Enabling distributed kinetic Monte Carlo simulations for catalysis and materials science

Michail Stamatakis, University College London

Main contributors: Ilektra Christidi, Roland Guichard, Giannis Savva, Srikanth Ravipati, Miguel Pineda

Edinburgh Parallel Computing Centre • ARCHER Webinar
Wednesday 25th September 2019

Catalytic Materials Design

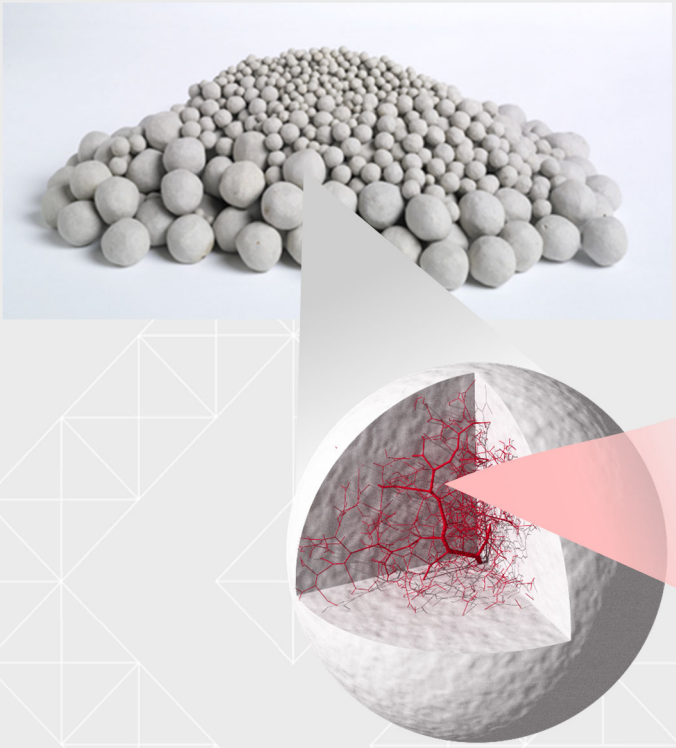

Convection/diffusion of reactants towards (or products from) the surface of the pellet

Diffusion through the porous network of the catalytic pellet

Possible **bulk phase reactions**…

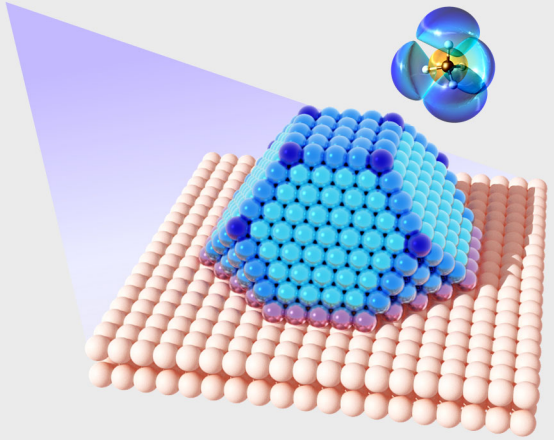

Adsorption and surface reactions on the catalytically active phase

Materials modelling opportunities:

- Property prediction
- Materials discovery/engineering
- Unit (reactor) & process design

The Kinetic Monte Carlo Approach

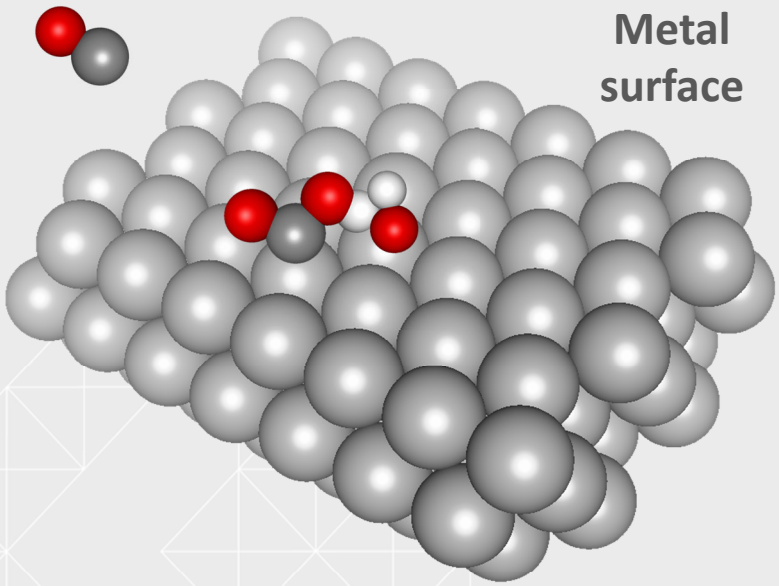

$CO_{(gas)} + OH^* \rightleftarrows COOH^*$

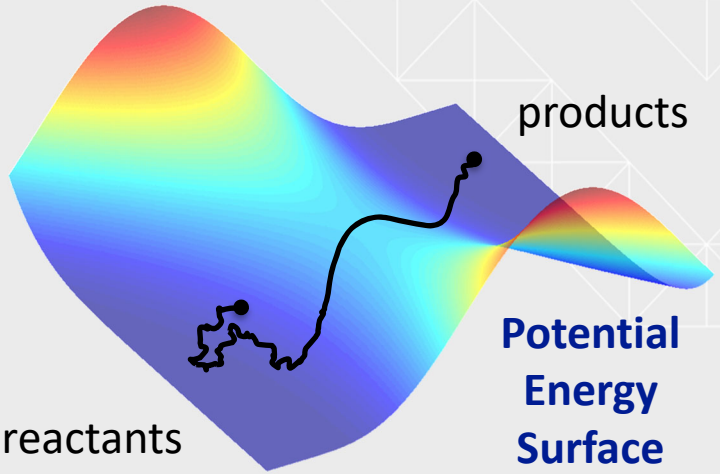

- Instead of simulating dynamics, KMC[1] focuses on rare events
- Simulates reactions much faster than Molecular Dynamics
- Incorporates spatial information contrary to micro-kinetic models[2]

[1] M. Neurock and E. W. Hansen, Comput. Chem. Eng. 22, S1045 (1998); K. Reuter and M. Scheffler, Phys. Rev. Lett. 90: 046103 (2003); M. Stamatakis, J. Phys. Condens. Matter. 27: 013001 (2015).
[2] J. A. Dumesic et al., The Microkinetics of Heterogeneous Catalysis. (American Chemical Society, 1993).

The Kinetic Monte Carlo Approach

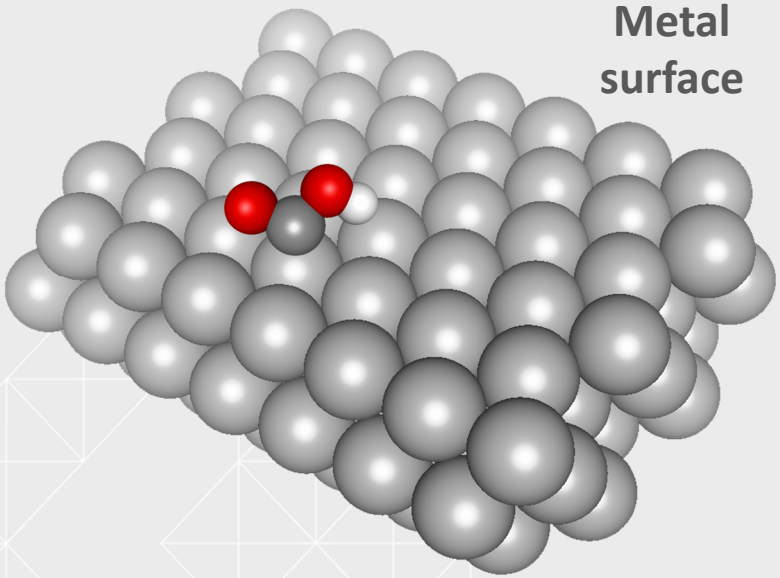

$CO_{(gas)} + OH^* \rightleftarrows COOH^*$

$$k_{rxn} = \kappa \cdot \frac{k_B \cdot T}{h} \cdot \frac{Q^{\ddagger}}{Q_{rxn}} \cdot e^{-\frac{\Delta E^{\ddagger}}{k_B \cdot T}}$$

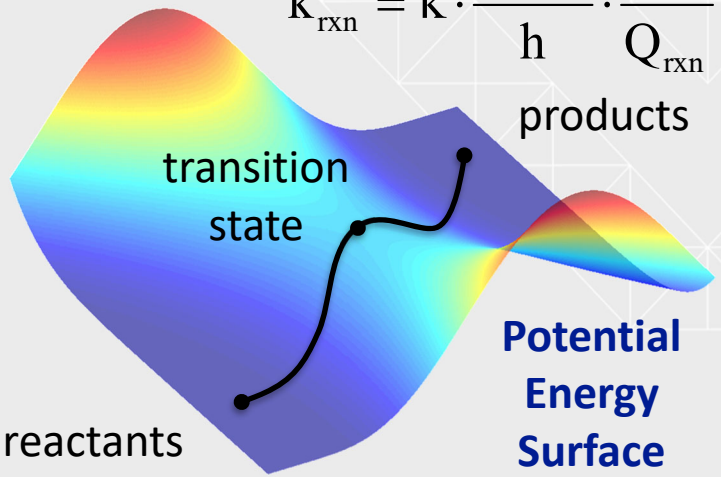

- Instead of simulating dynamics, KMC[1] focuses on rare events
- Simulates reactions much faster than Molecular Dynamics
- Incorporates spatial information contrary to micro-kinetic models[2]

[1] M. Neurock and E. W. Hansen, Comput. Chem. Eng. 22, S1045 (1998); K. Reuter and M. Scheffler, Phys. Rev. Lett. 90: 046103 (2003); M. Stamatakis, J. Phys. Condens. Matter, 27: 013001 (2015).

[2] J. A. Dumesic et al., The Microkinetics of Heterogeneous Catalysis. (American Chemical Society, 1993).

KMC Algorithm Flowchart

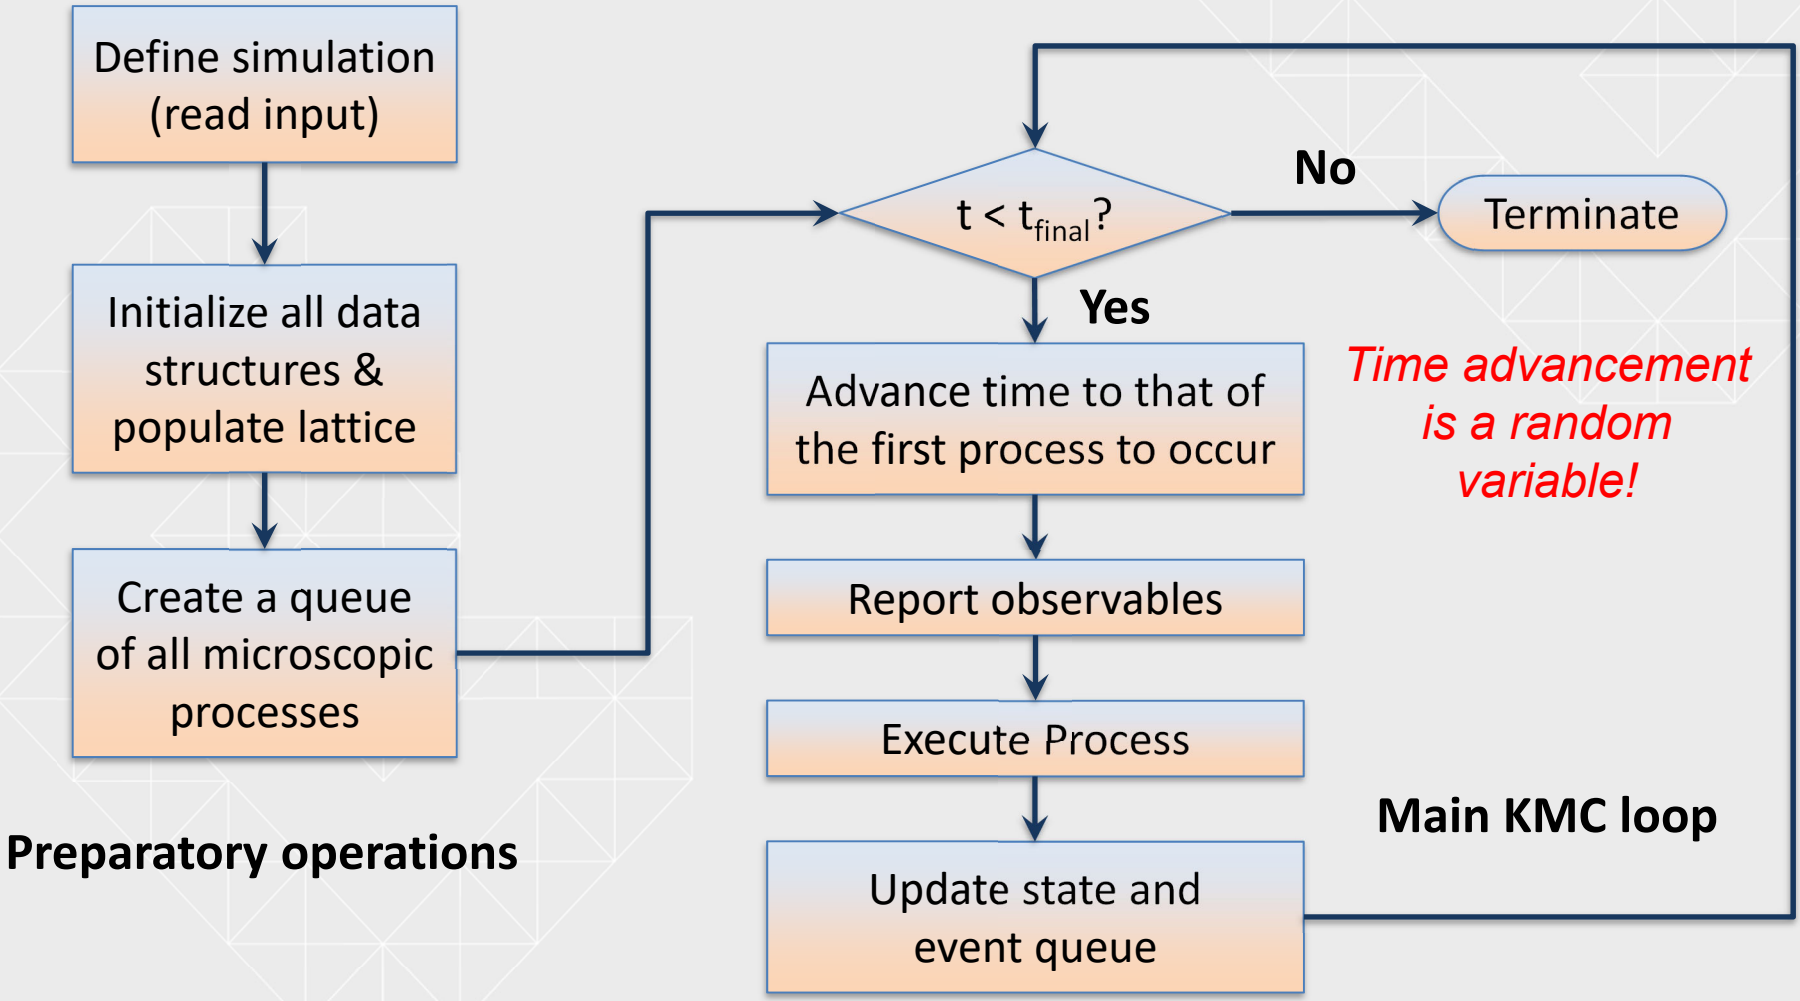

Typical KMC Output

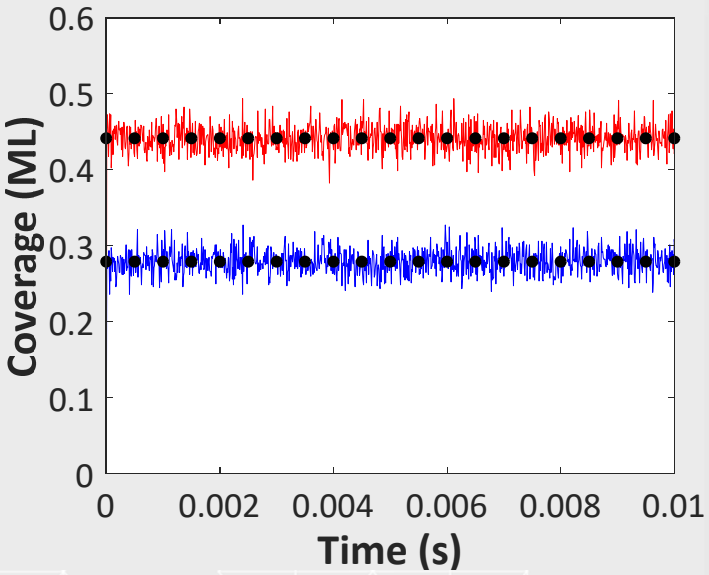

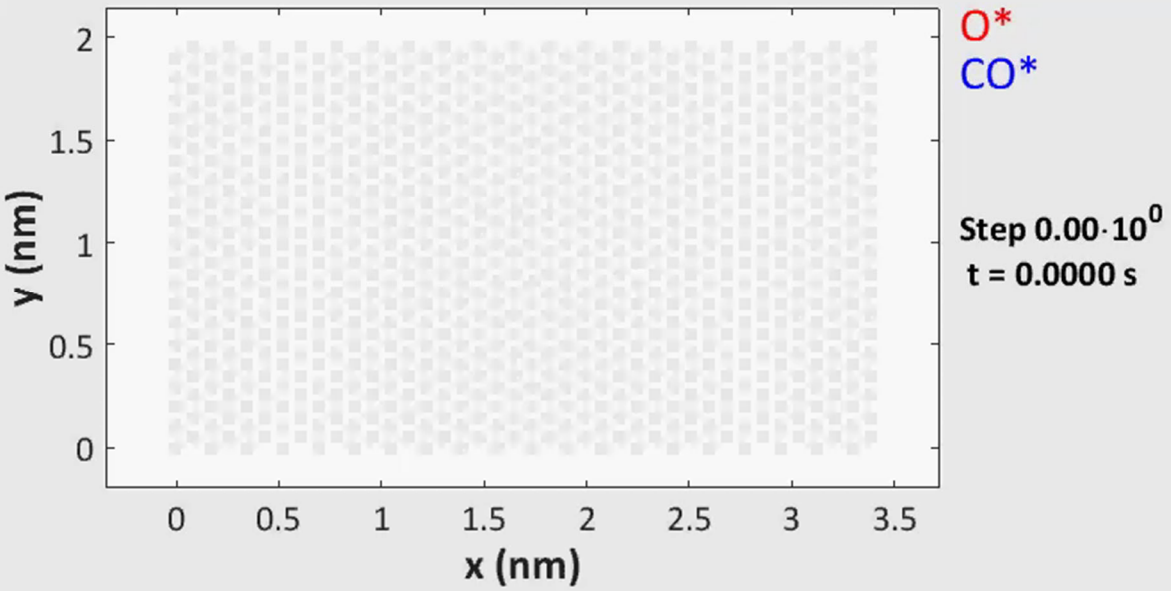

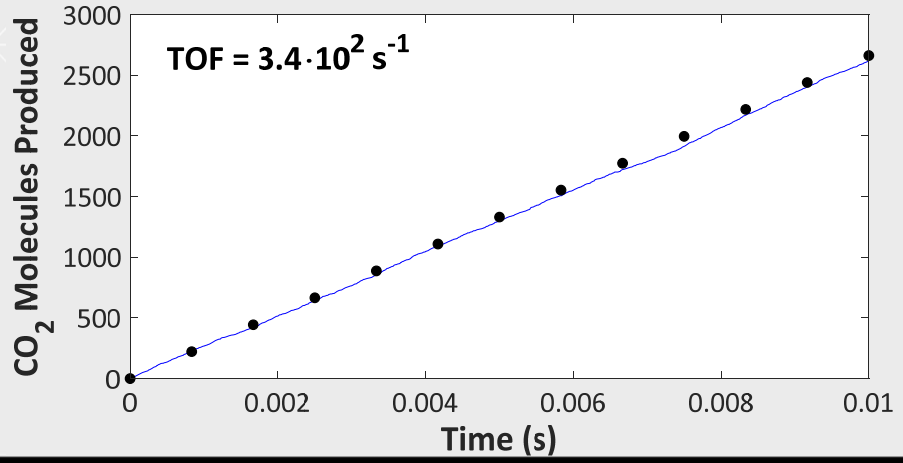

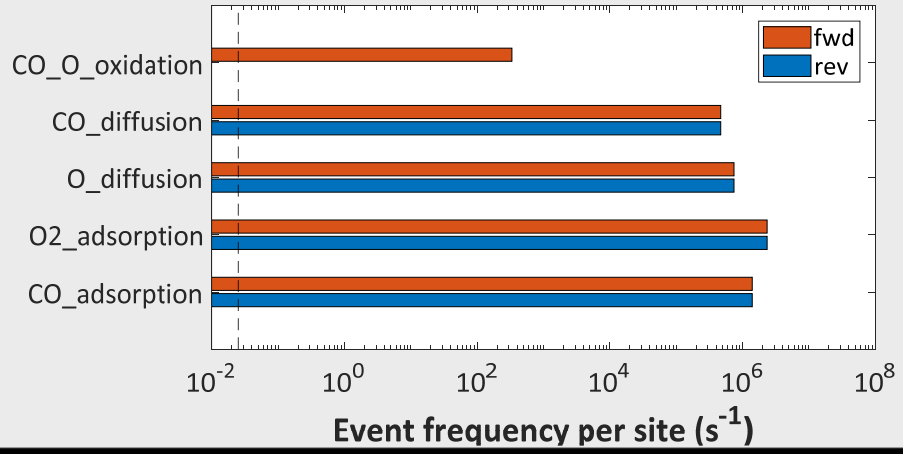

Our Approach to Kinetic Simulation

Zacros
Advanced lattice-KMC made easy

- For computational catalysis and surface science
- Captures detailed energetics and complex reaction mechanisms
- Features keyword-based syntax and troubleshooting

http://zacros.org

M. Stamatakis and D. G. Vlachos. J. Chem. Phys. 134, 214115 (2011).
J. Nielsen, M. d'Avezac, J. Hetherington and M. Stamatakis. J. Chem. Phys. 139, 224706 (2013).
https://xip.uclb.com/i/software/material_modelling/Zacros.html

Why Distributed Simulations?

Obtaining higher accuracy:

- Errors of observables in Monte Carlo drop as $1/\sqrt{N_{samples}}$
- Larger lattices ⇒ more reactions per unit time ⇒ more samples…

  Could run multiplicates… **but not always!**

Capturing relevant physics:

- Phenomena evolving at large spatial scales:
  - ✓ Catalyst reconstruction
  - ✓ Pattern formation

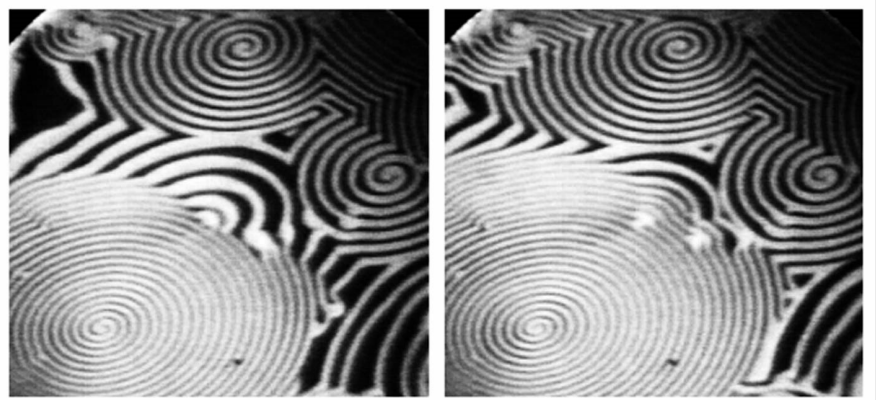

Spirals with wavelengths of 100s of atomic diameters

S. Nettesheim, A. Vonoertzen, H. H. Rotermund and G. Ertl. J. Chem. Phys., 98, 9977-9985 (1993).

Efficient Distributed KMC: Non-Trivial!

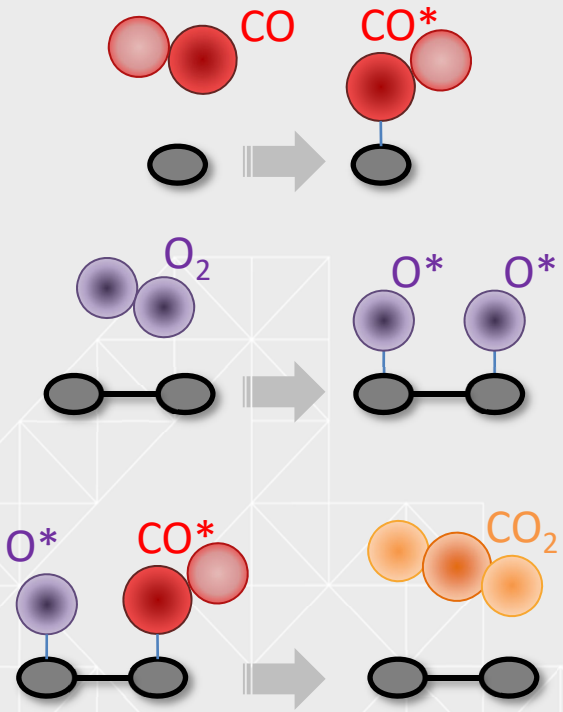

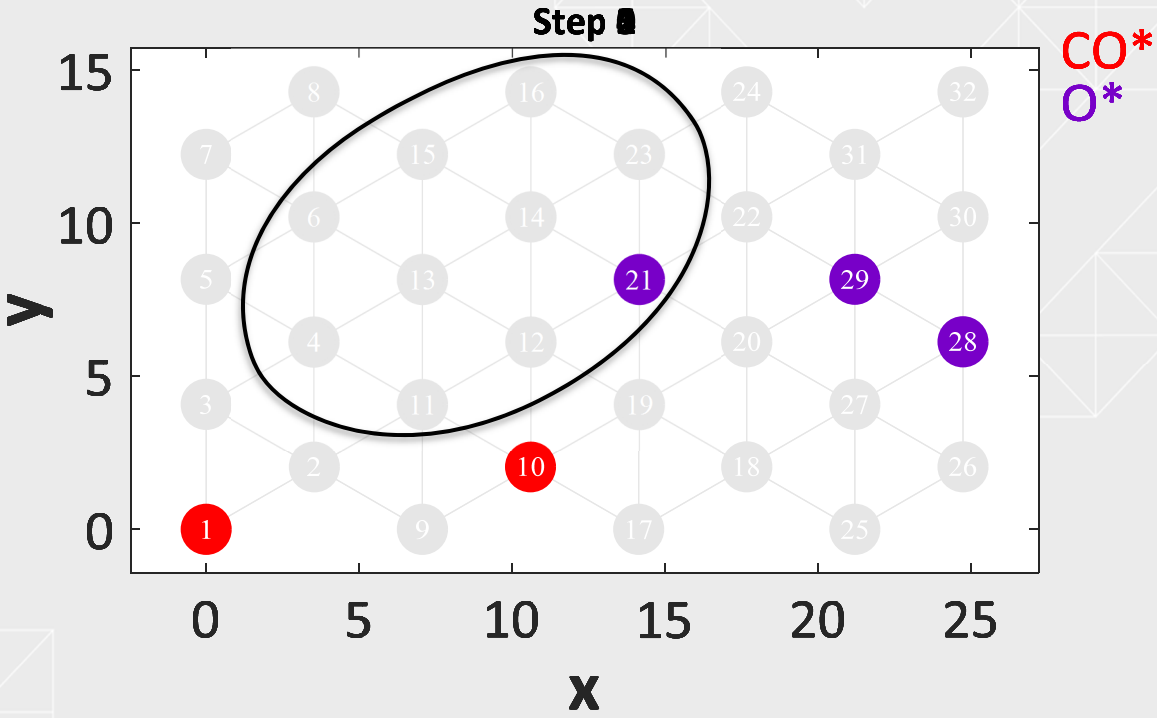

- **The KMC algorithm is inherently serial**, simulating history of events that have causal relationships
- **Optimised KMC algorithms need only perform local updates** due to the finite range of interactions or reactions…

How about Domain Decomposition?

- Define subdomains, each assigned to an MPI process
- Events at **internal** sites are simulated “privately” and asynchronously
- Events occurring at “boundaries” (**halos**) have to be communicated via **messages** to neighbouring subdomain(s)

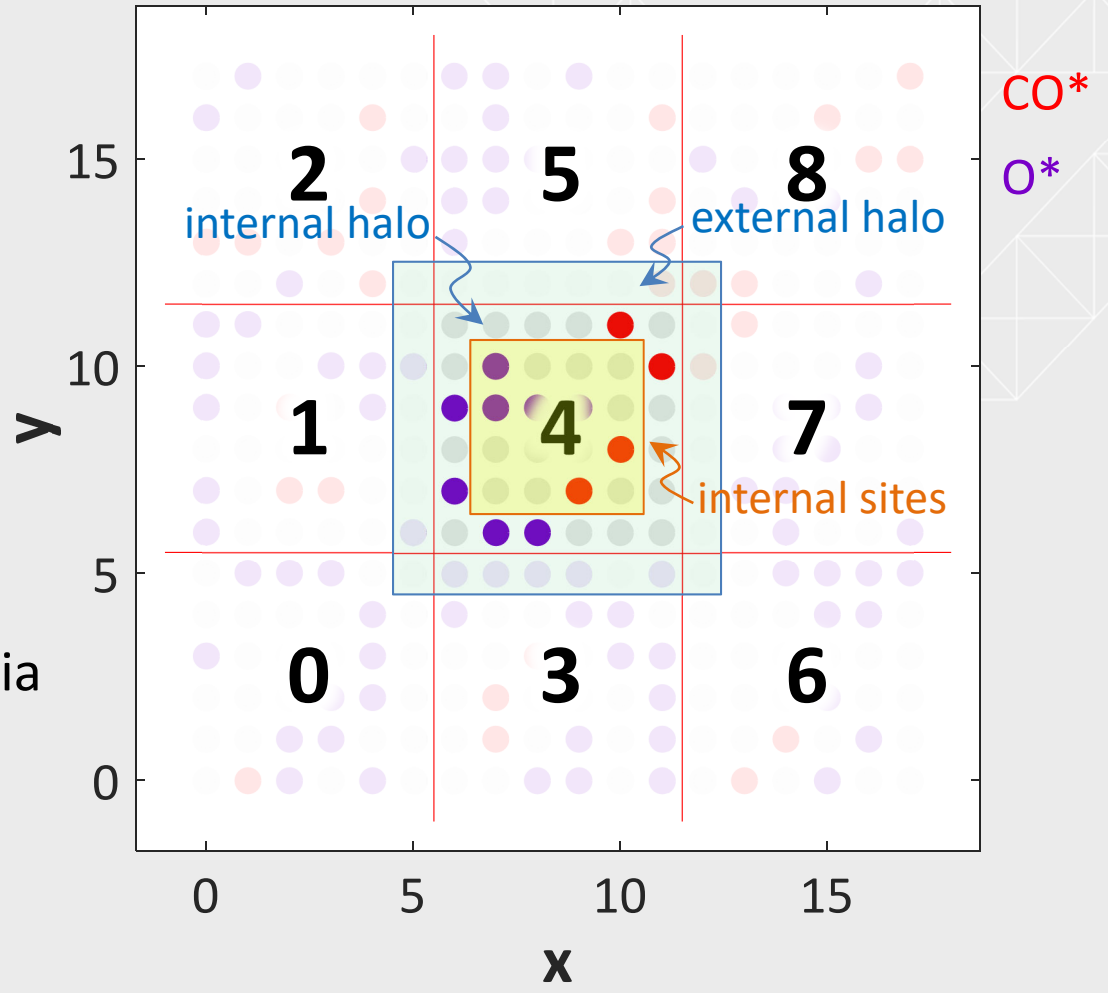

Maintaining Causality

- The asynchronous nature of the simulation can lead to **violations of causality**!

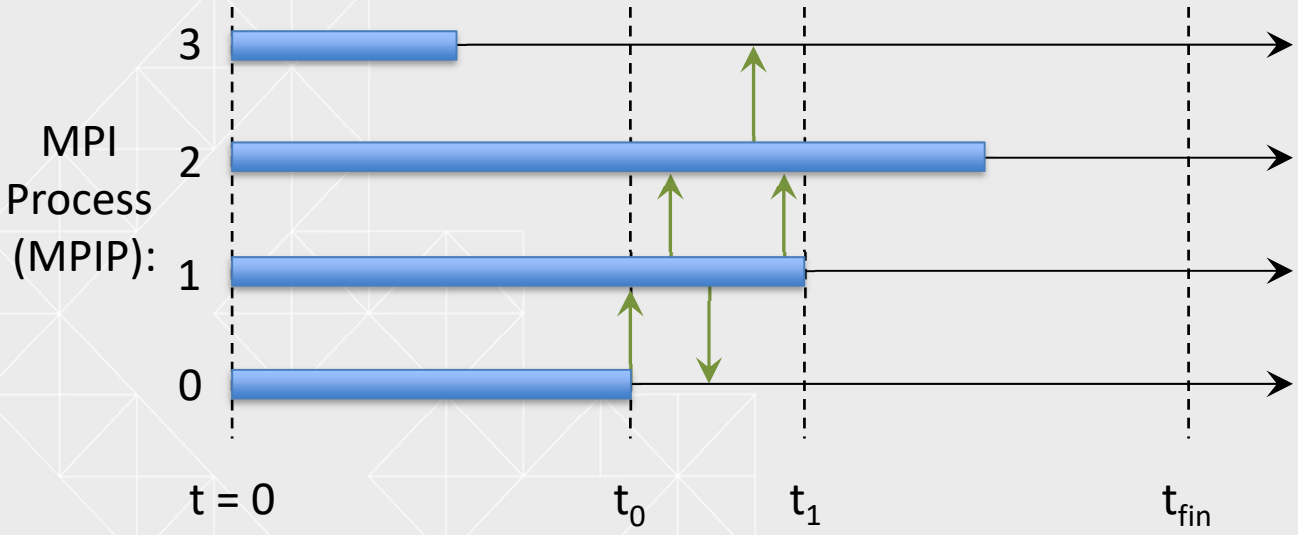

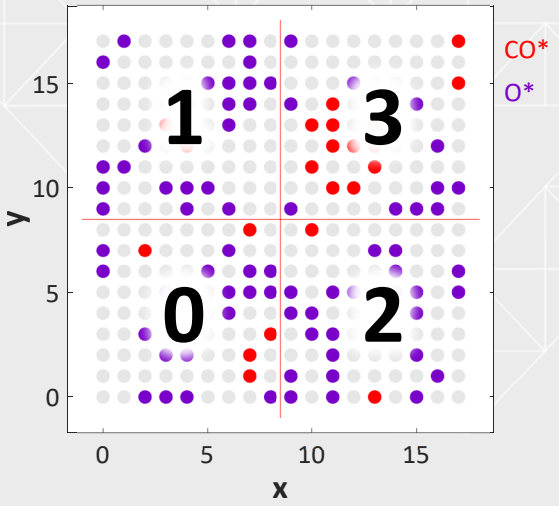

↑: messages

- At KMC time t_1 MPIP 0 sends a message to MPIP 1
- **Too late!** MPIP 1 is already at t_1… Need **rollback** mechanism!

Maintaining Causality

- The asynchronous nature of the simulation can lead to **violations of causality**!

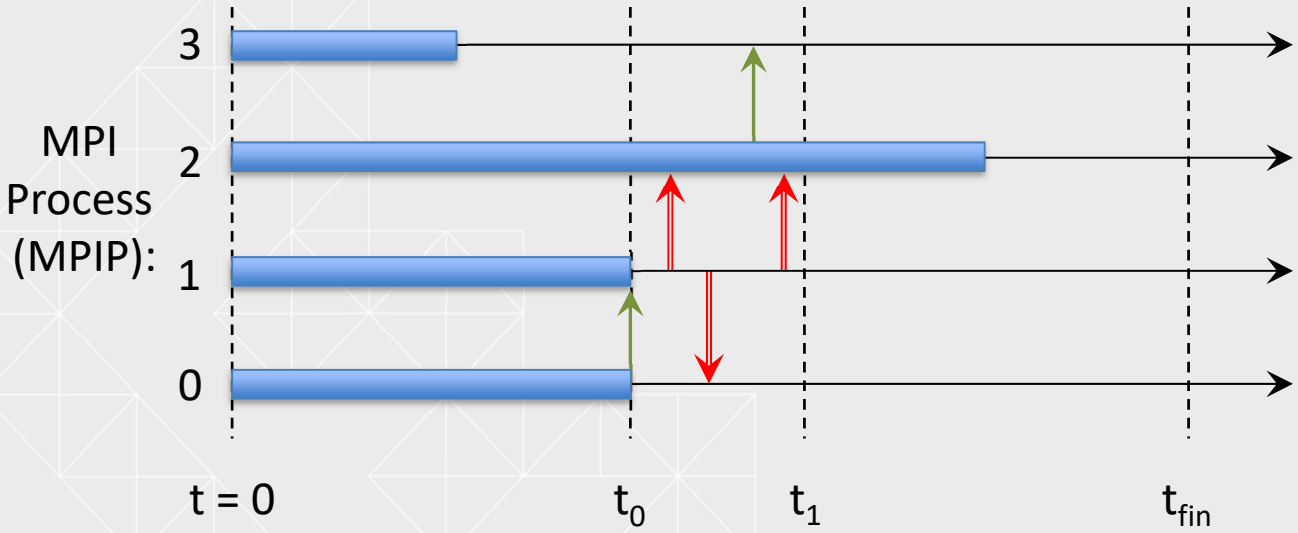

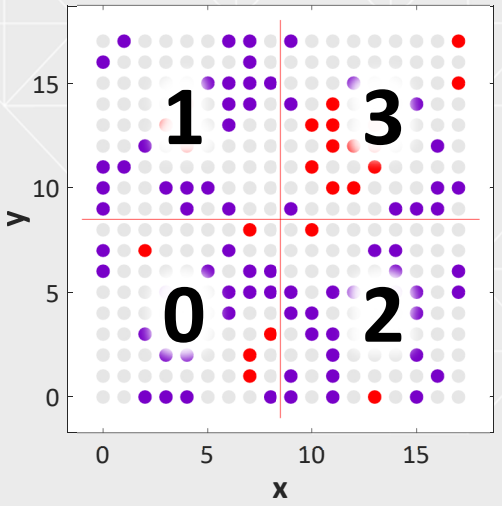

↑ : messages

↑ : anti-messages

- **MPIP 1 rolls-back to t_0** and sends **anti-messages** as necessary (instructing MPIPs to which it had sent messages to "undo things")

Maintaining Causality

- The asynchronous nature of the simulation can lead to **violations of causality**!

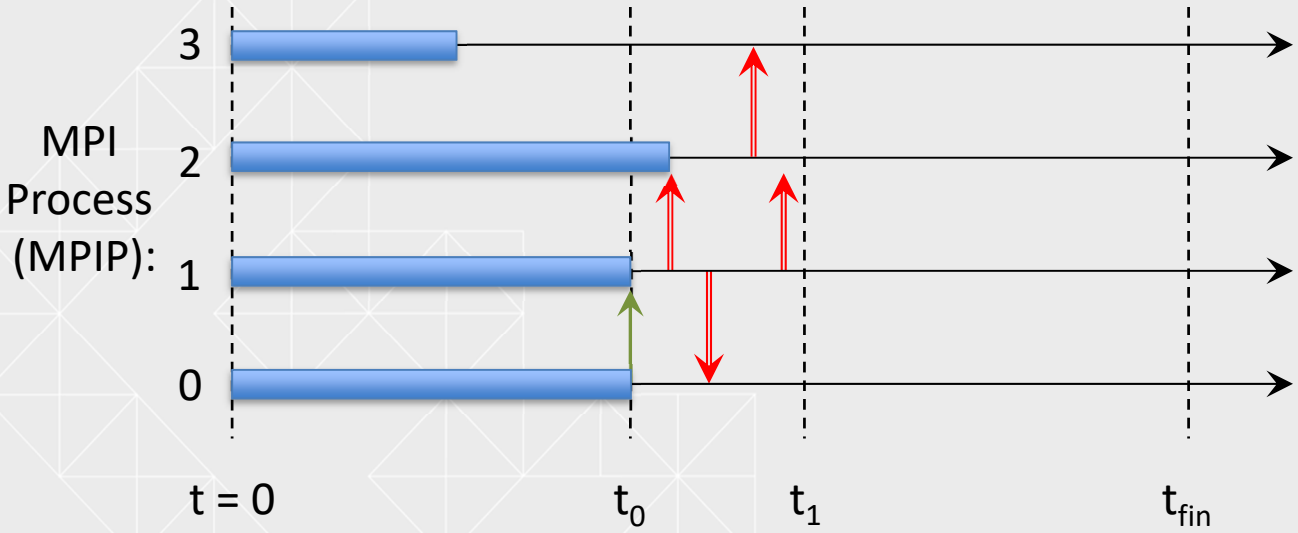

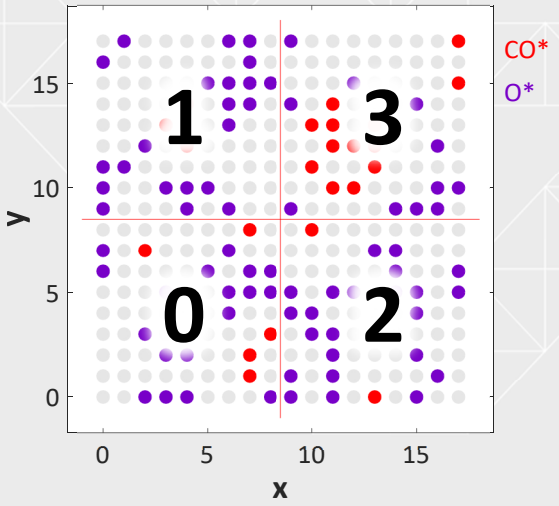

↑: messages

↑: anti-messages

- **MPIP 2 now has to roll-back too**, and send an anti-message to MPIP 3…
- **When does the roll-back cascade stop? How to keep track of rollbacks?**

The Time-Warp Algorithm

- Proposed by Jefferson in ’85 in his paper “Virtual Time”[1]
- **Underlying Principle**: *if event A causes event B, then* $t_{real,A} < t_{real,B}$
- Elegant algorithm to systematically deal with arbitrarily large cascades of rollbacks with **local operations**:
  - Taking **snapshots**
  - **Restoring** the state of the simulation at an earlier time
  - **Sending** and **receiving** messages or anti-messages
  - **Executing messaged actions**

… in addition to the usual KMC operations

[1] D. R. Jefferson, D. R. ACM Transactions on Programming Languages and Systems, 7: 404-425 (1985).

Time-Warp: Conceptual Implementation

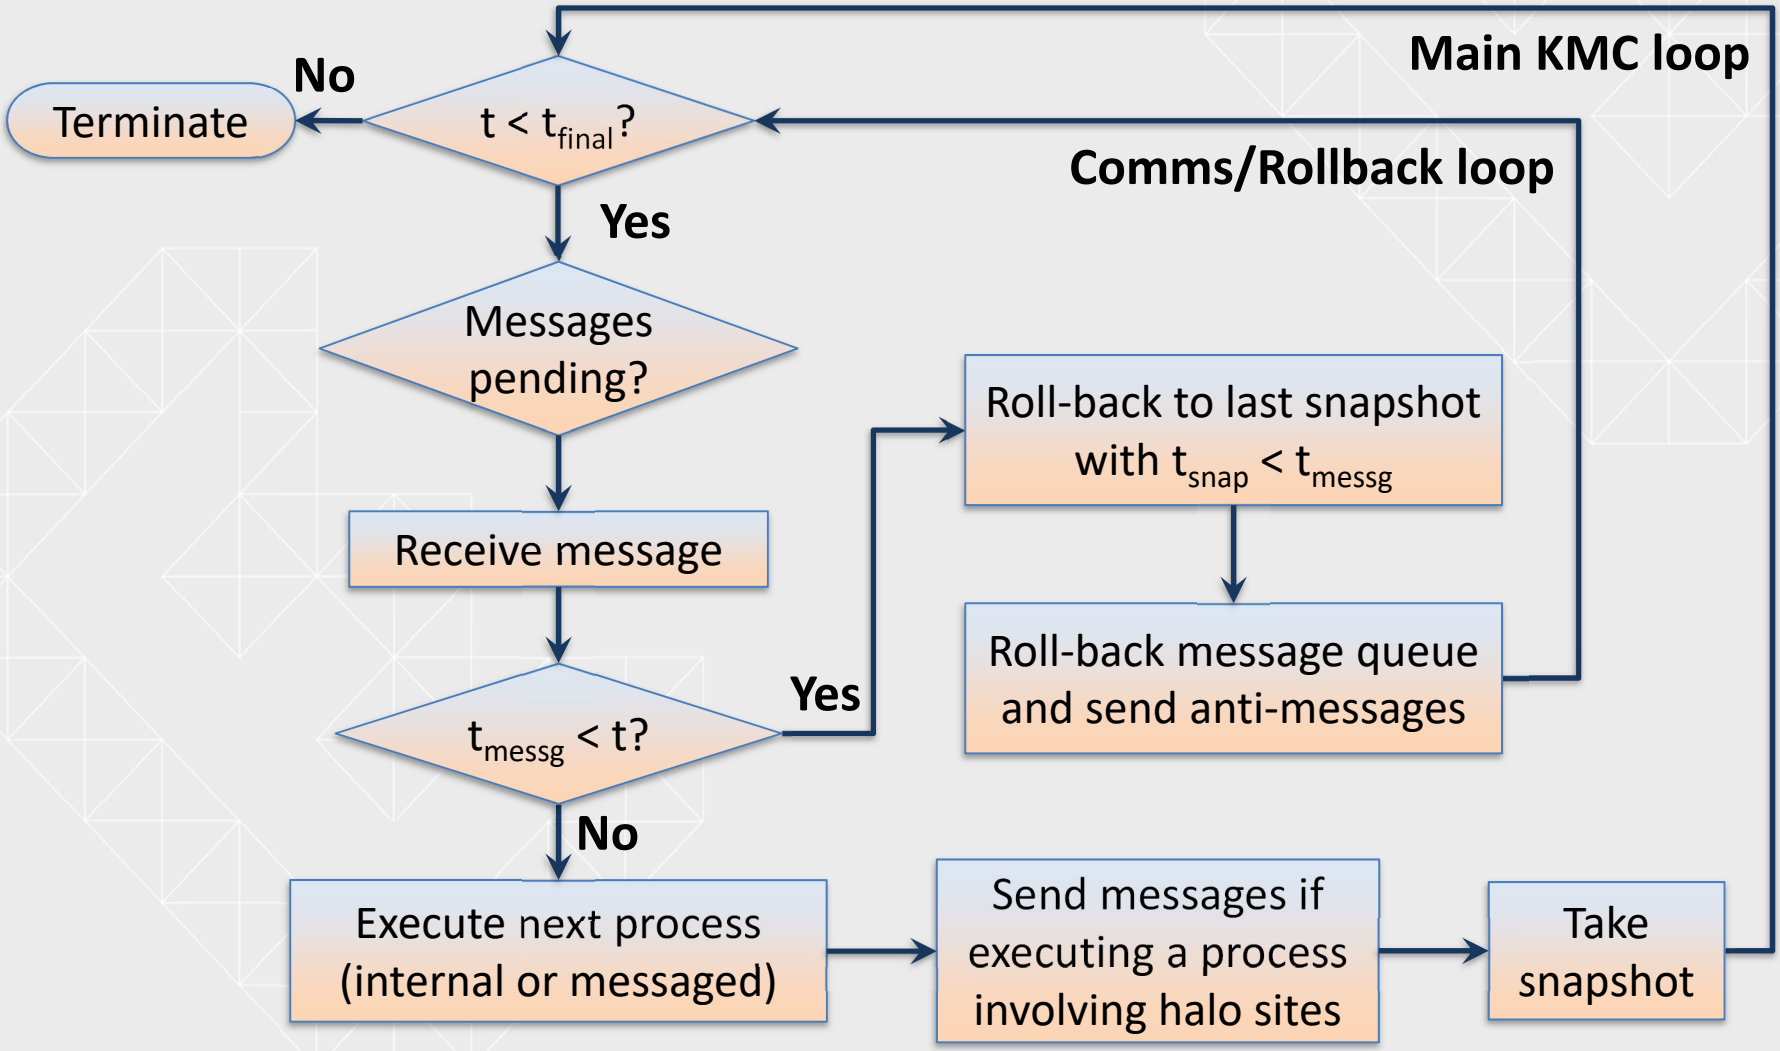

Validating the Implementation

- Recall the underlying principle of Time-Warp algorithm:
  - ✓ *if event A causes event B, then* $t_{real,A} < t_{real,B}$
- Enables construction of a **serial algorithm** that **emulates the MPI run**

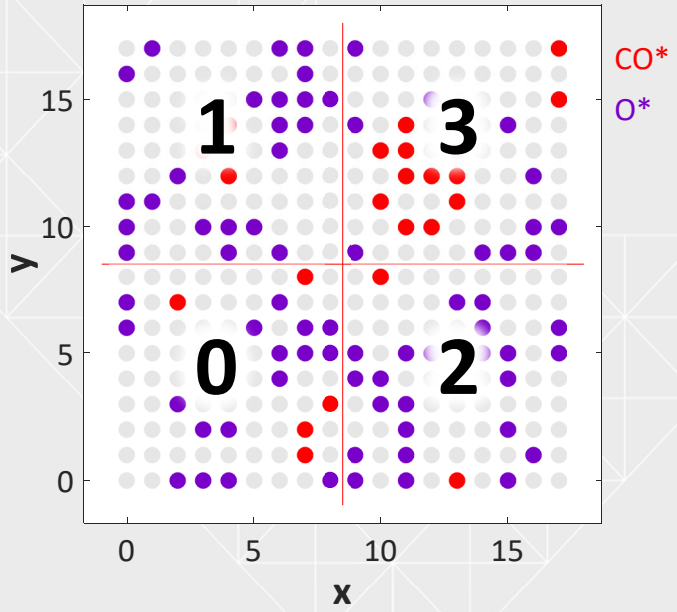

- Decompose the domain for *n* MPIPs, as in the actual MPI run
- Initialise *n* random sequences (identical to those used in the MPI run)
- Simulate most imminent event serially
- Use the appropriate random sequence when updating KMC state

Setup of Validation Simulations

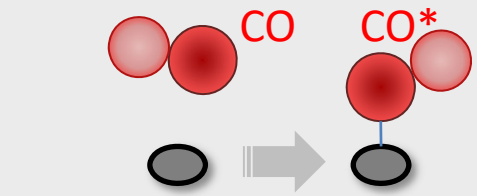

CO reversible adsorption

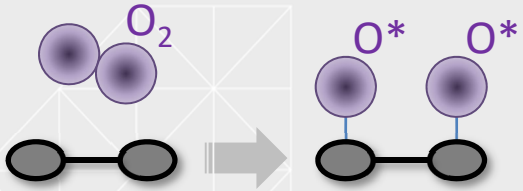

CO reversible adsorption

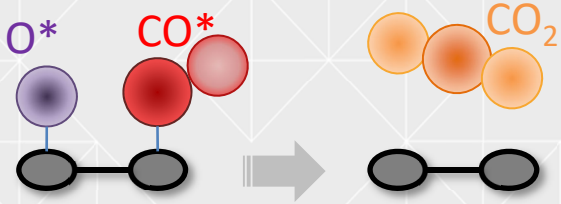

CO-O reversible oxidation

+ O* and CO* diffusions

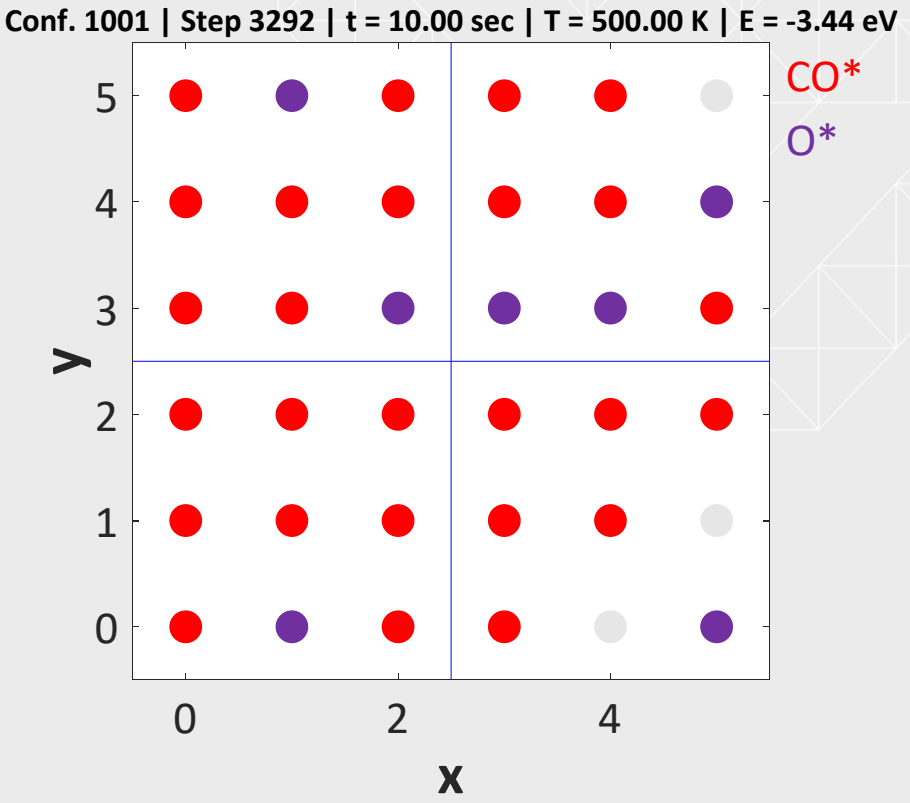

- CO oxidation simulations with 2 species, 10 elementary events and 2 energetic clusters, on a 6 × 6 lattice with 4 MPI processes

Validation Results

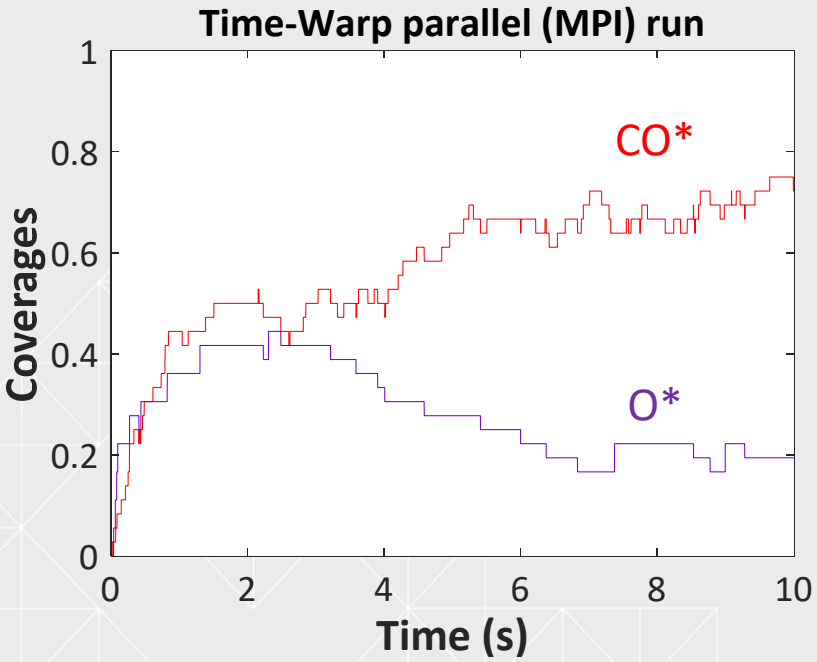

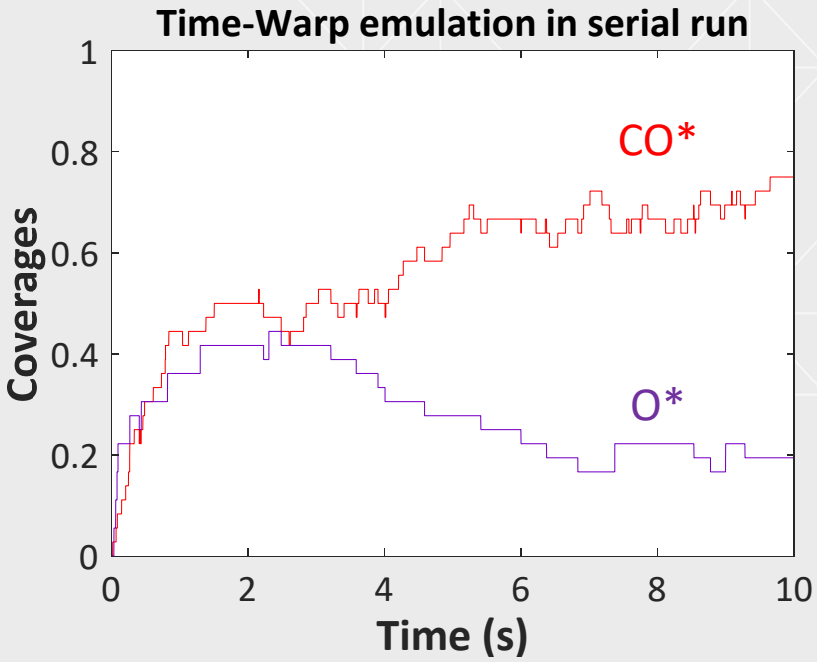

- Comparison of **Time-Warp MPI** and **parallel-emulation** runs:
  - ✓ **identical results**, down to the stochastic fluctuations!
- Also true for larger lattices and MPI process configurations

Validation Results

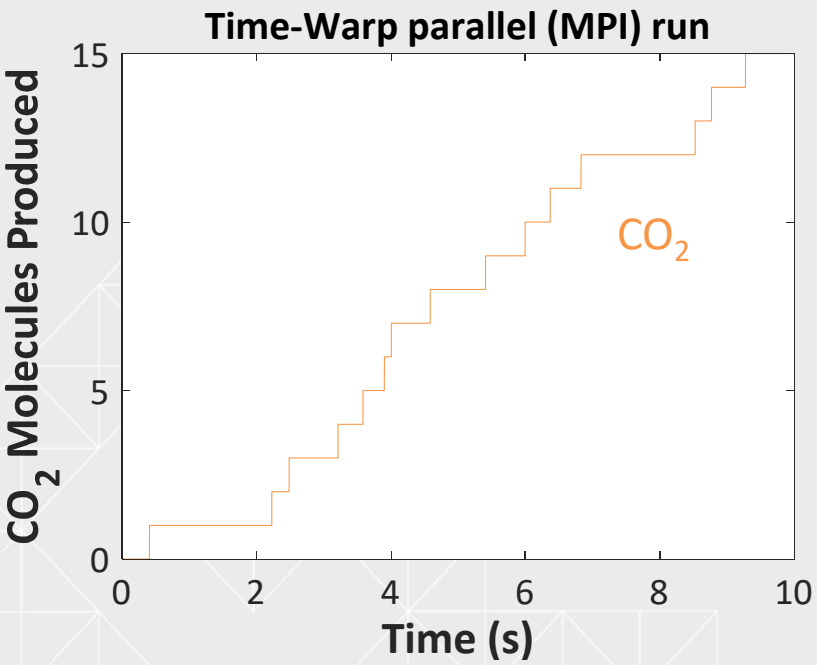

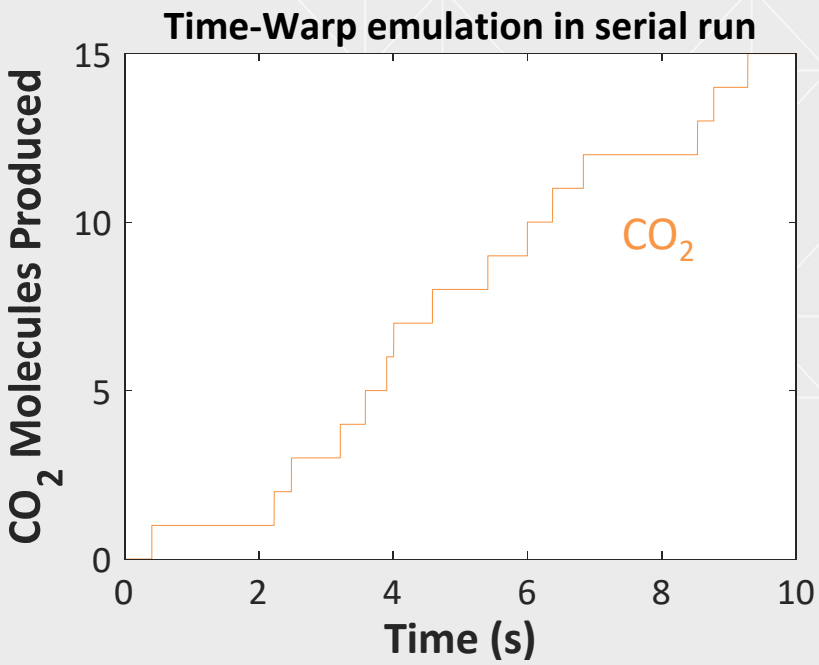

- Comparison of **Time-Warp MPI** and **parallel-emulation** runs:
  - ✓ **identical results**, down to the stochastic fluctuations!
- Also true for larger lattices and MPI process configurations

Performance Benchmarks

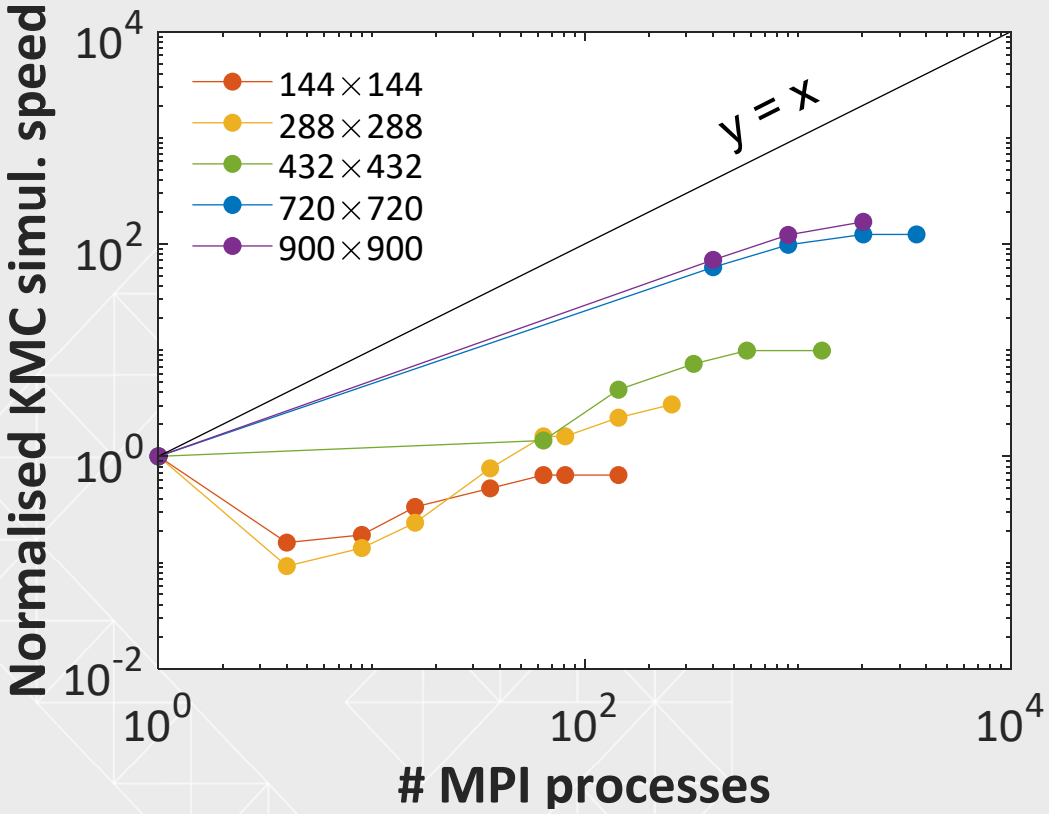

- Used system with reversible CO adsorption and diffusion
- KMC simulation speed

$$= \frac{\{\text{KMC time advancement}\}}{\{\text{required clock time}\}}$$

- Normalisation with respect to 1 MPI process run

- Time-Warp becomes progressively more efficient for larger lattices

Performance Benchmarks

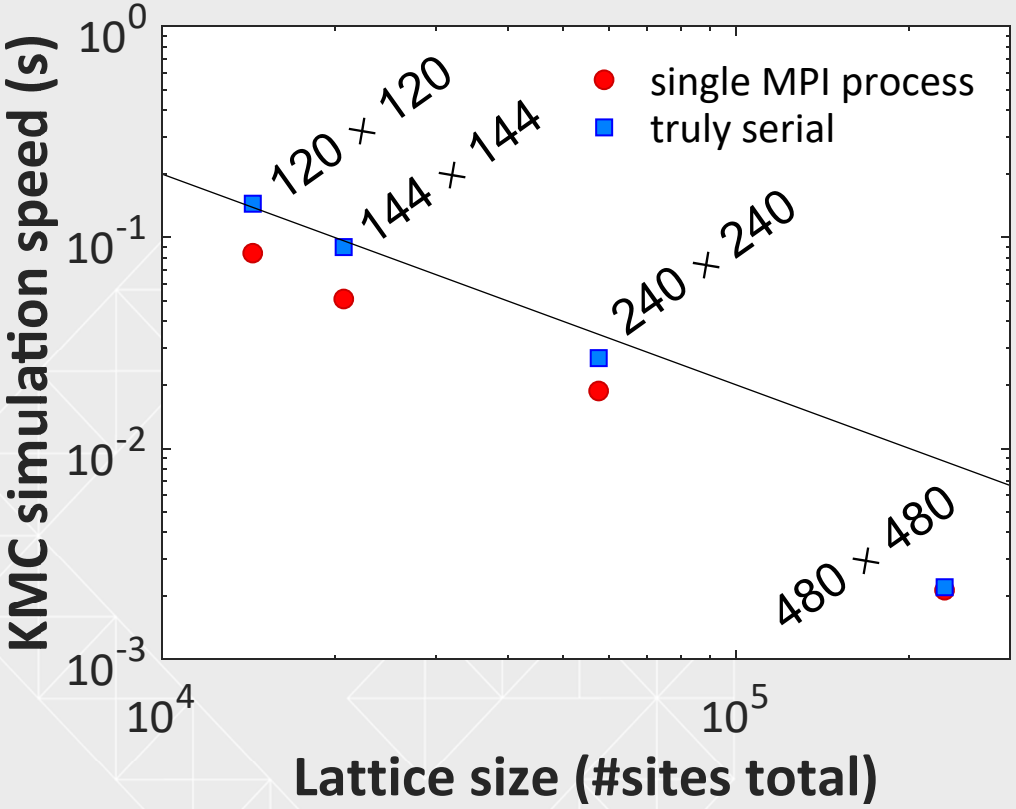

- Used system with reversible CO adsorption and diffusion
- KMC simulation speed

$$= \frac{\{\text{KMC time advancement}\}}{\{\text{required clock time}\}}$$

- Normalisation with respect to 1 MPI process run

- Overheads of Time-Warp algorithm become negligible for larger lattices

Conclusions

Kinetic Monte Carlo simulations:

- powerful approach towards understanding heterogeneous catalysts

Large-scale simulations needed:

- Higher accuracy predictions
- New physics: **reconstruction**, **pattern formation** on catalytic surfaces

KMC code Zacros Advanced lattice-KMC made easy **now implements Time Warp algorithm:**

- **First-of-its kind prototype** for massively parallel simulations (MPI)
- **Exact algorithm** yielding **reproducible** and **validated** output

Acknowledgments

- **Research Software Development Group** (UCL)
  Ilektra Christidi, Roland Guichard, Jens Nielsen
- Srikanth Ravipati, Miguel Pineda, Giannis Savva

➢ Funding:

EPSRC
Engineering and Physical Sciences Research Council

LEVERHULME TRUST

epcc

European Commission

➢ Computational Resources:

Grace Myriad @ UCL

archer