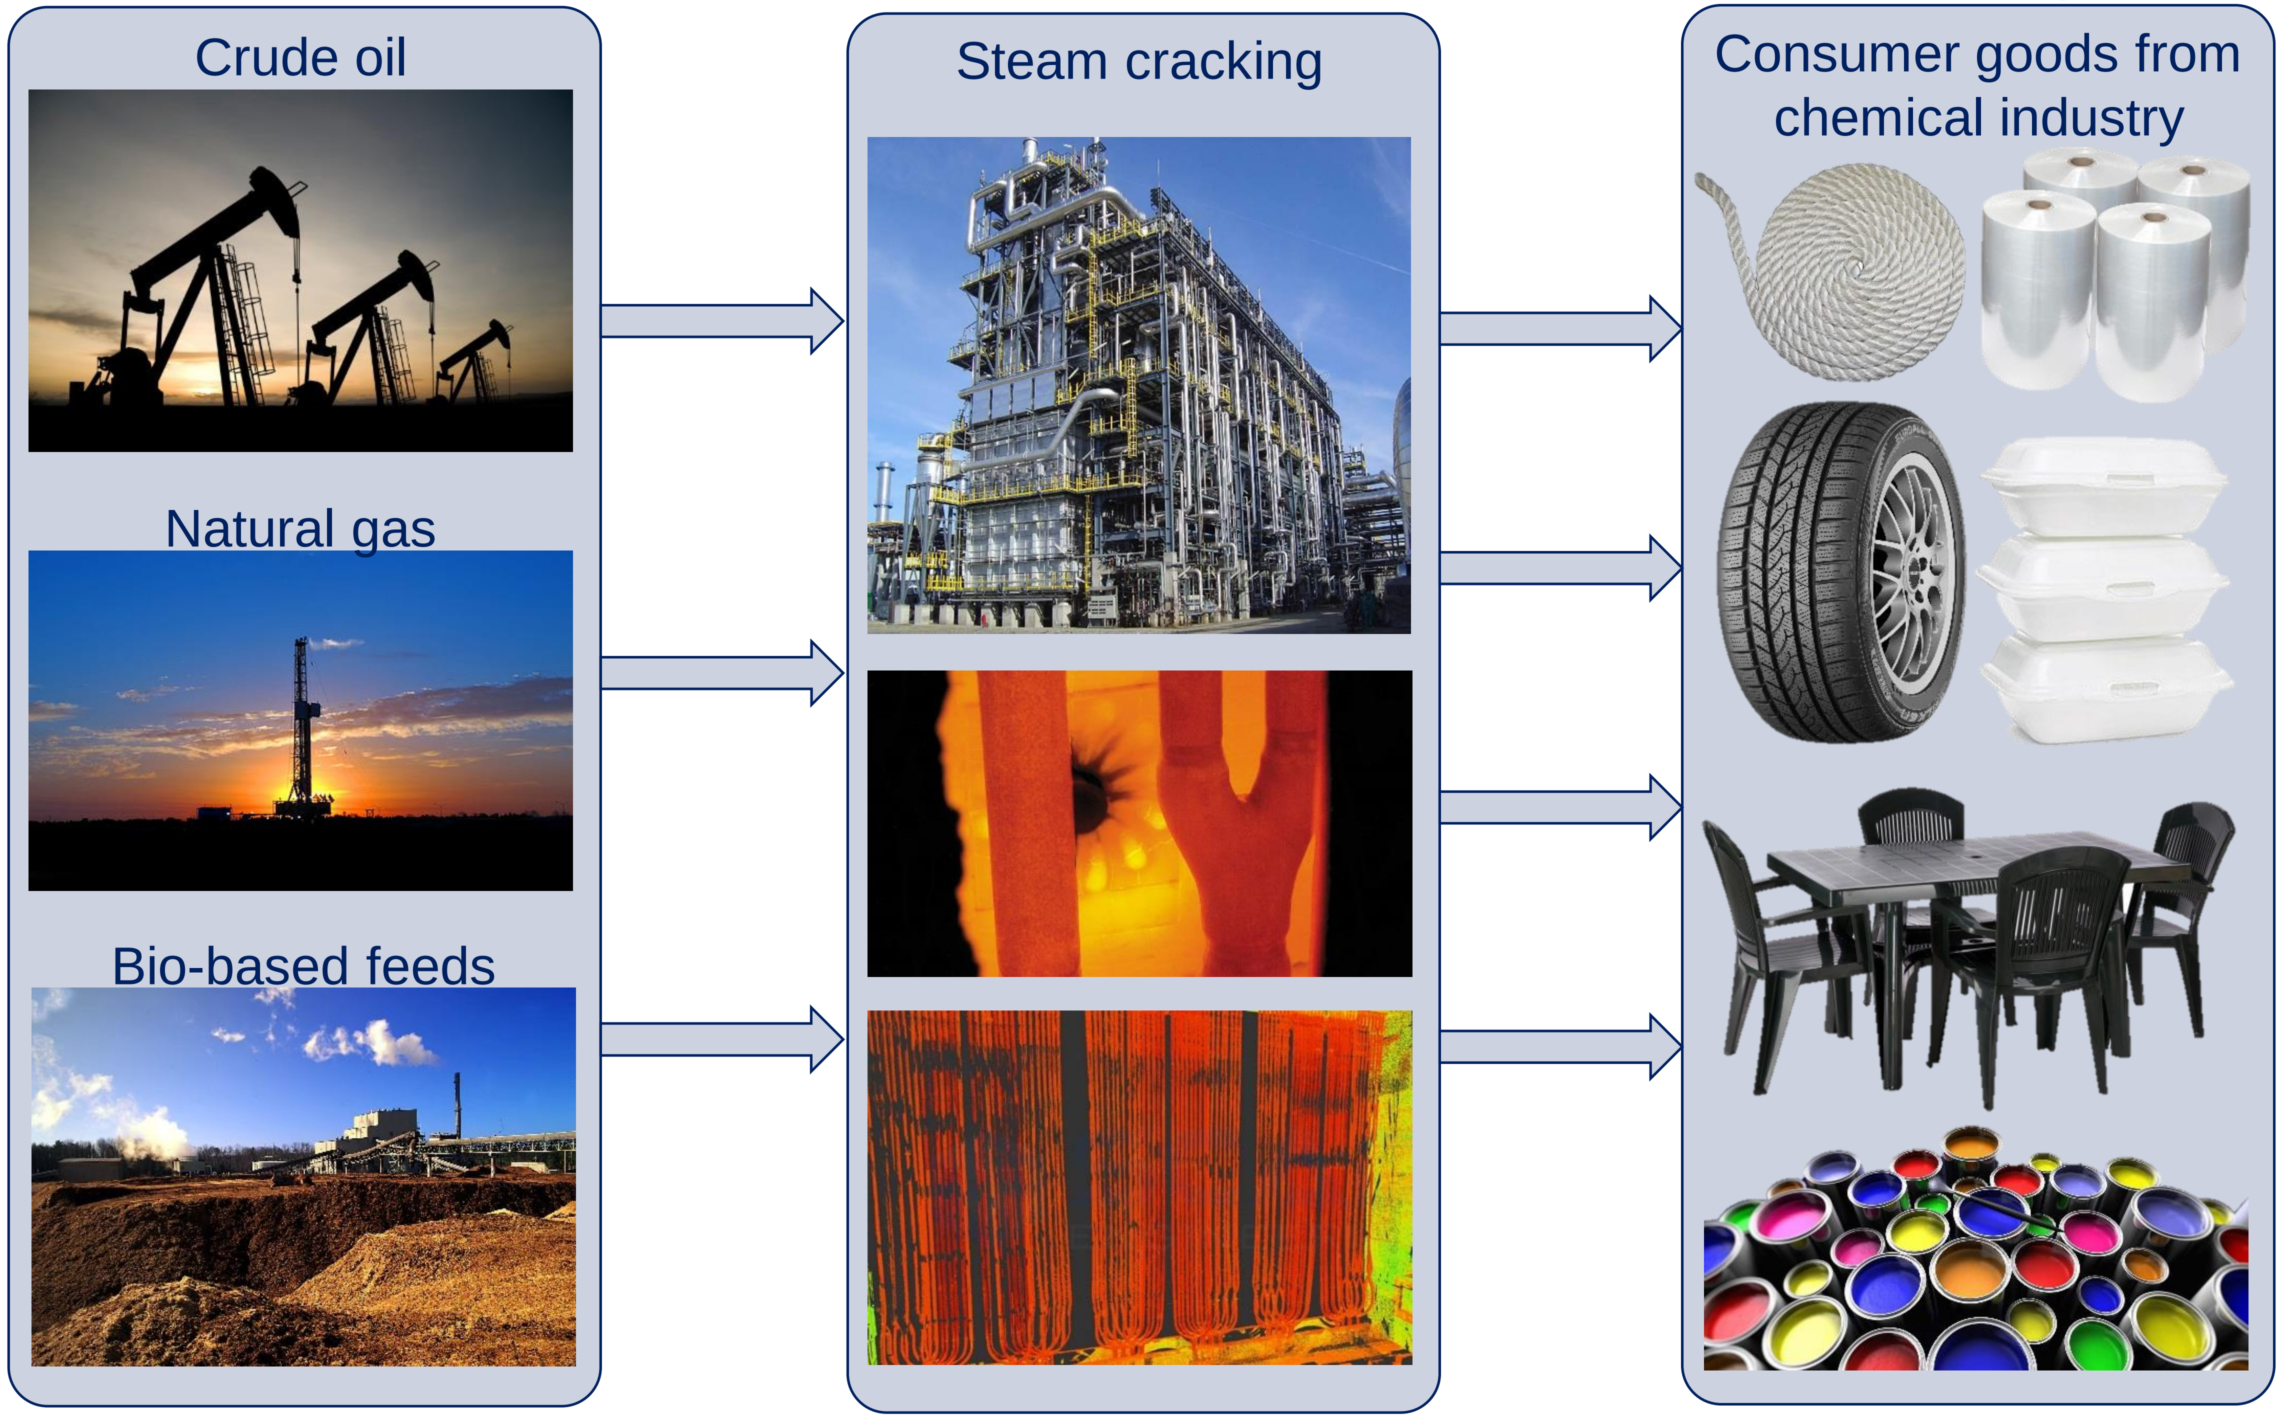

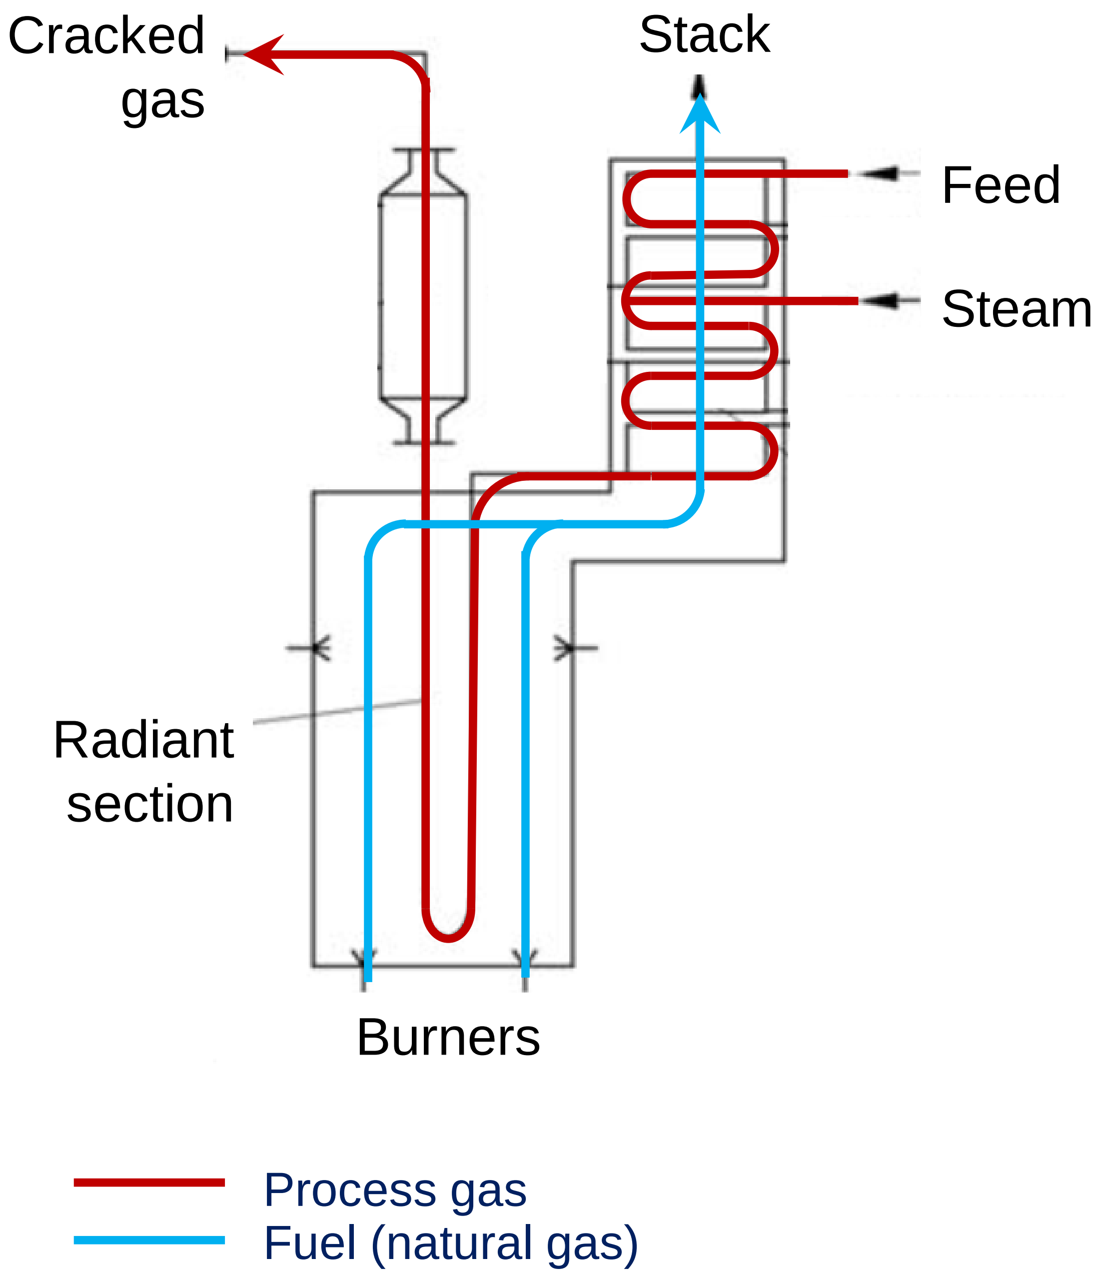

atozforex.com; pnnl.org; districtenergy.org; scade.fr; schmidt-clemens.de; Linde Group; IHS Chemical Insight

COKE FORMATION IN STEAM CRACKING

Endothermic process at temperatures of 800–900 °C

Deposition of a carbon layer on the reactor surface

- Reduced thermal efficiency
- High pressure causes loss of product selectivity
- Coil carburization and thermal stress

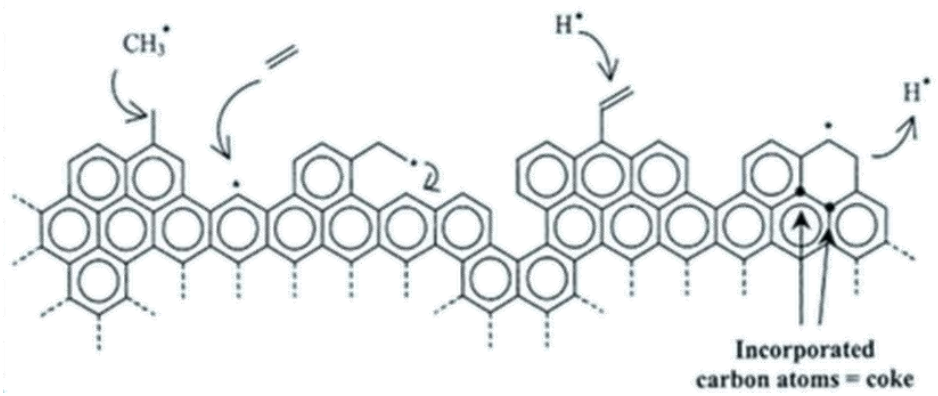

$$r_C = \sum_i c_i \cdot A_i \cdot \exp\left(\frac{-E_{a,i}}{RT_{int}}\right)$$

Coil cracking due to differences in thermal expansion rate

Hot spots due to inhomogeneous coke formation

Coke reduction method: 3D reactor technology

Nova Chemicals, 2002; Linde Group; Muños et al., 2013; Albright et al., 1988; Muños et al., 2014

GHENT UNIVERSITY

LCT DRIVING CHEMICAL TECHNOLOGY

CHEMICAL ENGINEERING IN 21ST CENTURY

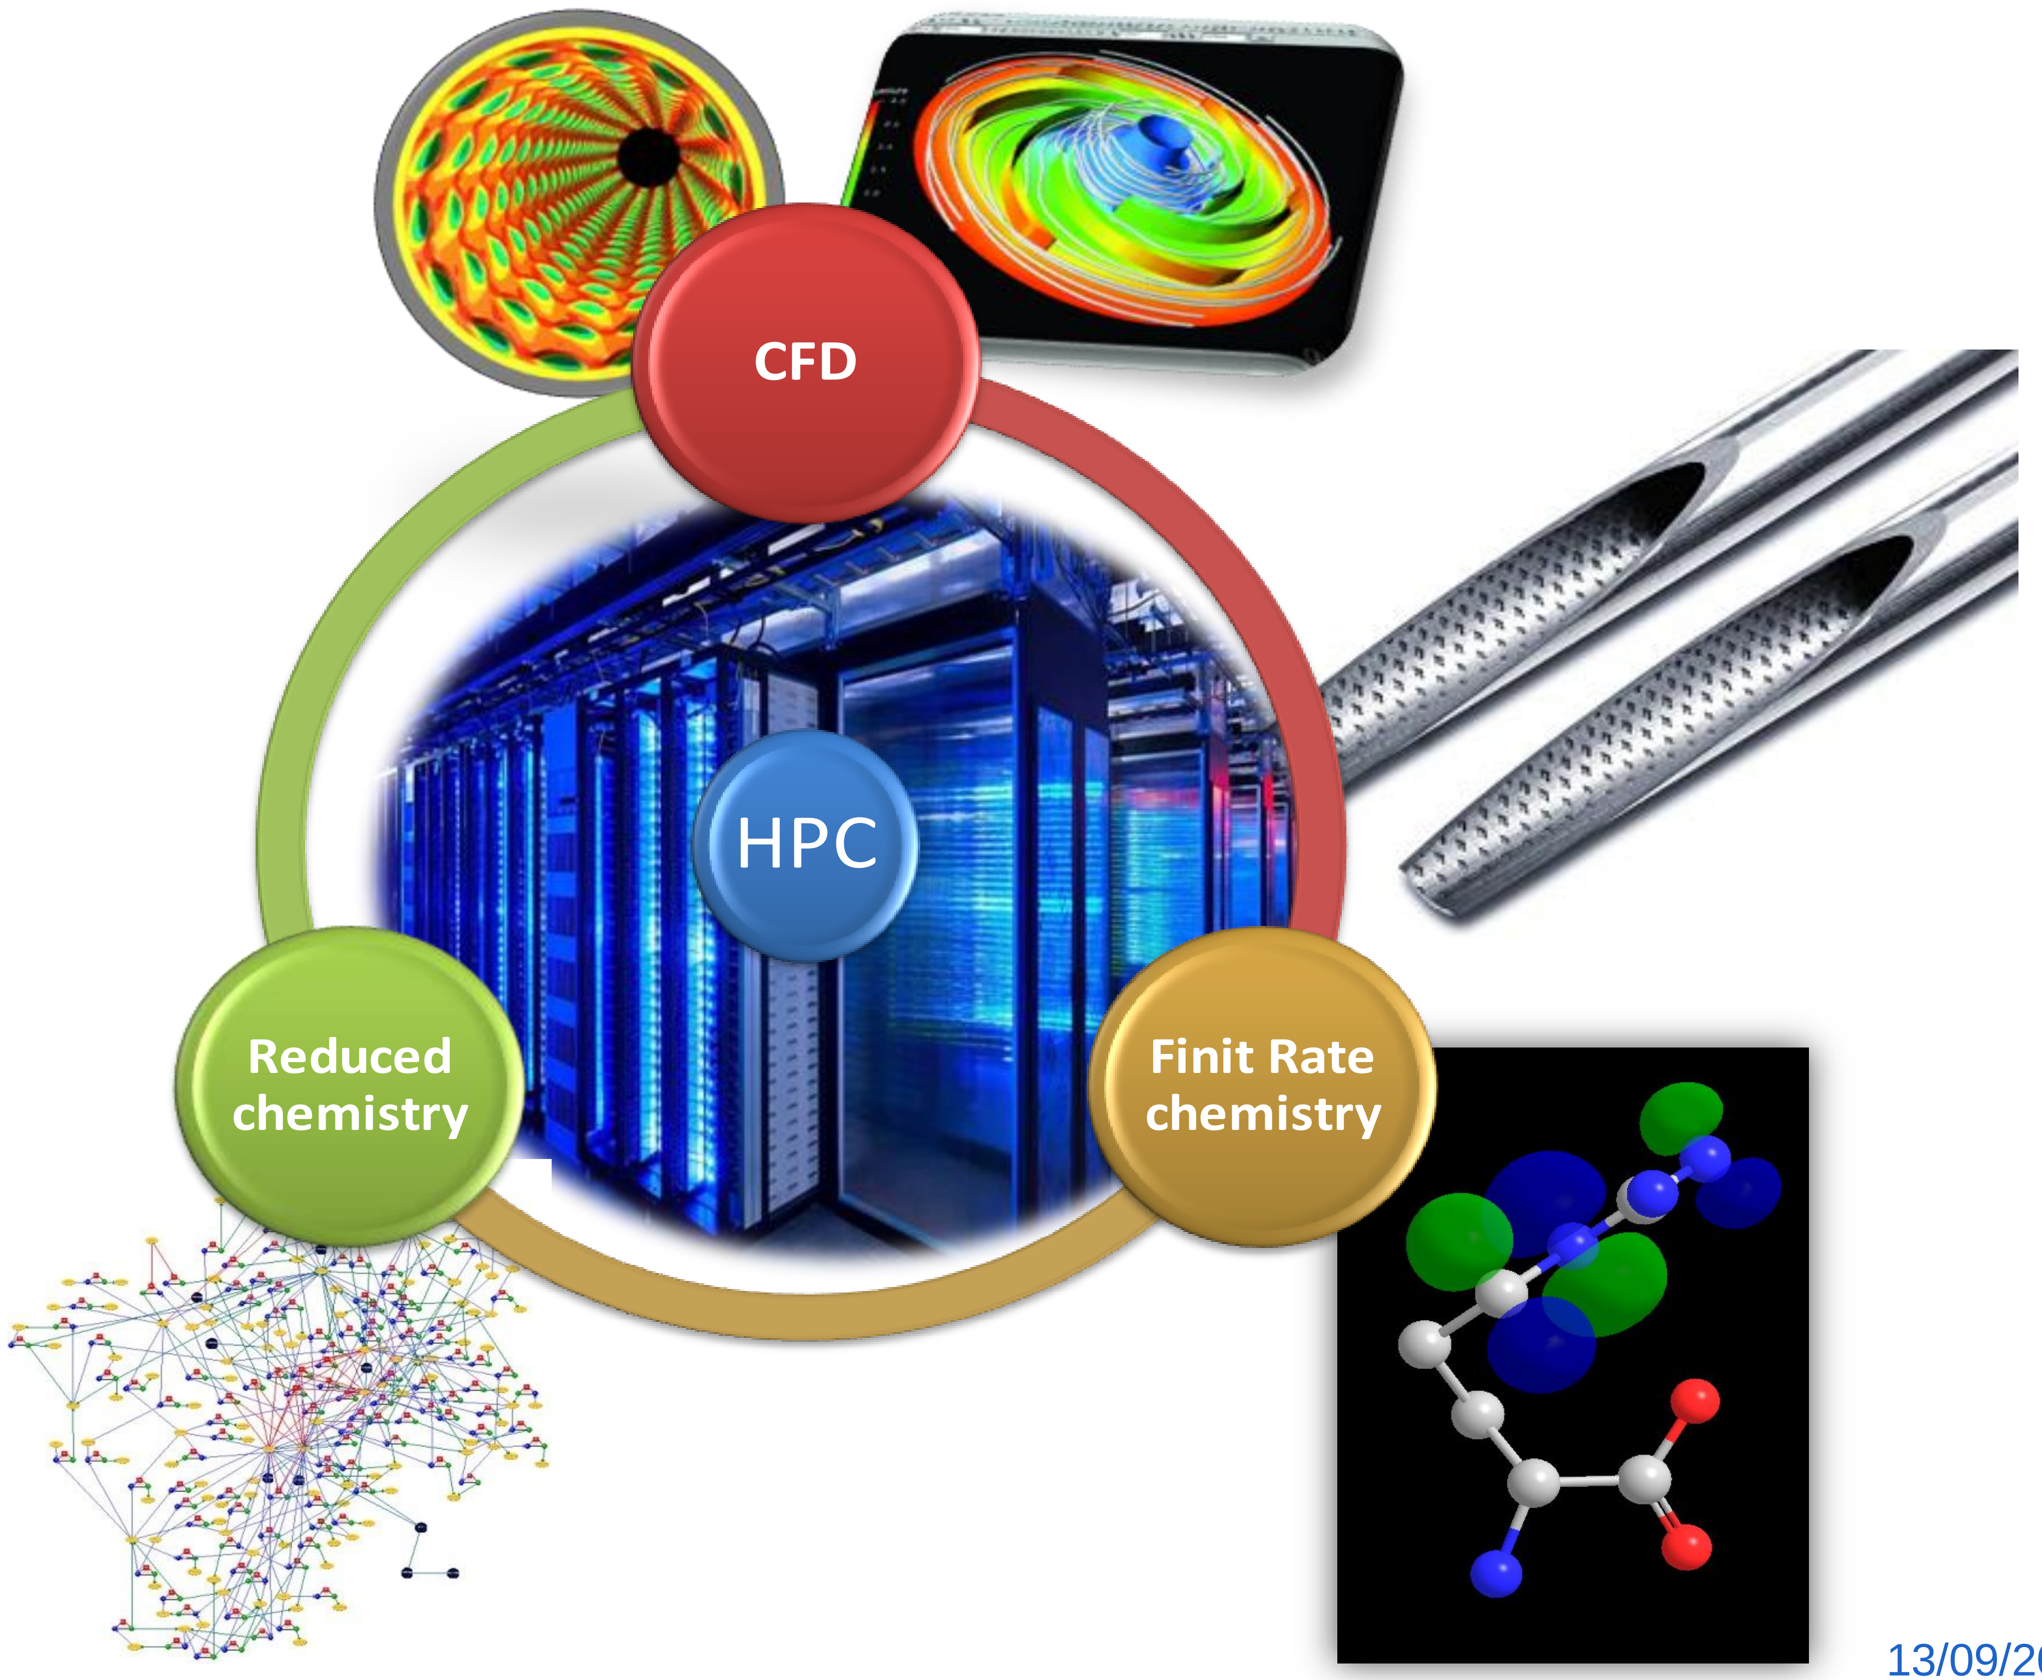

GHENT UNIVERSITY

LCT DRIVING CHEMICAL TECHNOLOGY

ENHANCED REACTOR DESIGNS

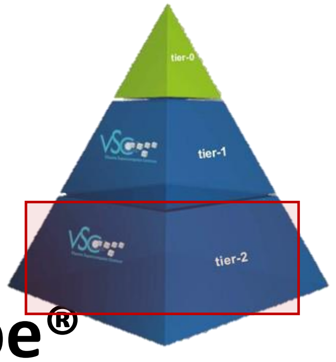

KBR / S+C **Straight / rifled fins**	Kubota **(Slit- / X-) MERT®**	Lummus/Sinopec **IHT®**	Technip **Swirl Flow Tube®**
1988, 2002, 2011	1996, 2003, 2009	2009	2011

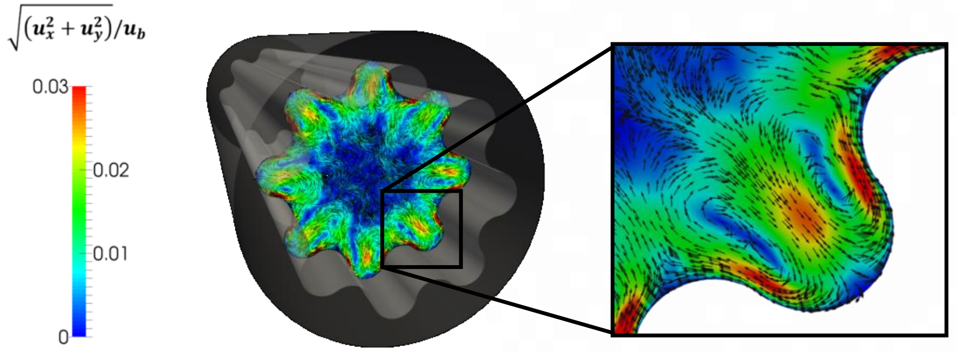

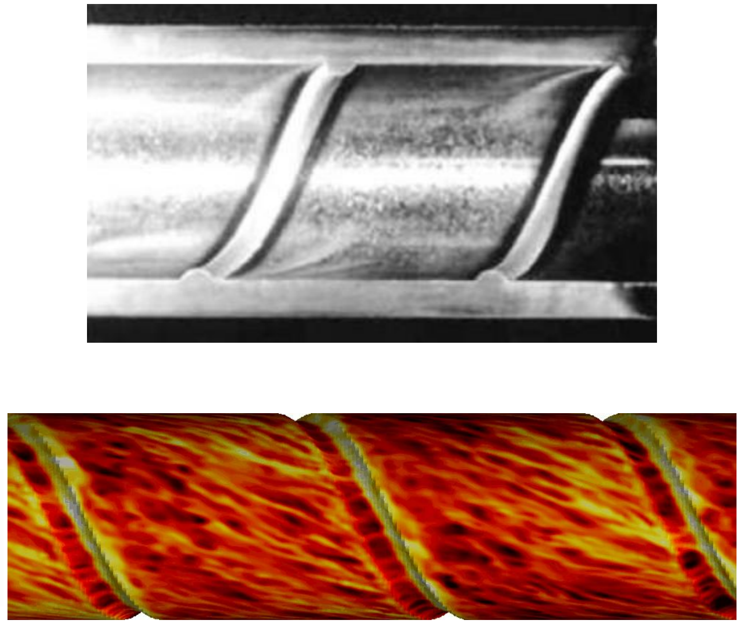

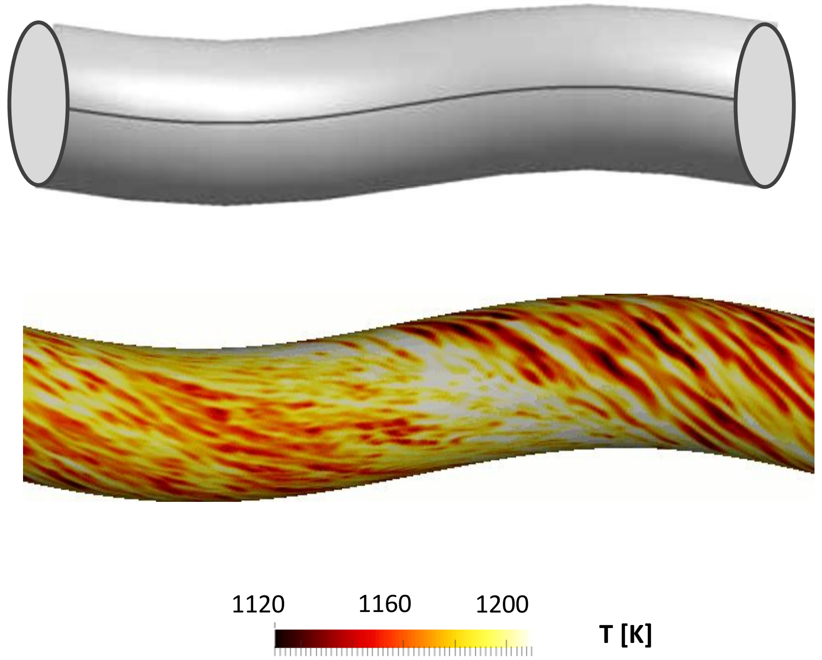

(*) Van Cauwenberge et al. 2015, Chem. Eng. J.

FURNACE MODELLING

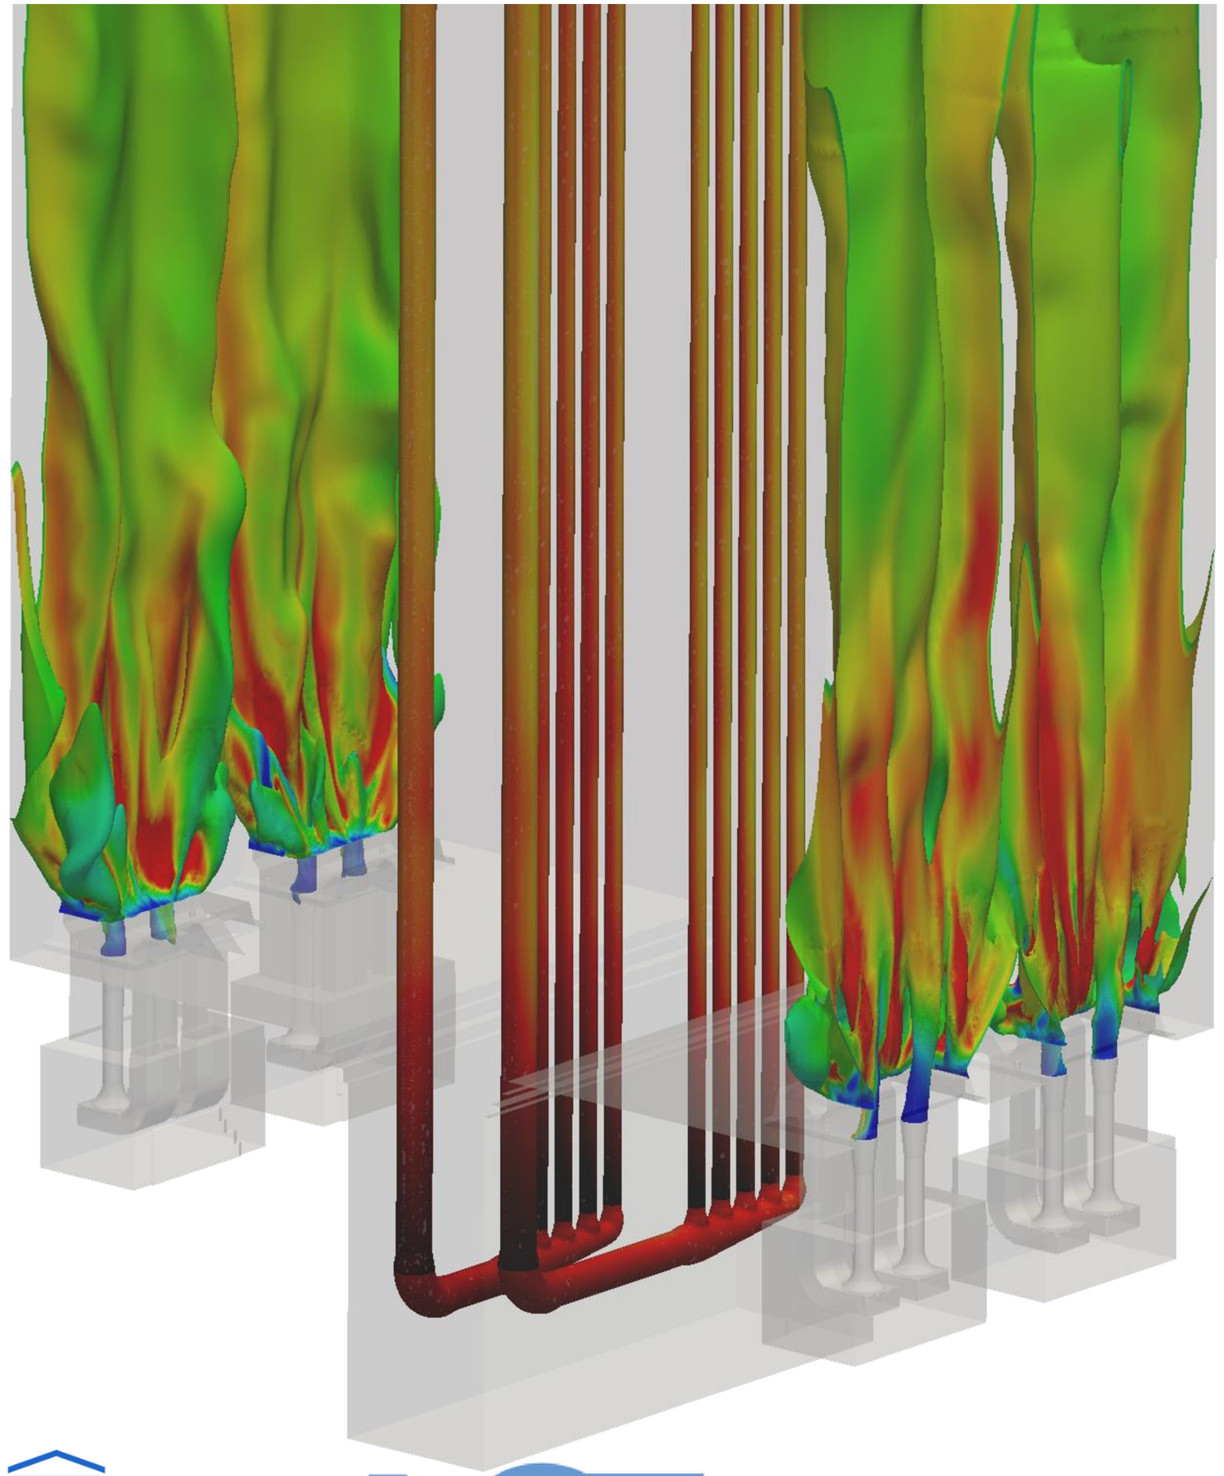

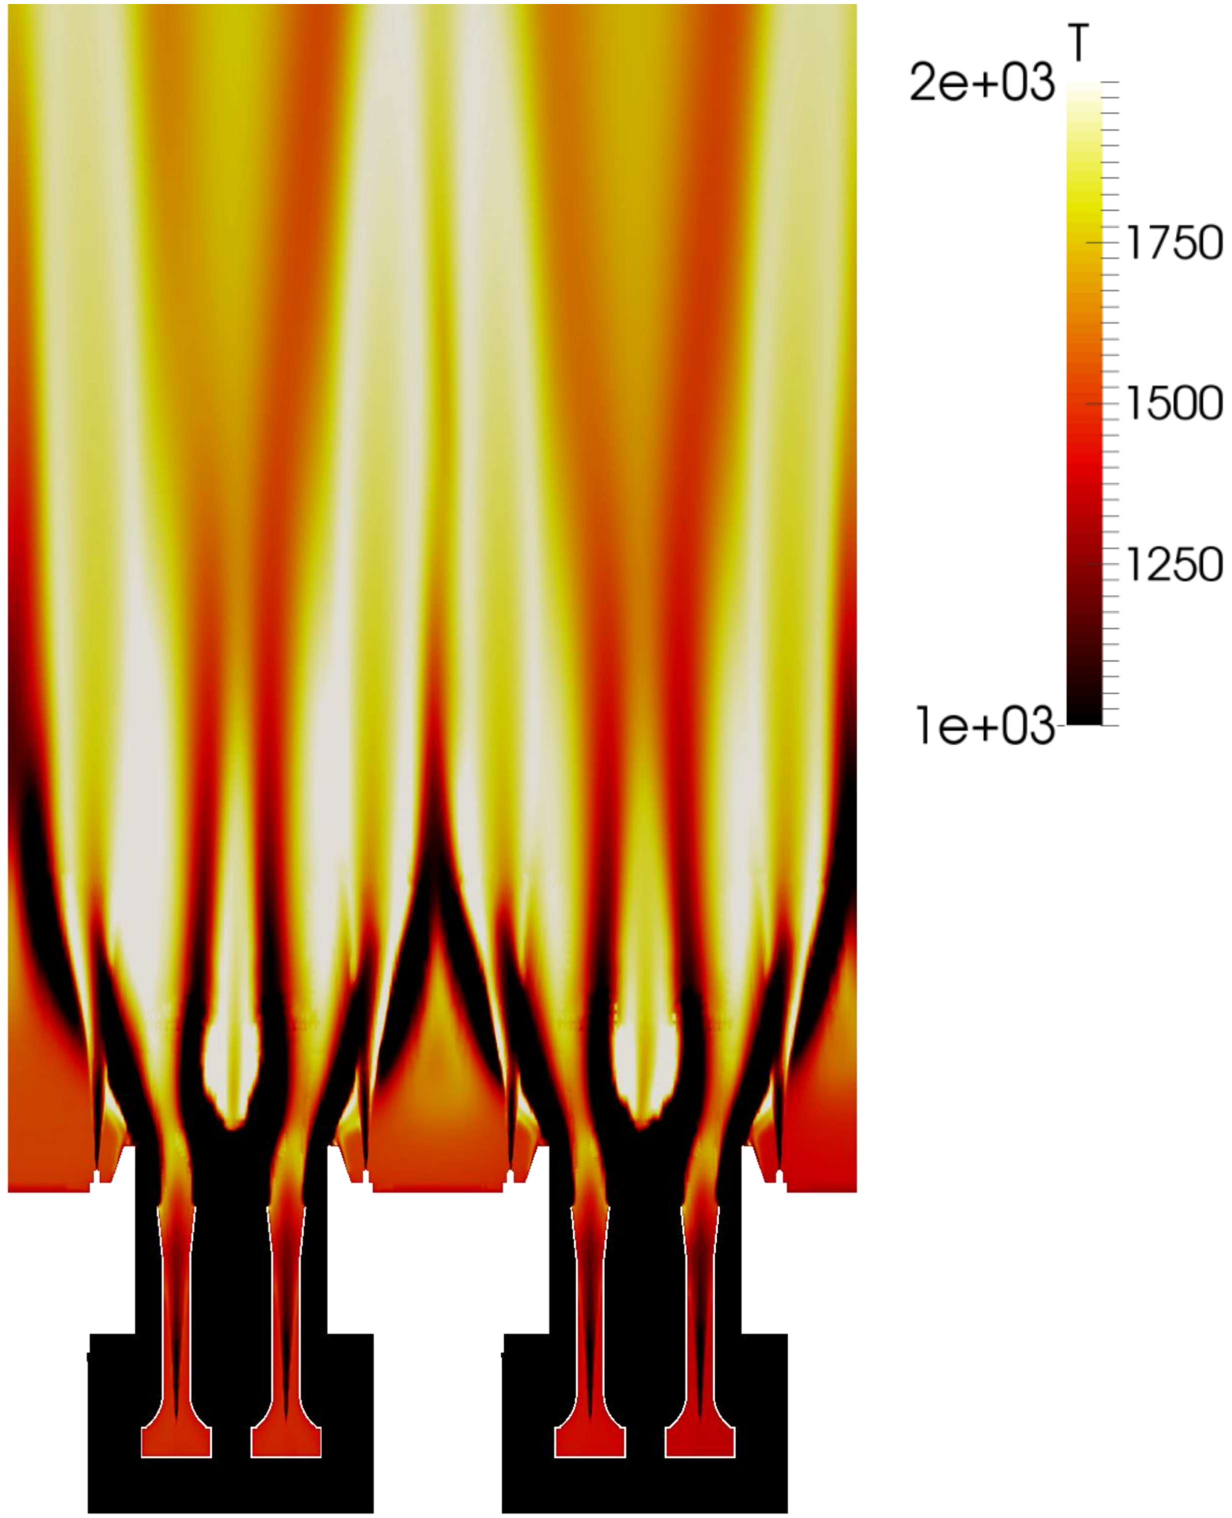

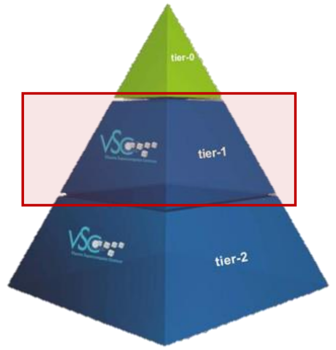

GHENT UNIVERSITY

LCT DRIVING CHEMICAL TECHNOLOGY

DYNAMIC RUN LENGTH SIMULATION

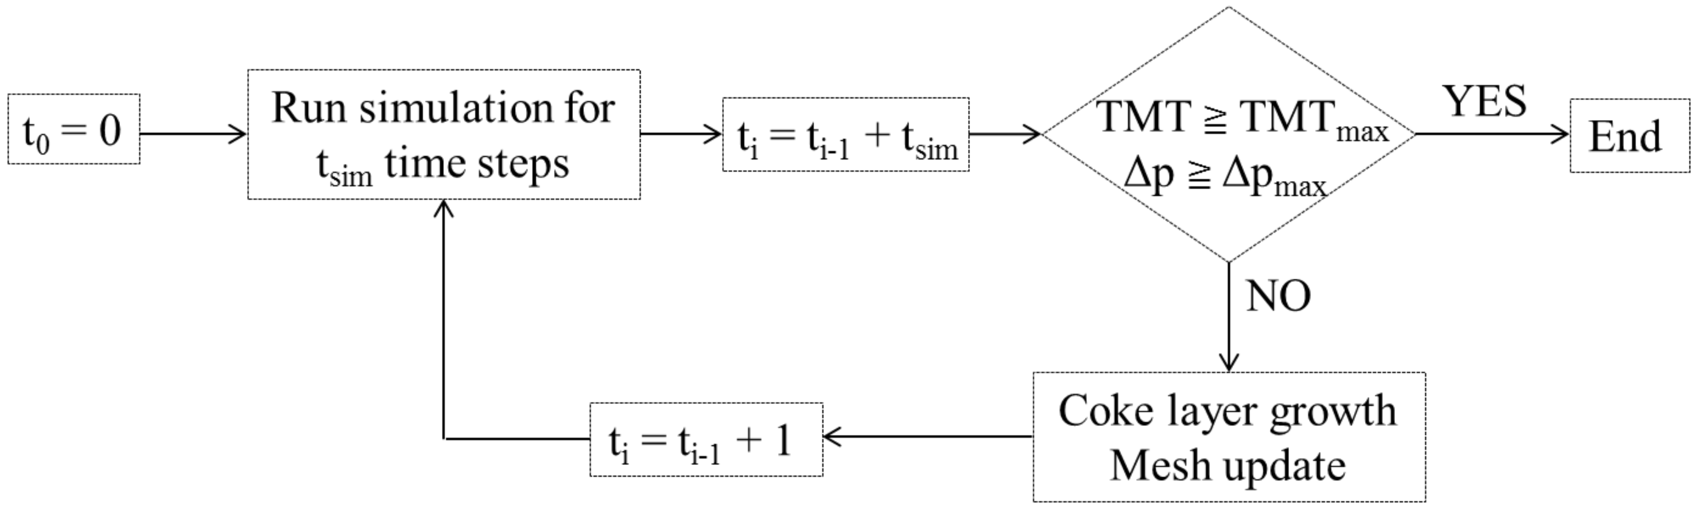

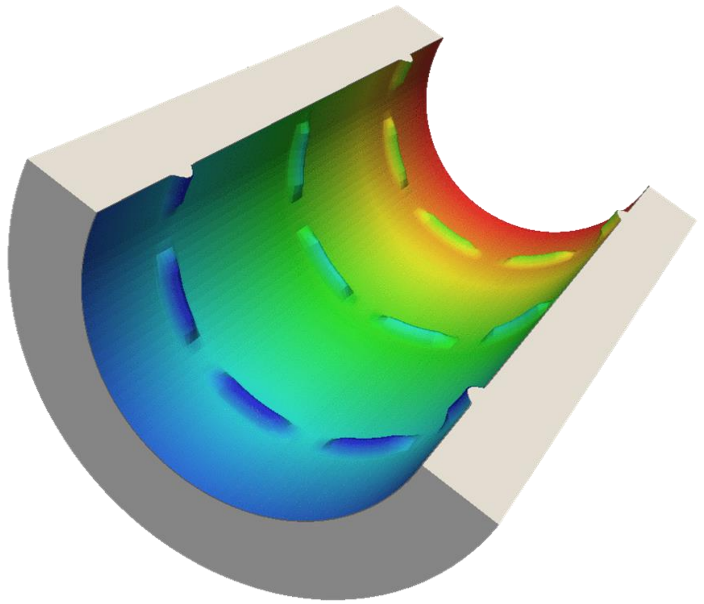

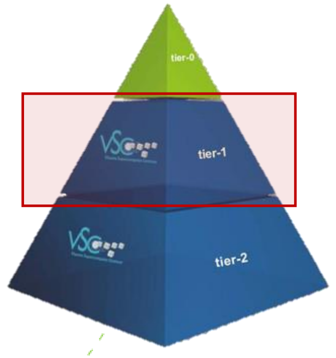

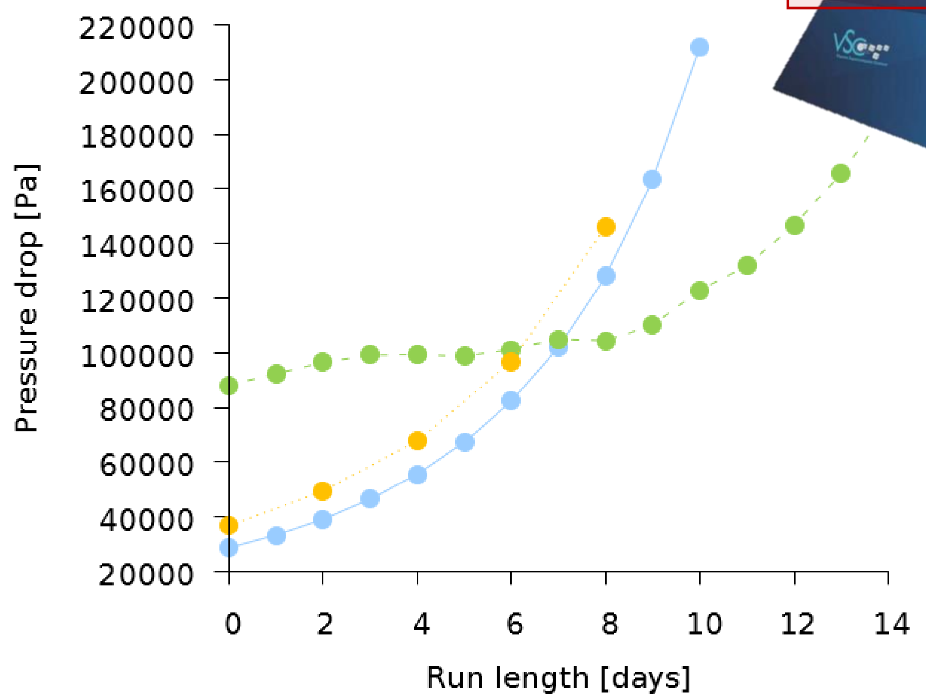

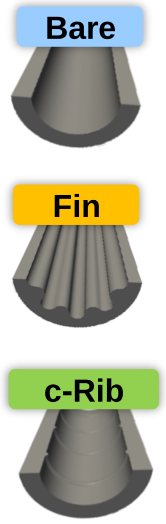

Propane Millisecond
29 species, 151 reactions

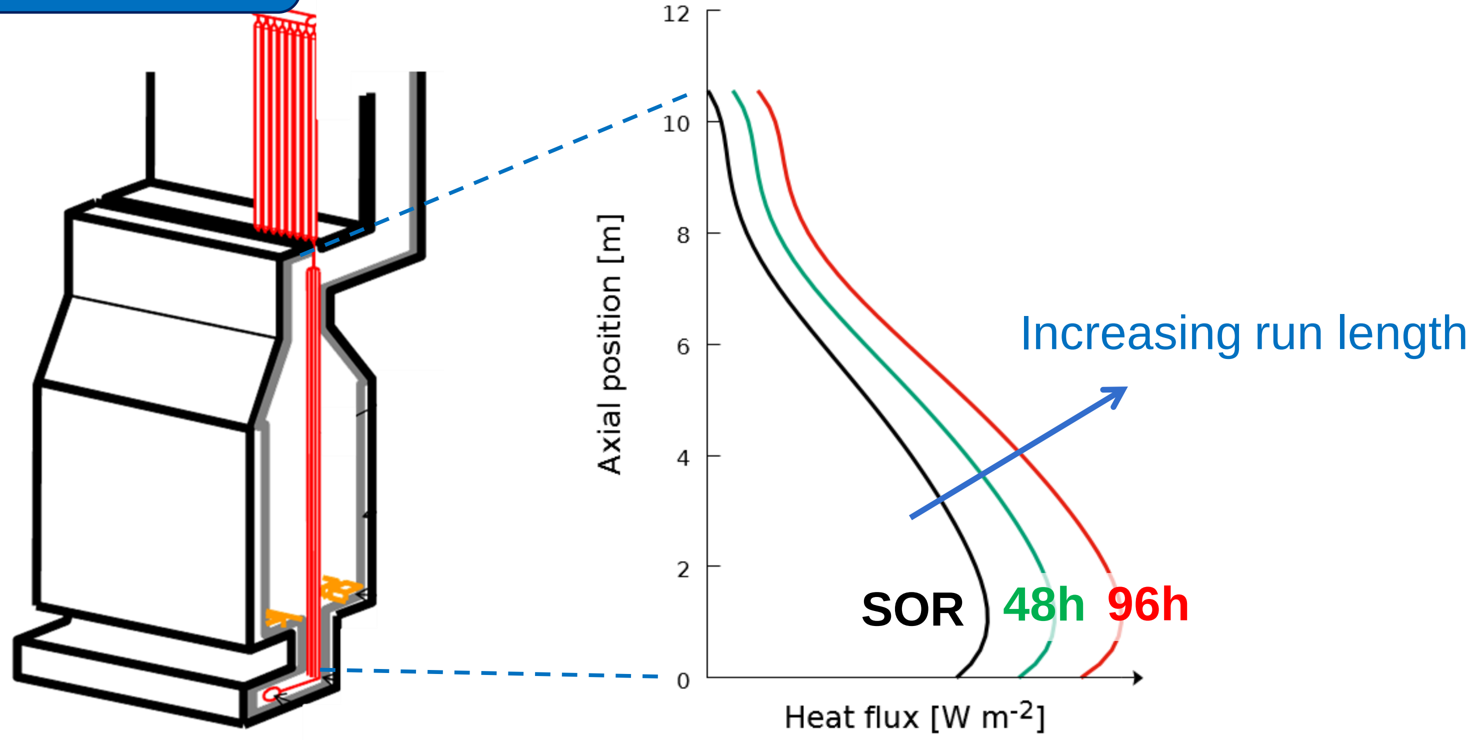

SOR (0 hrs)

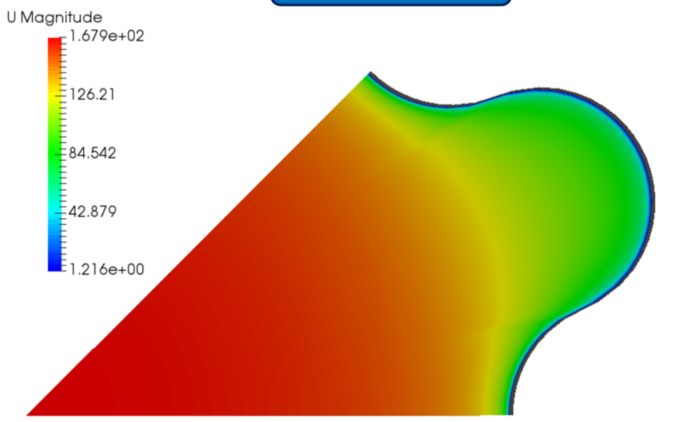

48 hrs

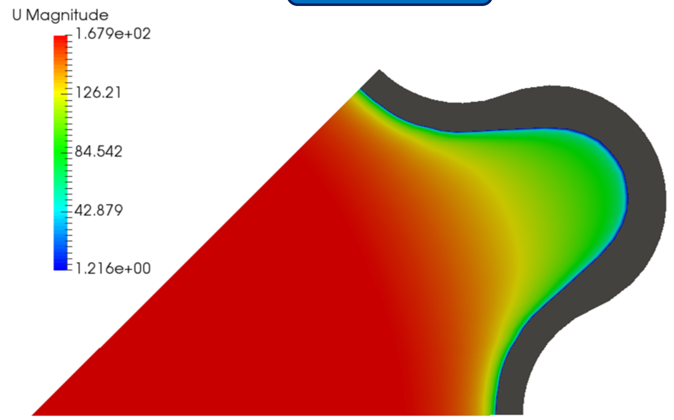

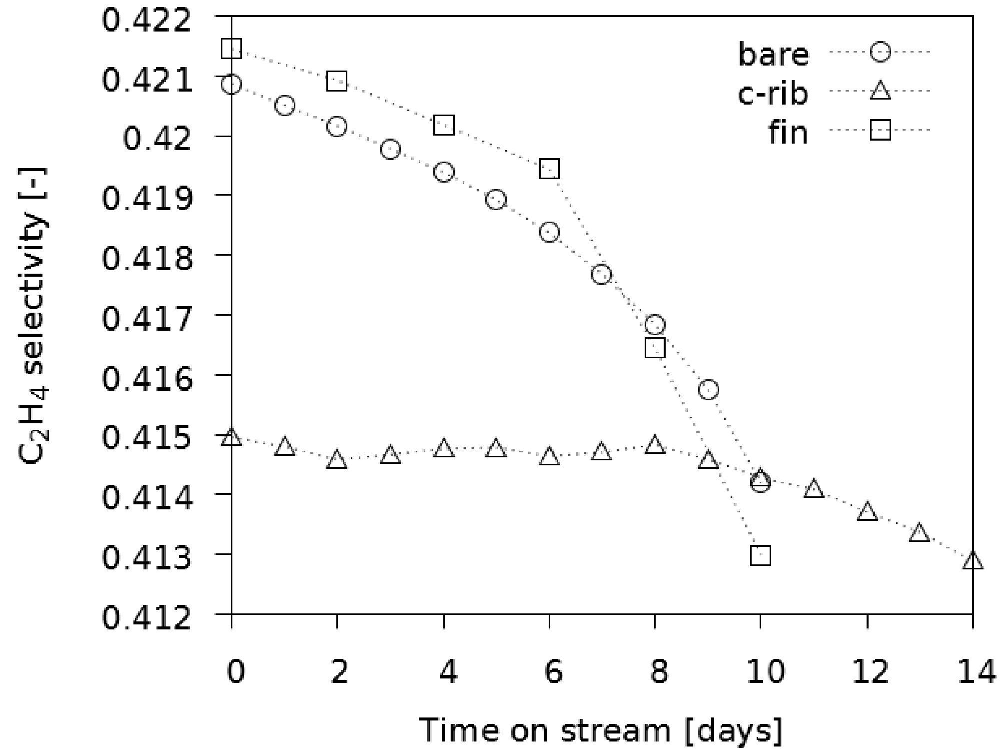

GHENT UNIVERSITY

LCT DRIVING CHEMICAL TECHNOLOGY

Vandewalle, L. A.; Van Cauwenberge, D. J.; Dedeyne, J. N.; Van Geem, K. M.; Marin, G. B., Dynamic simulation of fouling in steam cracking reactors using CFD. Chem. Eng. J. 2017.

OPENFOAM ON TIER-1

Hardware overview

	Tier-1b – BrENIAC	Tier-1a – Muk	Tier-2 – Golett	Tier-2 – Swalot
# nodes	580	528	200	128
CPU	2 x 14c Intel E5-2680v4	2 x 8c Intel E5-2670	2 x 12c Intel E5-2680v3	2 x 10c Intel E5-2660v3
Memory	128 GiB (435) 256 GiB (145)	64 GiB	64 GiB	128 GiB
Interconnect	EDR IB (11.75 GB/s)	FDR IB (6.5 GB/s)	FDR-10 IB (5.0 GB/s)	FDR IB (6.5 GB/s)
Access	Project-based, free	Project-based, paying/free	Open, free	Open, free

OPENFOAM ON TIER-1

Information on project access
https://www.vscentrum.be/en/access-and-infrastructure/project-access-tier1.

- Project proposal in a single document (maximum 17 pages)
- Scientific relevance is demonstrated by framing the calculation time in an approved project
- Next cut-off date for proposals: October 2, 2017.
- Possibility of requesting a starting grant (continuous call)
- FWO bears all the cost but the number of nodedays is limited

- Nearly identical in use compared to UGent Tier-2 machines (modules, scheduler, job-scripts)
- Major difference: accounting system to keep track of consumed nodedays
- Connection between BrENIAC (@KULeuven) and UGent via BelNet (1 Gbps).

OPENFOAM ON TIER-1

Scaling on Tier-1b

muk – small: OpenFOAM/2.2.0-ictce-4.1.13 / Intel MPI v4.1.0
BrENIAC – small: OpenFOAM/2.2.0-intel-2016a / Intel MPI v5.1.3
BrENIAC – big: OpenFOAM/2.2.0-intel-2016a / Intel MPI v5.1.3

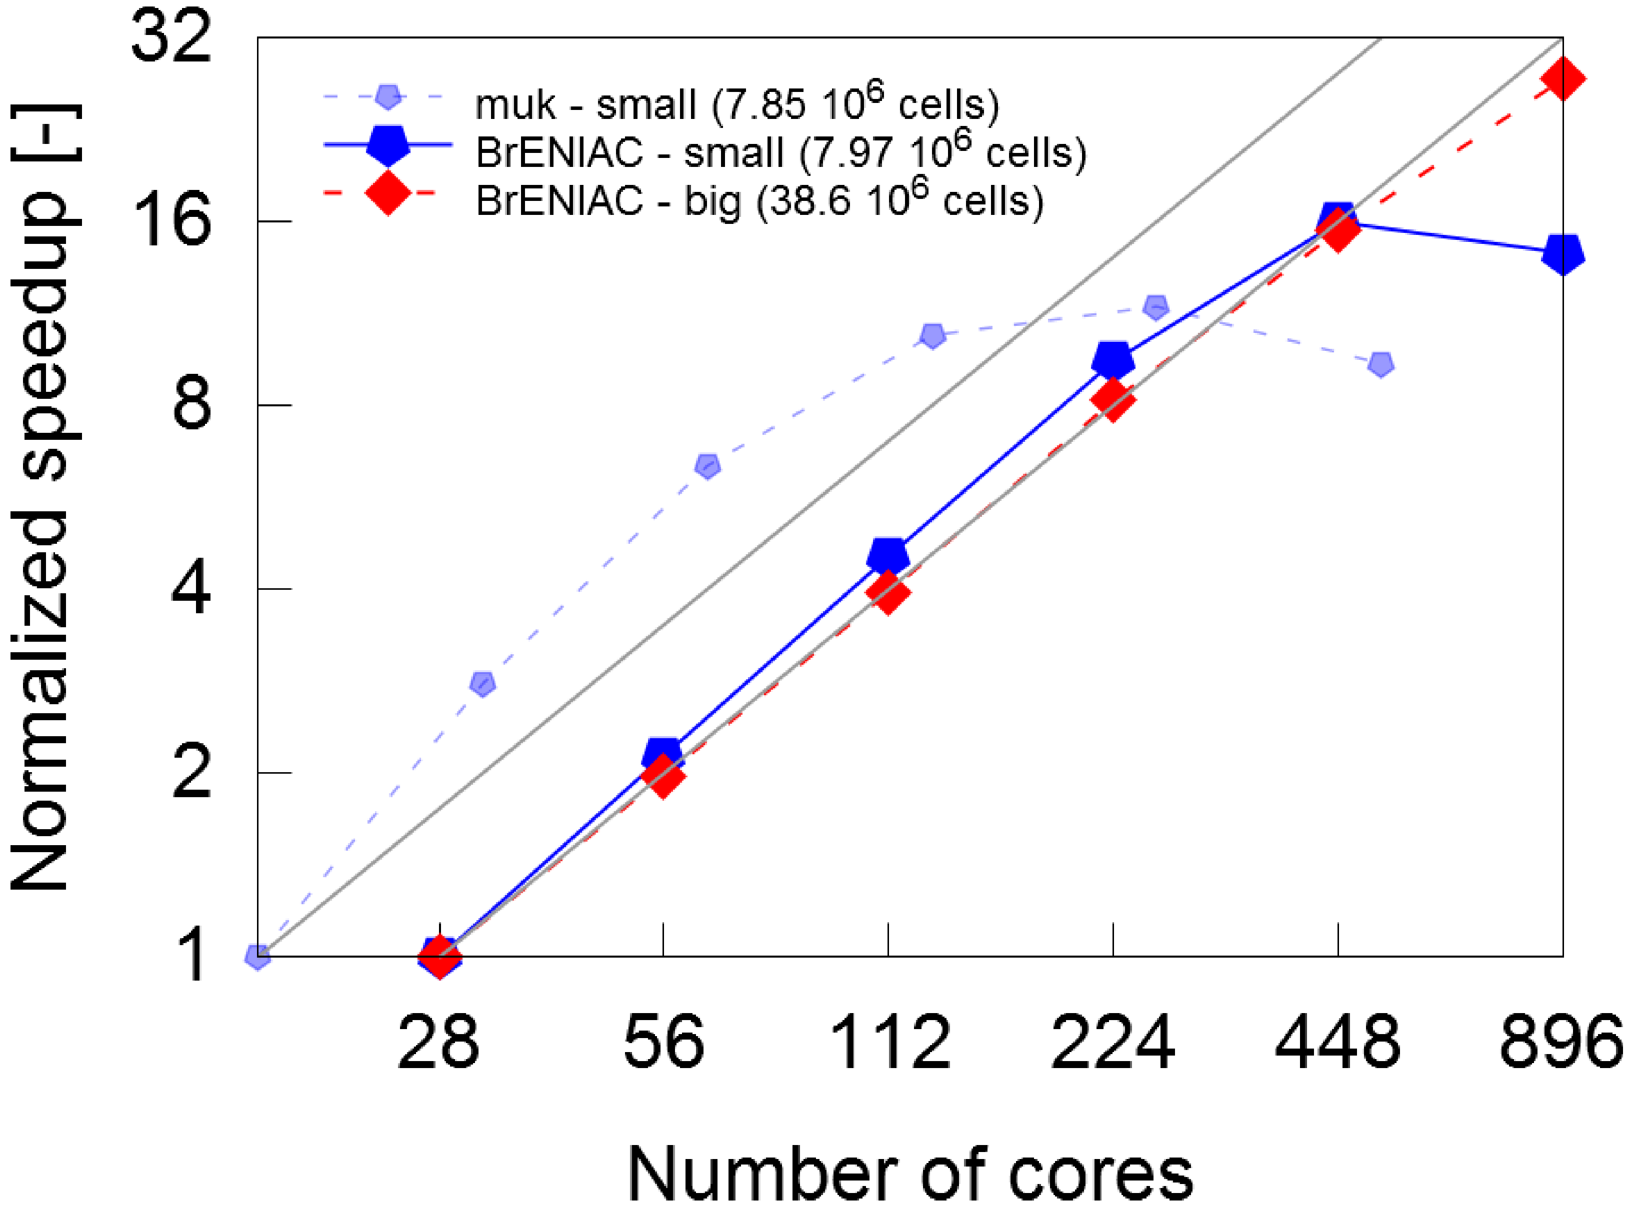

Better scaling compared to Tier-1a – Muk

Fast interconnect (EDR IB) reduces wall-clock time and maintains efficiency while scaling on more cores

Bottleneck: pre- and postprocessing

Remote desktop on Tier-1b login node with GPU via NoMachine client.

GHENT UNIVERSITY

LCT DRIVING CHEMICAL TECHNOLOGY

OPENFOAM ON TIER-1

Scaling on Tier-1b

BrENIAC – small: OpenFOAM/2.2.0-intel-2016a / Intel MPI v5.1.3
BrENIAC – medium: OpenFOAM/2.2.0-intel-2016a / Intel MPI v5.1.3
BrENIAC – big: OpenFOAM/2.2.0-intel-2016a / Intel MPI v5.1.3

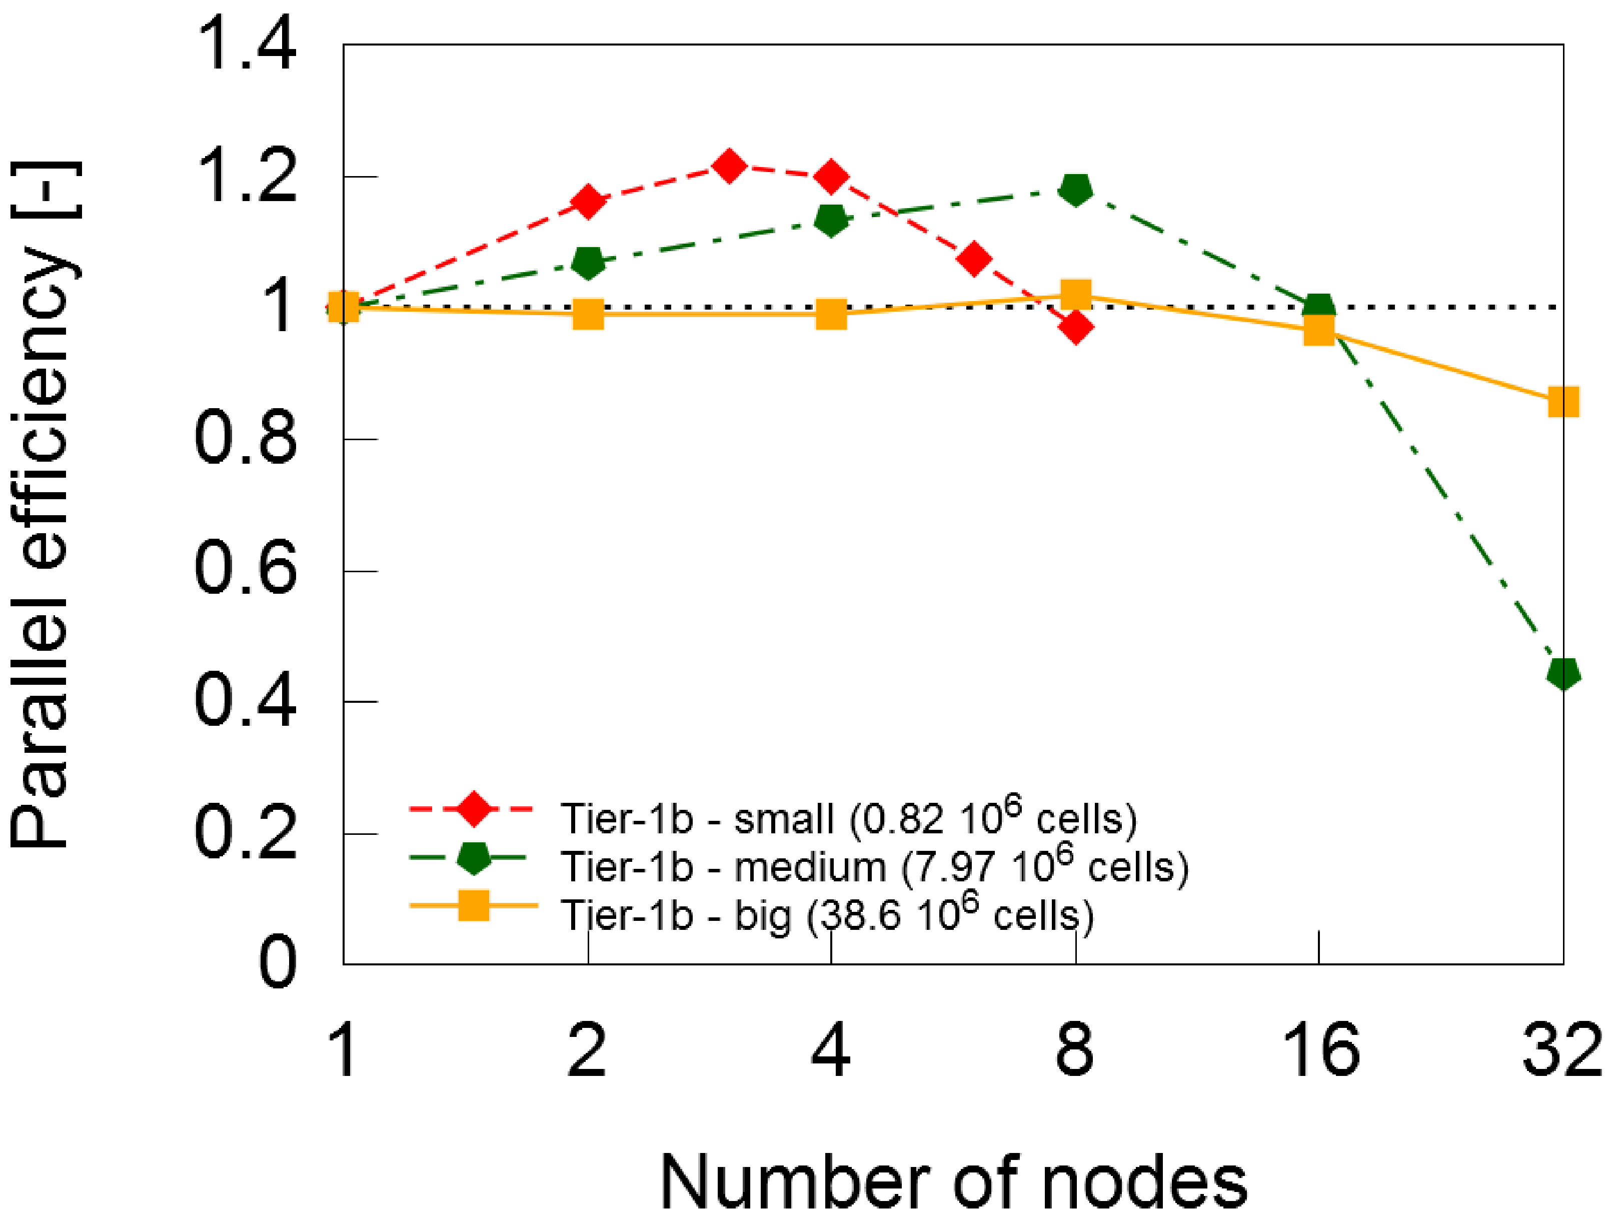

Super-linear scalability due to cache effect and better accomodation of memory patterns across multiple nodes

The choice of decomposition method (scotch, simple, metis, etc.) is important for the number of processor faces

Use `preservePatches` for periodic cases

GHENT UNIVERSITY

LCT DRIVING CHEMICAL TECHNOLOGY

Axtmann, G.; Rist, U., Scalability of OpenFOAM with Large Eddy Simulations and DNS on High-Performance Systems. In *High Performance Computing in Science and Engineering ´16: Transactions of the High Performance Computing Center, Stuttgart (HLRS) 2016*, Nagel, W. E.; Kröner, D. H.; Resch, M. M., Eds. Springer International Publishing: Cham, 2016; pp 413-424.

IMPROVED SCALING FOR OPENFOAM

Meshers generated with native meshers (blockMesh, snappyHexMesh) calculate quicker than third-party meshes

Use *renumberMesh* to decrease bandwidth and increase speed for third-party meshes

Use OpenFOAM versions compiled with recent compiler toolchains

```
------------------------------------ /apps/leuven/broadwell/2016a/modules/all ------------------
OpenFOAM/2.2.0-intel-2016a      OpenFOAM/2.3.1-intel-2016a      OpenFOAM/4.1-intel-2016a
OpenFOAM/2.2.2-intel-2016a      OpenFOAM/3.0.1-intel-2016a      OpenFOAM-Extend/3.2-intel-2016a
   OpenFOAM/2.3.1-intel-2017a      OpenFOAM/3.0.1-intel-2016b      OpenFOAM/4.1-intel-2017a (D)

--------------------------------------- /apps/gent/SL6/sandybridge/modules/all ---------------
   OpenFOAM-Extend/3.2-intel-2016a    OpenFOAM/2.3.1-intel-2015a
   OpenFOAM/2.1.1-ictce-4.0.10        OpenFOAM/2.4.0-intel-2015b
   OpenFOAM/2.2.0-ictce-4.1.13        OpenFOAM/3.0.0-intel-2015b-eb-deps-Python-2.7.10
   OpenFOAM/2.2.2-intel-2015a         OpenFOAM/3.0.1-intel-2016b
   OpenFOAM/2.2.2-intel-2016a         OpenFOAM/4.0-intel-2016b
   OpenFOAM/2.3.0-intel-2014b
```

Tier-1b

IMPROVED SCALING FOR OPENFOAM

Use checkpointing at reasonable time intervals, consider file compression for large cases

`writeControl timeStep;`		
`writeInterval 1000;`	Writing out every 2-3 hours for a 72h job is sufficient
`purgeWrite 2;`	Only keep the last # time steps, earlier time steps are removed
`writeFormat ascii;`	
`writePrecision 6;`	
`writeCompression on;`	For large cases, consider to write out the files in .gz format

OpenFOAM® optimized

IMPROVED SCALING FOR OPENFOAM

Run your job from the appropriate location, excessive I/O on low-bandwidth locations will seriously slow down your job

- **$VSC_DATA**: not meant for calculations, only long-term storage
- **$VSC_SCRATCH**: default scratch on 15k disks
- **$VSC_SCRATCH_NODE**: /tmp location on local node, only accessible as long as the jobs is running, suited for single-node jobs

Disable runTimeModifiable in controlDict to avoid excessive stat() calls at every time step

```
runTimeModifiable false;
```

See best practices document on the HPC UGent support site: http://www.ugent.be/hpc/en/support

IMPROVED SCALING FOR OPENFOAM

NEW in OpenFOAM v5.0: `collated` file format

> "the data for each decomposed field (and mesh) is collated into a **single** file that is written (and read) on the master processor. The files are stored in a single directory named `processors`."
>
> "The file writing can be threaded allowing the simulation to continue running while the data is being written to file."

More information: https://openfoam.org/news/parallel-io/

IMPROVED SCALING FOR OPENFOAM

CASE STUDY (OpenFOAM 1.7)

Metadata handling becomes a bottleneck when scaling on a large number of cores due to an increasing volume of small files

Number of processes	64	128	256	512	1024
Compute time [s]	686	801	890	1161	2248
Cumulative metadata [s]	64	202	274	389	892
Metadata share [%]	9%	25%	31%	34%	39%

Number of processes	64	128	256	512	1024
# files created	512	1024	2048	4096	8192
# files read	1089	2177	4353	9729	17409
Average file size	597K	317K	163K	84K	47K
# stat() calls	500,000	1,000,000	2,000,000	4,400,000	8,500,000

Curie

owned by GENCI

operated into the TGCC by CEA

B510 bullx nodes

2 x 8c Intel E5-2680

64 GiB

local SSD disk

QDR IB Full Fat Tree

LUSTRE storage (150 GB/s)

GHENT UNIVERSITY

LCT DRIVING CHEMICAL TECHNOLOGY

Lindi, B., I/O-profiling with Darshan. Norwegian University of Science and Technology (NTNU): http://www.prace-ri.eu/IMG/pdf/IO-profiling_with_Darshan-2.pdf
Moylesa, M., Nash, P., Girotto, I., Performance Analysis of Fluid-Structure Interactions using OpenFOAM, Irish Centre for High End Computing: http://www.prace-ri.eu/IMG/pdf/Performance_Analysis_of_Fluid-Structure_Interactions_using_OpenFOAM.pdf

IMPROVED SCALING FOR OPENFOAM

CASE STUDY (OpenFOAM 1.7)

stat()-calls are use to check the timestamps of files to check for updates. Disabling 'runTimeModifiable' reduces the number of stat()-calls drastically

Number of processes	64	128	256	512
Compute time [s]	542	381	343	411
Cumulative metadata [s]	0.99	2.62	6.03	14.4
Metadata share [%]	0.20%	0.70%	1.80%	3.50%

Number of processes	64	128	256	512
# files created	512	1024	2048	4096
# files read	1089	2177	4353	9729
Average file size	597K	317K	163K	84K
# stat() calls	5,000	10,000	20,000	44,000

Curie

owned by GENCI

operated into the TGCC by CEA

B510 bullx nodes

2 x 8c Intel E5-2680

64 GiB

local SSD disk

QDR IB Full Fat Tree

LUSTRE storage (150 GB/s)

GHENT UNIVERSITY

LCT DRIVING CHEMICAL TECHNOLOGY

Lindi, B., I/O-profiling with Darshan. Norwegian University of Science and Technology (NTNU): http://www.prace-ri.eu/IMG/pdf/IO-profiling_with_Darshan-2.pdf
Moylesa, M., Nash, P., Girotto, I., Performance Analysis of Fluid-Structure Interactions using OpenFOAM, Irish Centre for High End Computing: http://www.prace-ri.eu/IMG/pdf/Performance_Analysis_of_Fluid-Structure_Interactions_using_OpenFOAM.pdf

PREDICTION OF TURBULENT REACTIVE FLOWS

PRETREF: http://www.pretref.ugent.be/

Contact: dr. Georgios Maragkos (Georgios.Maragkos@UGent.be)

A project by **Ghent University** which aims to develop a flexible, open source **Large-Eddy Simulations (LES) Computational Fluid Dynamics (CFD)** code-base for multiscale modelling of several multidisciplinary applications.

Objectives defined in the following fields

1. Reduced chemistry
2. Sprays
3. Turbulent steady spray flames
4. Unsteady sprays, in internal combustion engines
5. Fire dynamics

ACKNOWLEDGEMENTS

Colleagues @LCT

David Van Cauwenberge

Jens Dedeyne

Pieter Reyniers

Moreno Geerts

Laurien Vandewalle

Shekhar Kulkarni

Pieter Verhees

Yu Zhang

LABORATORY FOR CHEMICAL TECHNOLOGY

Tech Lane Ghent Science Park – Campus A
Technologiepark 914, 9052 Ghent, Belgium

E info.lct@ugent.be
T 0032 9 331 17 57

https://www.lct.ugent.be

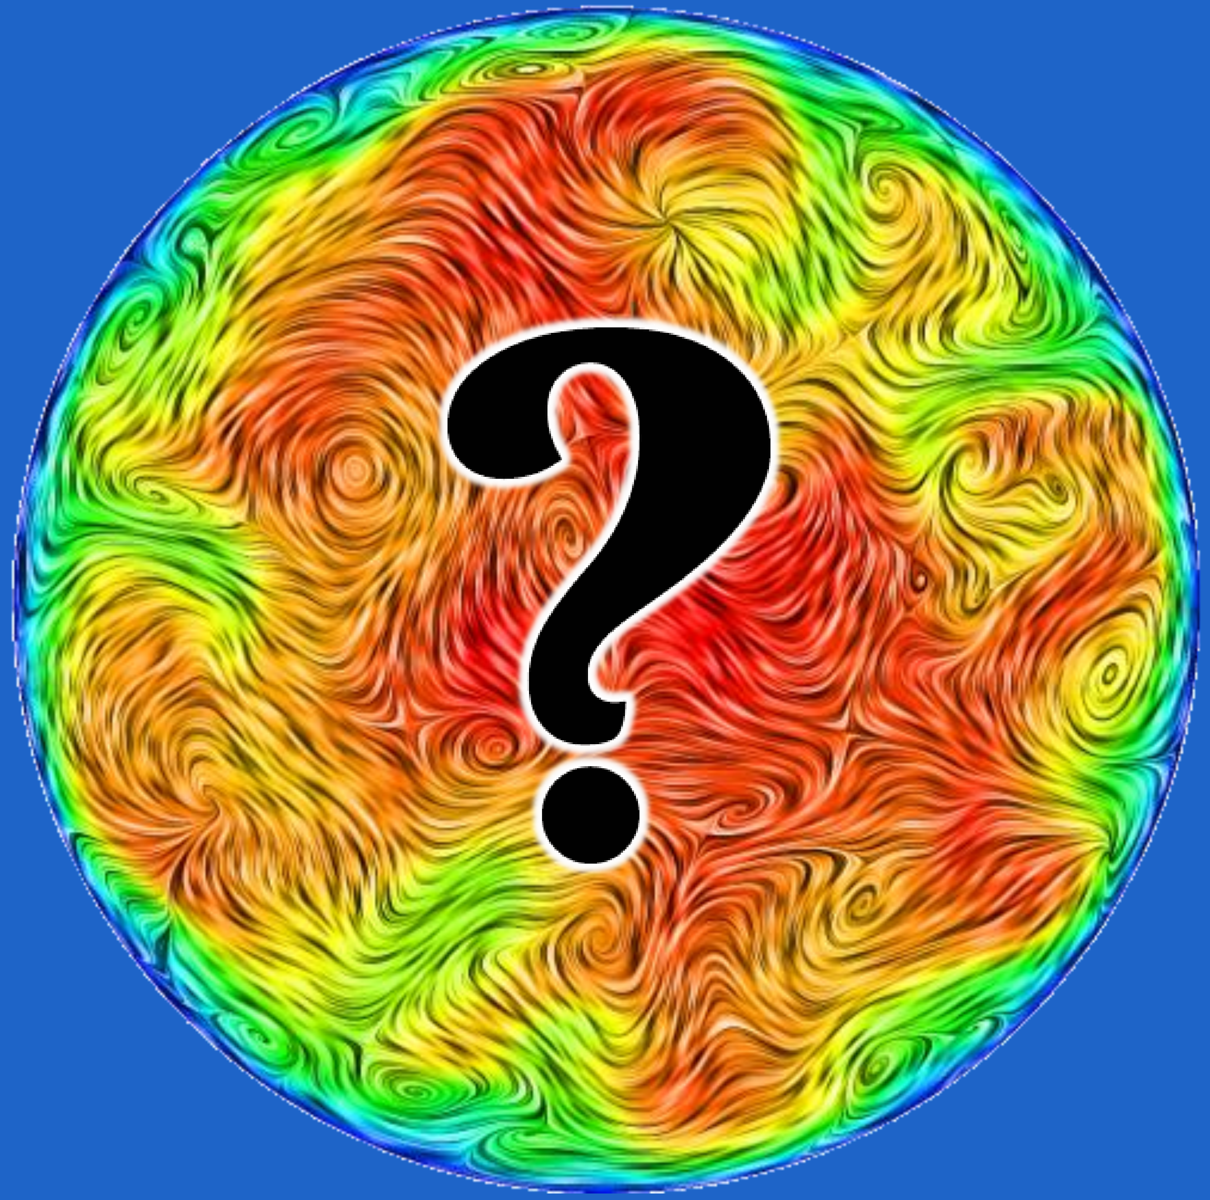

GHENT UNIVERSITY

LCT
DRIVING CHEMICAL TECHNOLOGY