Speculatively Parallelized Temperature Accelerated Dynamics (SpecTAD)

Arthur F. Voter

co-authors: Richard J. Zamora, Danny Perez

Theoretical Division
Los Alamos National Laboratory
Los Alamos, NM, USA

BRITS16 – Bridging time scale techniques…
Max Planck Institute for Physics of Complex Systems
Dresden, Germany
September 12, 2016

Acknowledgments

Stephan Eidenbenz (ISTI, LANL)
Nandu Santhi (CCS, LANL)
Sunil Thulasidasan (CCS, LANL)
Sue Mniszewski (CCS, LANL)

DOE Office of Basic Energy Sciences
Los Alamos LDRD

Accelerated molecular dynamics approach for infrequent-event systems

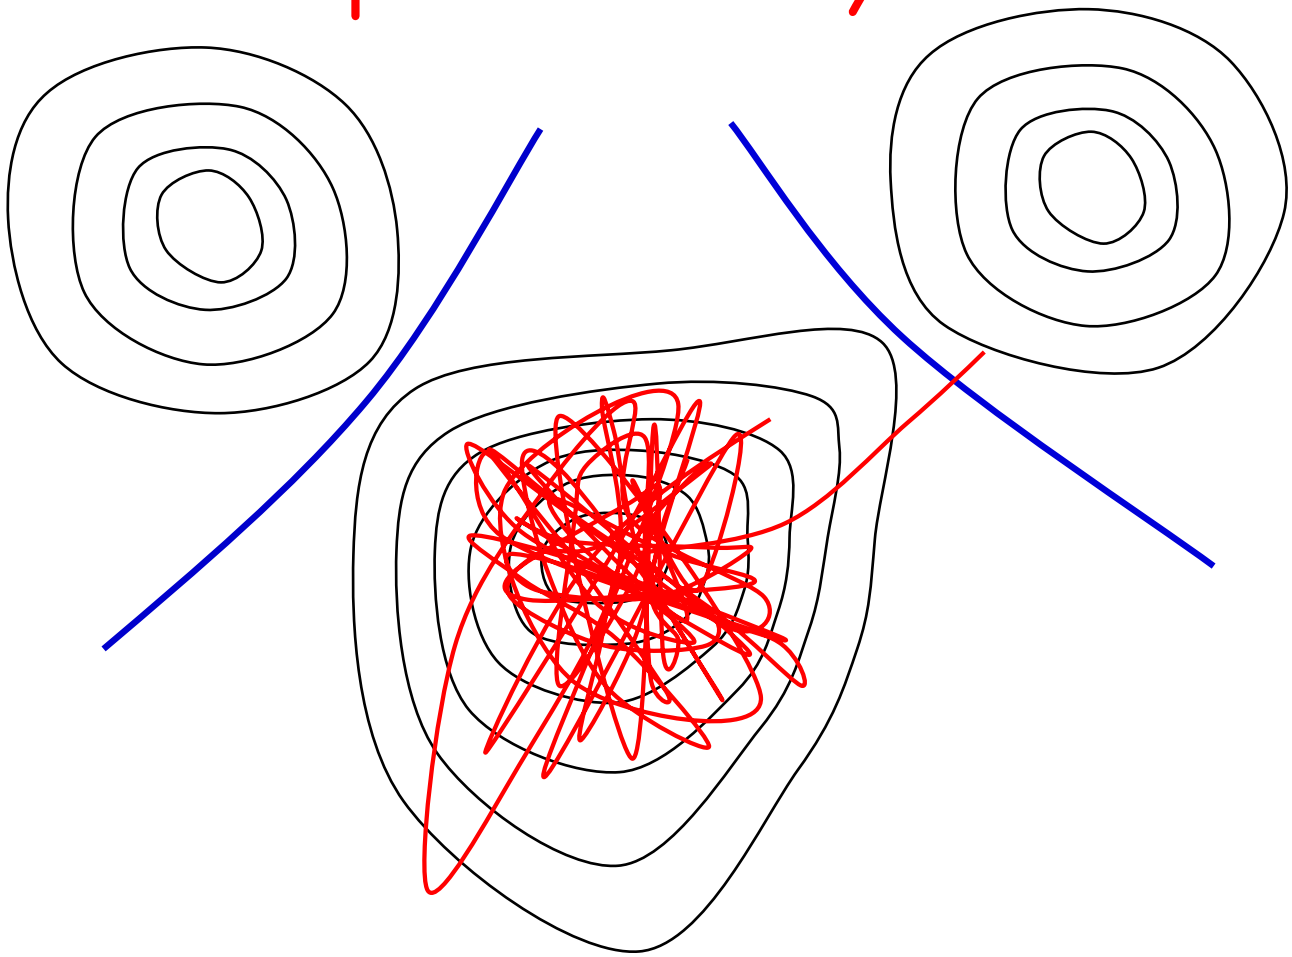

The trajectory finds an appropriate way out (i.e., proportional to the rate constant) without knowing about any of the escape paths except the one it first sees.

Accelerated molecular dynamics approach for infrequent-event systems

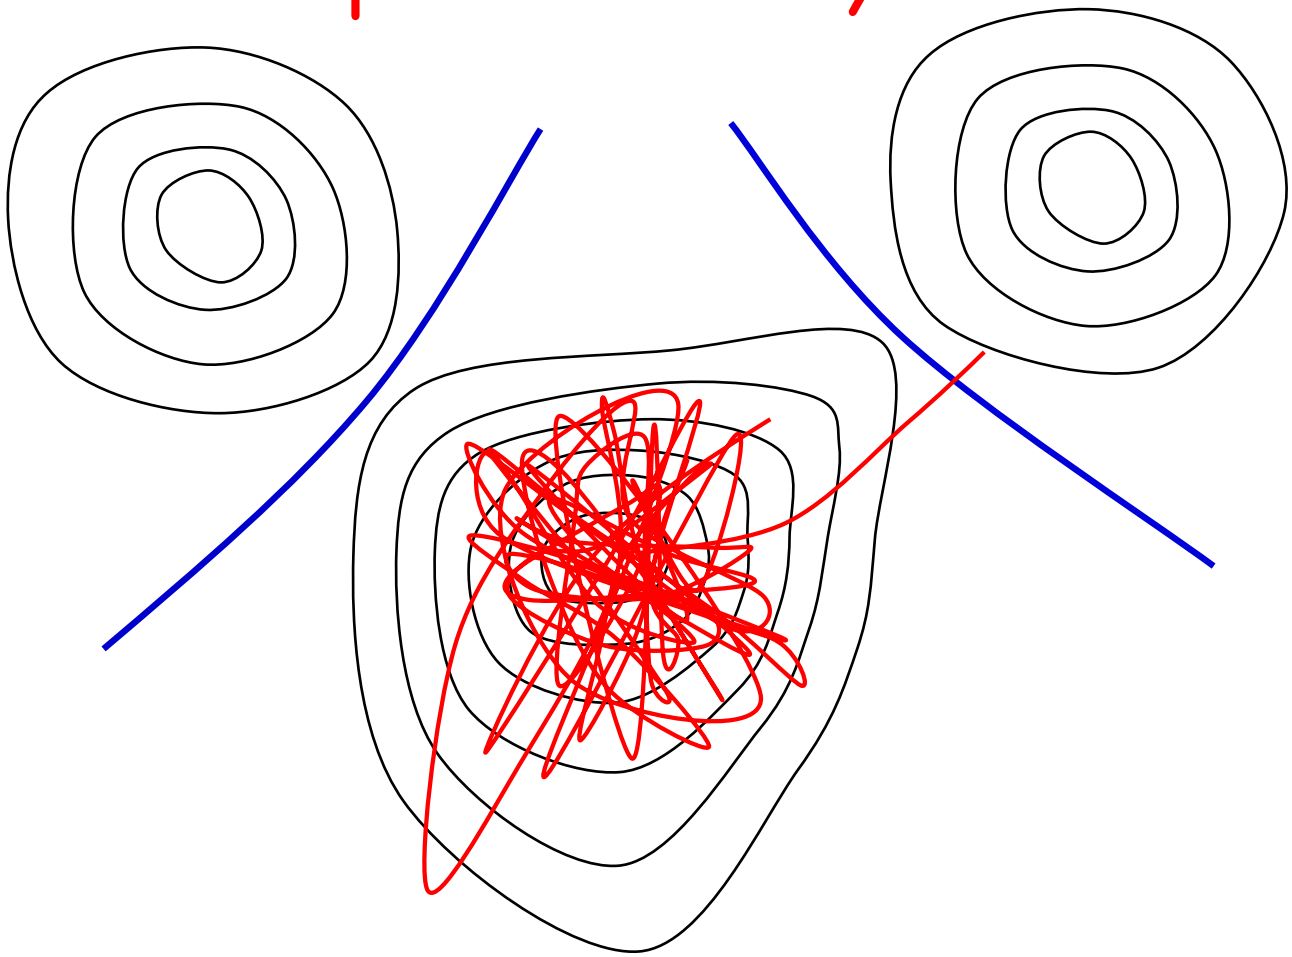

The trajectory finds an appropriate way out (i.e., proportional to the rate constant) without knowing about any of the escape paths except the one it first sees. In AMD, we trick the system into doing this more quickly, trying very hard not to corrupt the relative escape probabilities.

Temperature Accelerated Dynamics (TAD)

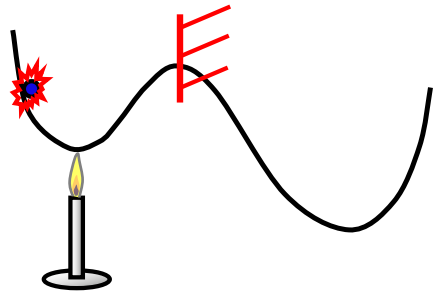

[Sørensen and Voter, J. Chem. Phys. 112, 9599 (2000)]

Temperature Accelerated Dynamics (TAD)

Concept:

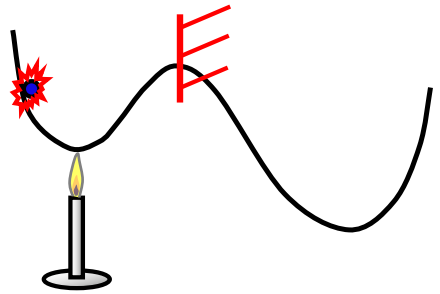

Raise temperature of system to make events occur more frequently. Filter out the events that should not have occurred at the lower temperature.

[Sørensen and Voter, J. Chem. Phys. 112, 9599 (2000)]

Temperature Accelerated Dynamics (TAD)

Concept:

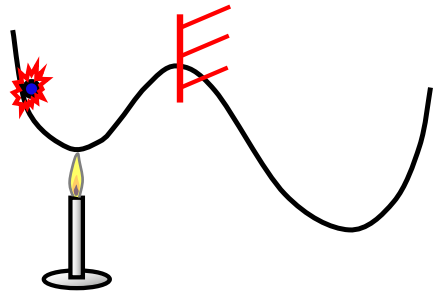

Raise temperature of system to make events occur more frequently. Filter out the events that should not have occurred at the lower temperature.

Assumptions:

- infrequent-event system
- transition state theory (no correlated events)
- harmonic transition state theory (gives Arrhenius behavior)

$$k = \nu_0 \exp[-\Delta E/k_B T]$$

- all preexponentials (ν_0) are greater than ν_{min}

[Sørensen and Voter, J. Chem. Phys. 112, 9599 (2000)]

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

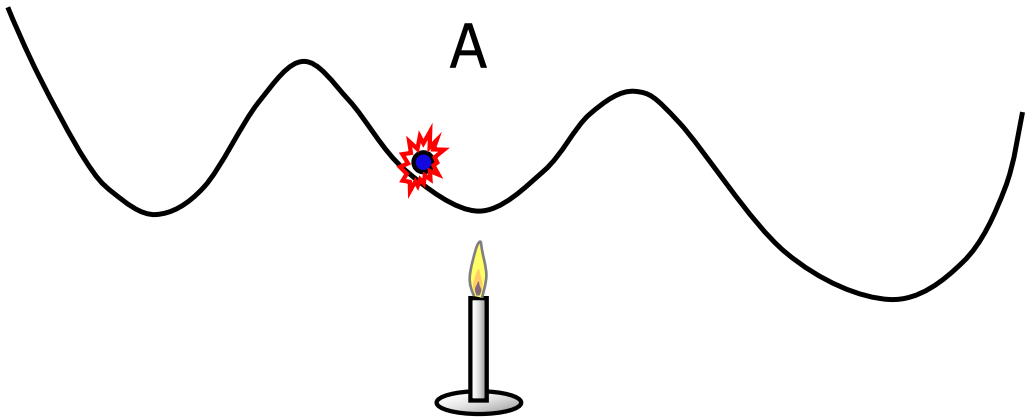

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

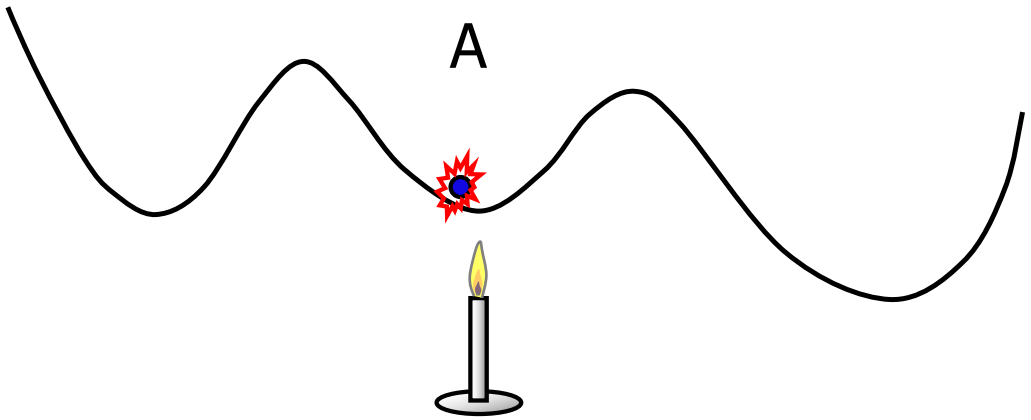

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

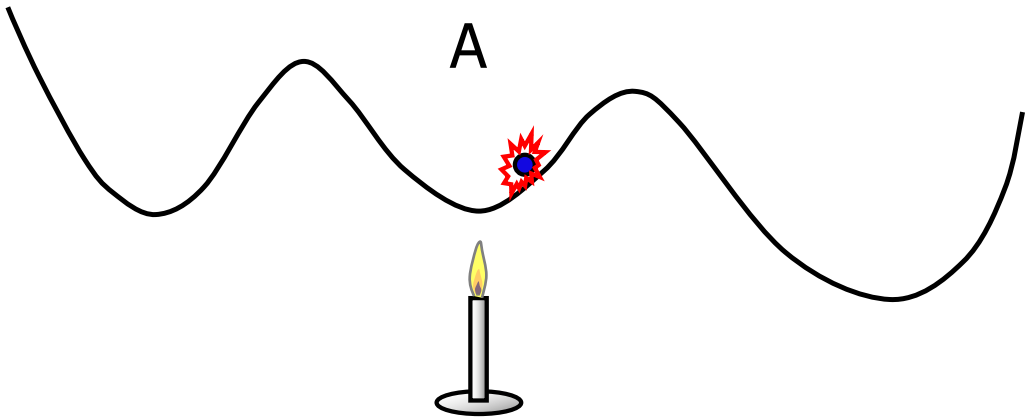

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

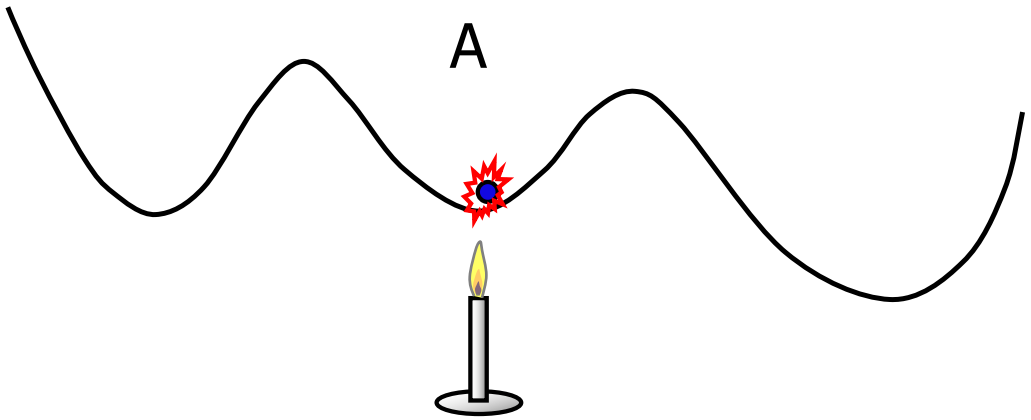

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

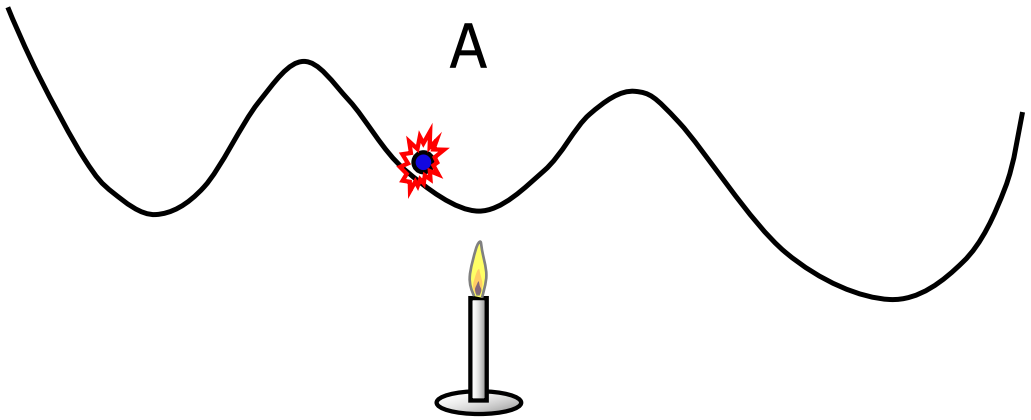

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

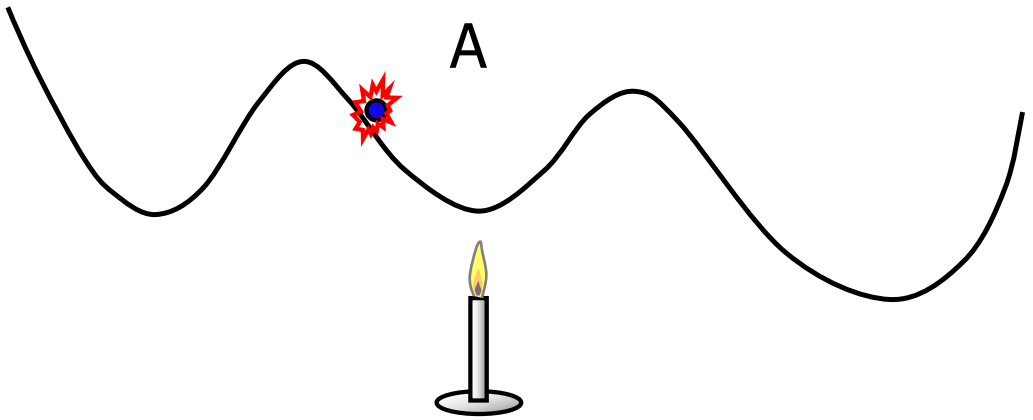

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

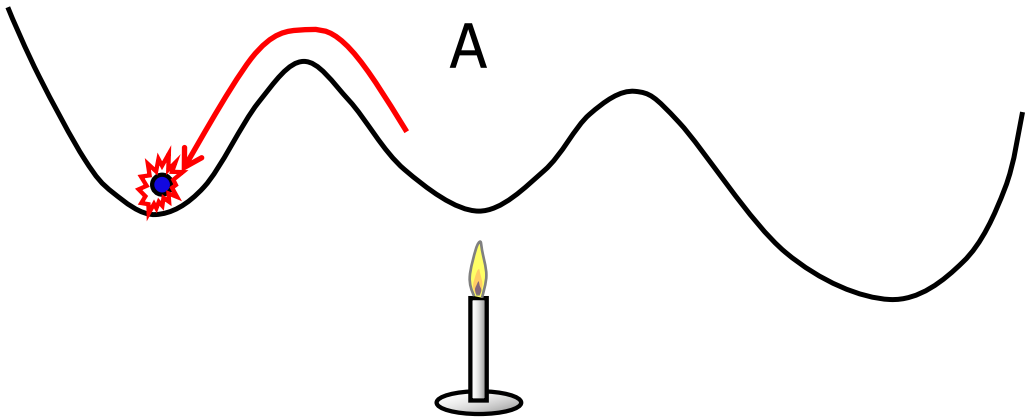

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

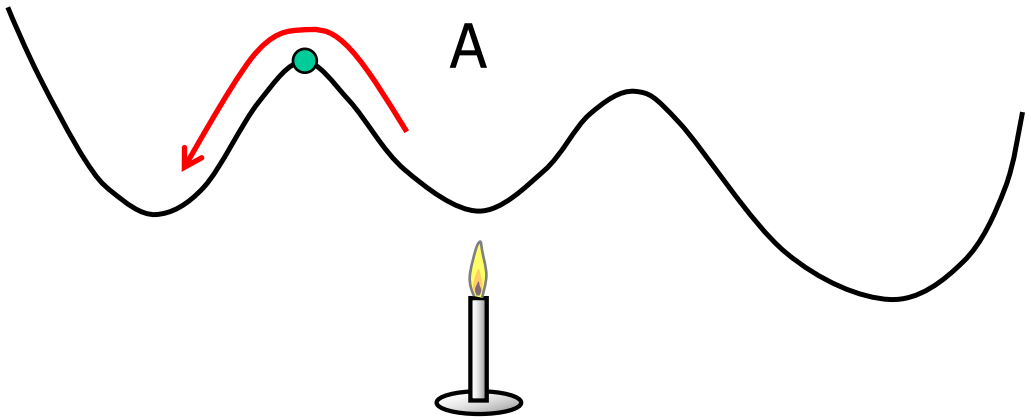

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

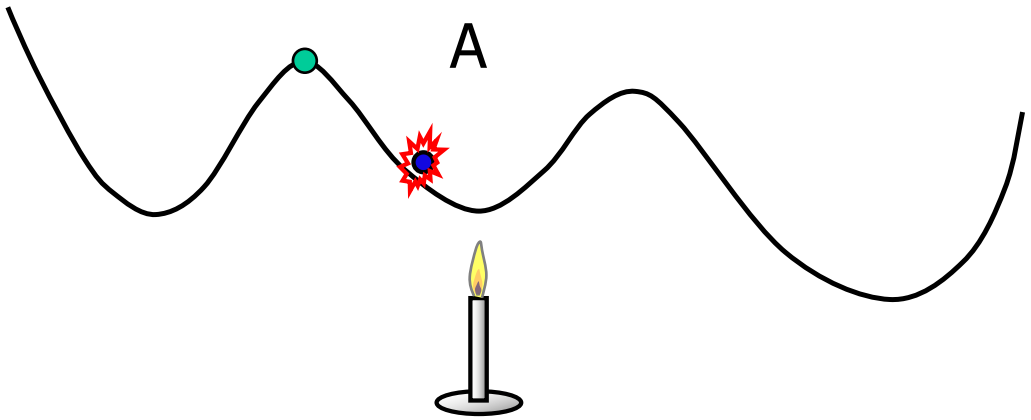

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

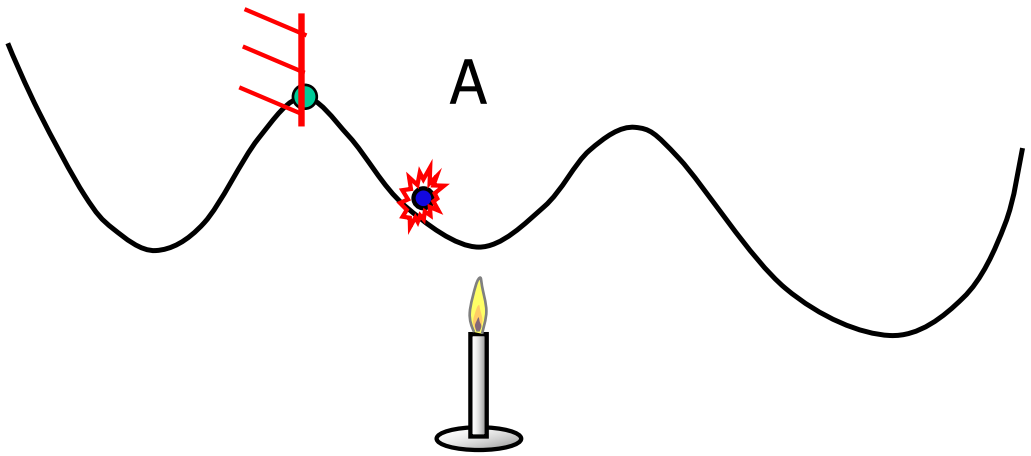

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

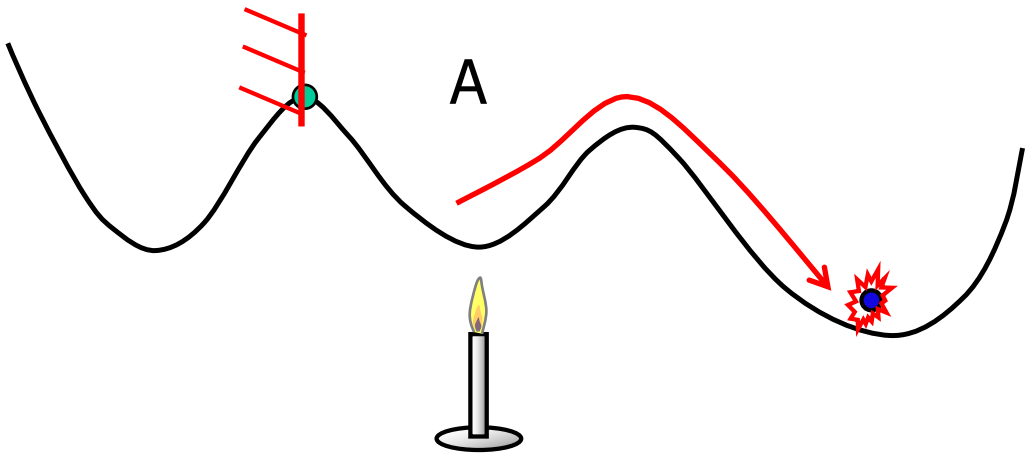

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

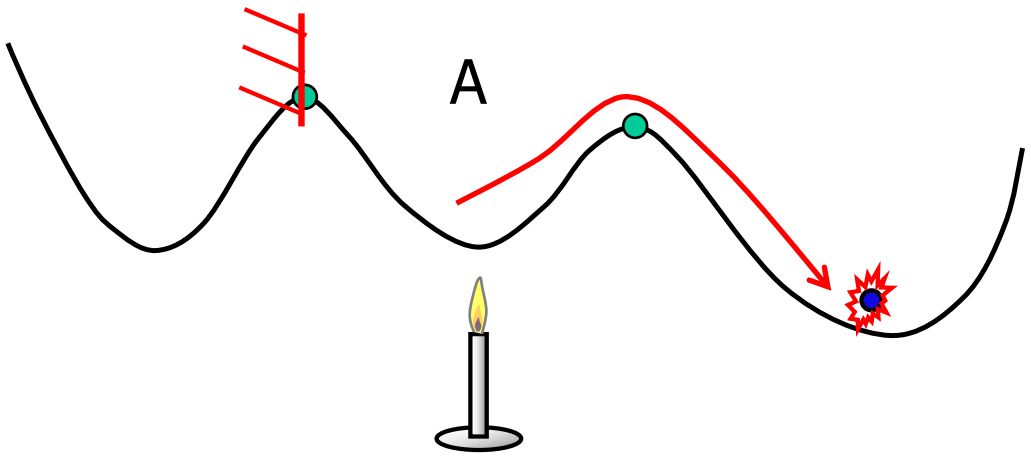

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

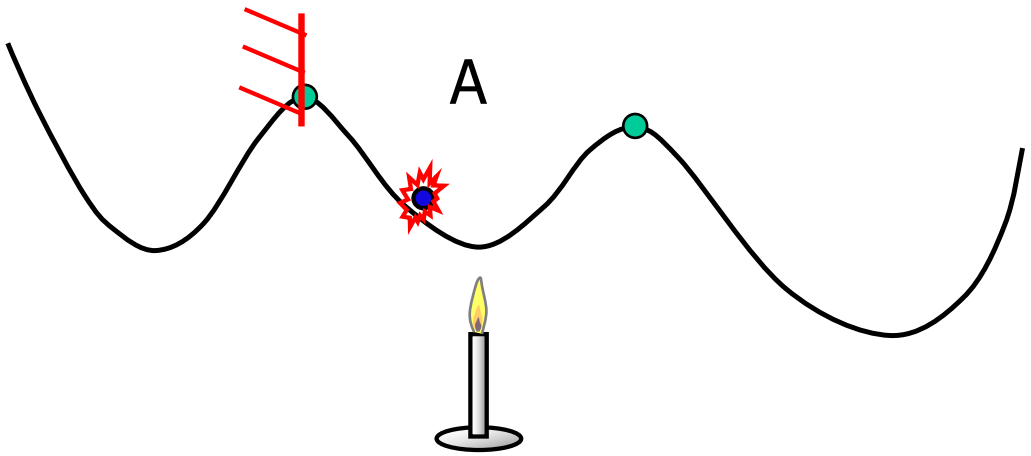

TAD Procedure

- Run MD at elevated temperature (T_{high}) in state A.
- Intercept each attempted escape from basin A
    - find saddle point (and hence barrier height)
      (e.g., using nudged elastic band method of Jonsson et al).
    - extrapolate to predict event time at T_{low}.
- Reflect system back into basin A and continue.
- When safe, accept transition with shortest time at T_{low}.
- Go to new state and repeat.

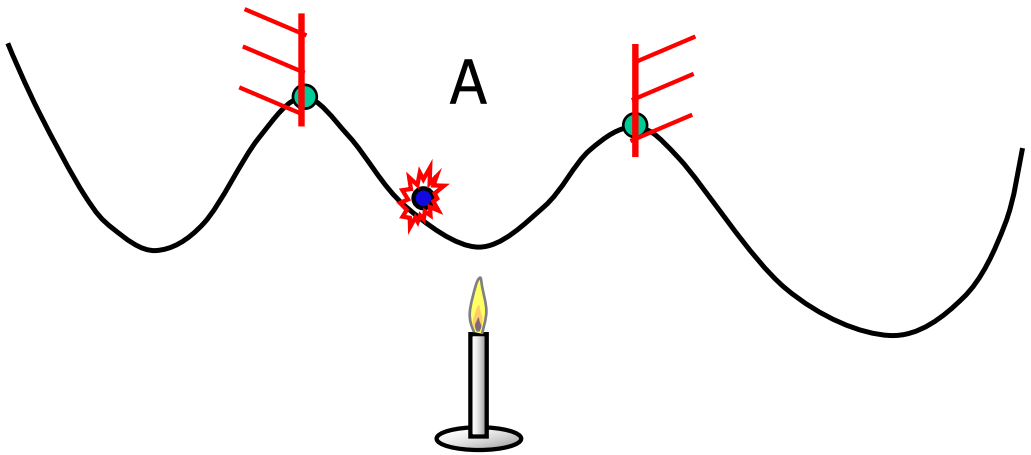

TAD temperature-extrapolated time

Because each rate is assumed to be Arrhenius,

$$k = \nu_0 \exp[-\Delta E/k_B T] ,$$

the time for each particular event at high T can be extrapolated to low T:

$$t_{low} = t_{high} \exp[\Delta E(1/k_B T_{low} - 1/k_B T_{high})] .$$

This time is sampled correctly from the exponential distribution at low T, mapped from the high T sample:

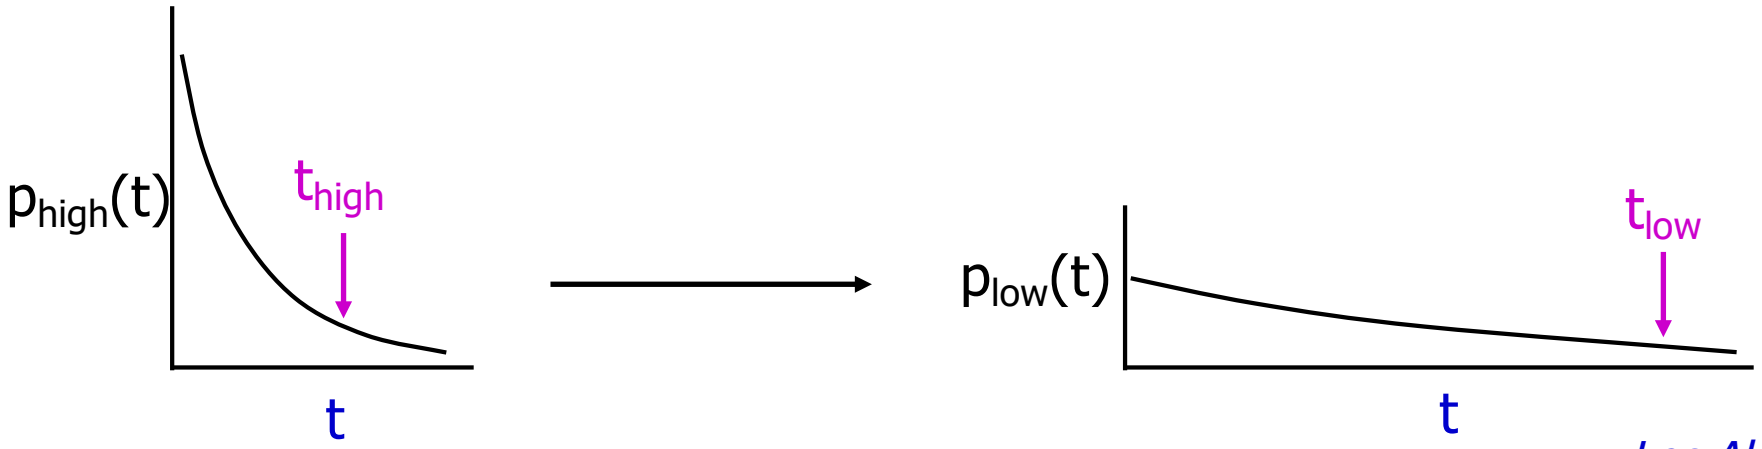

The Arrhenius view

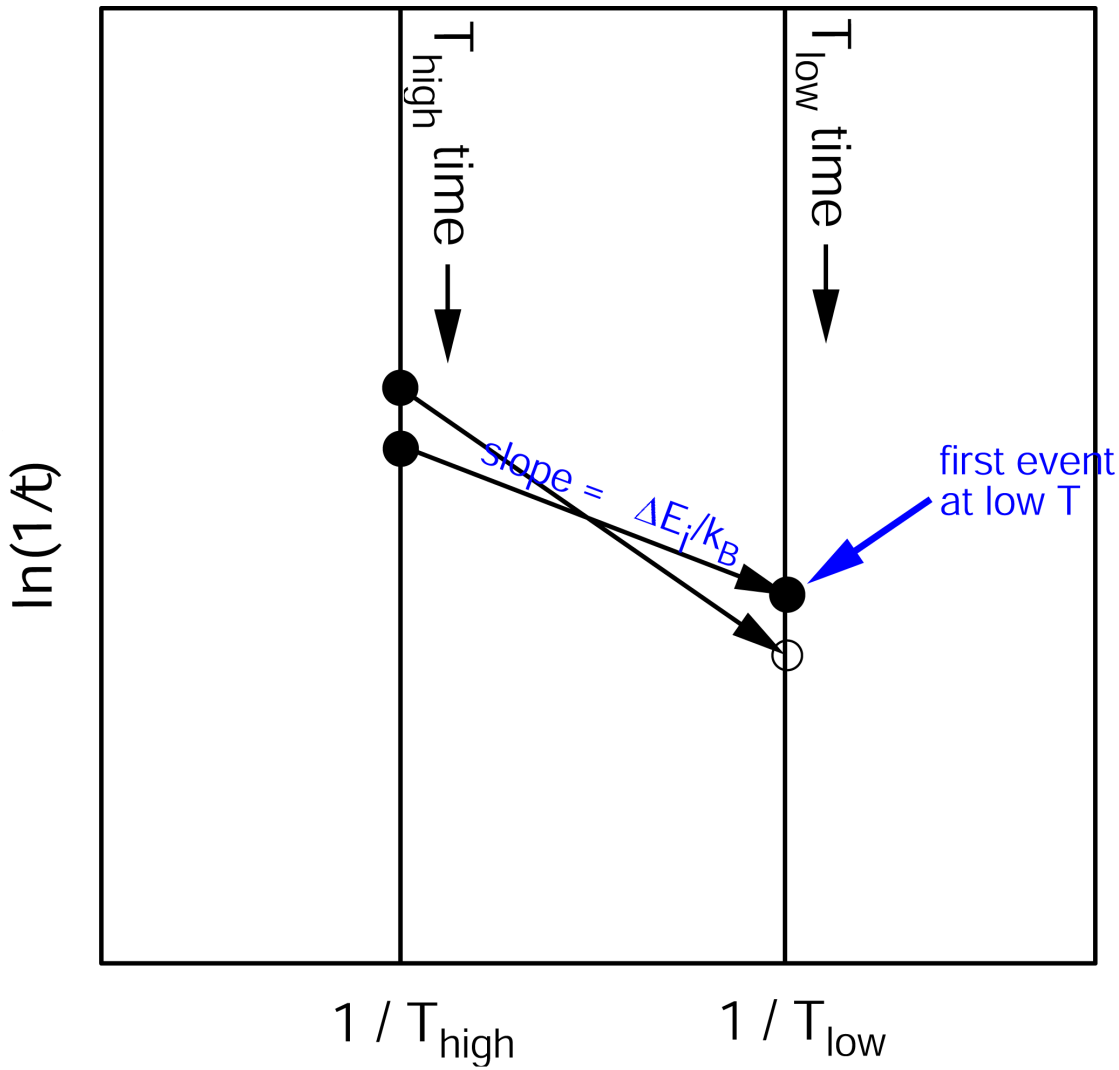

The Arrhenius view

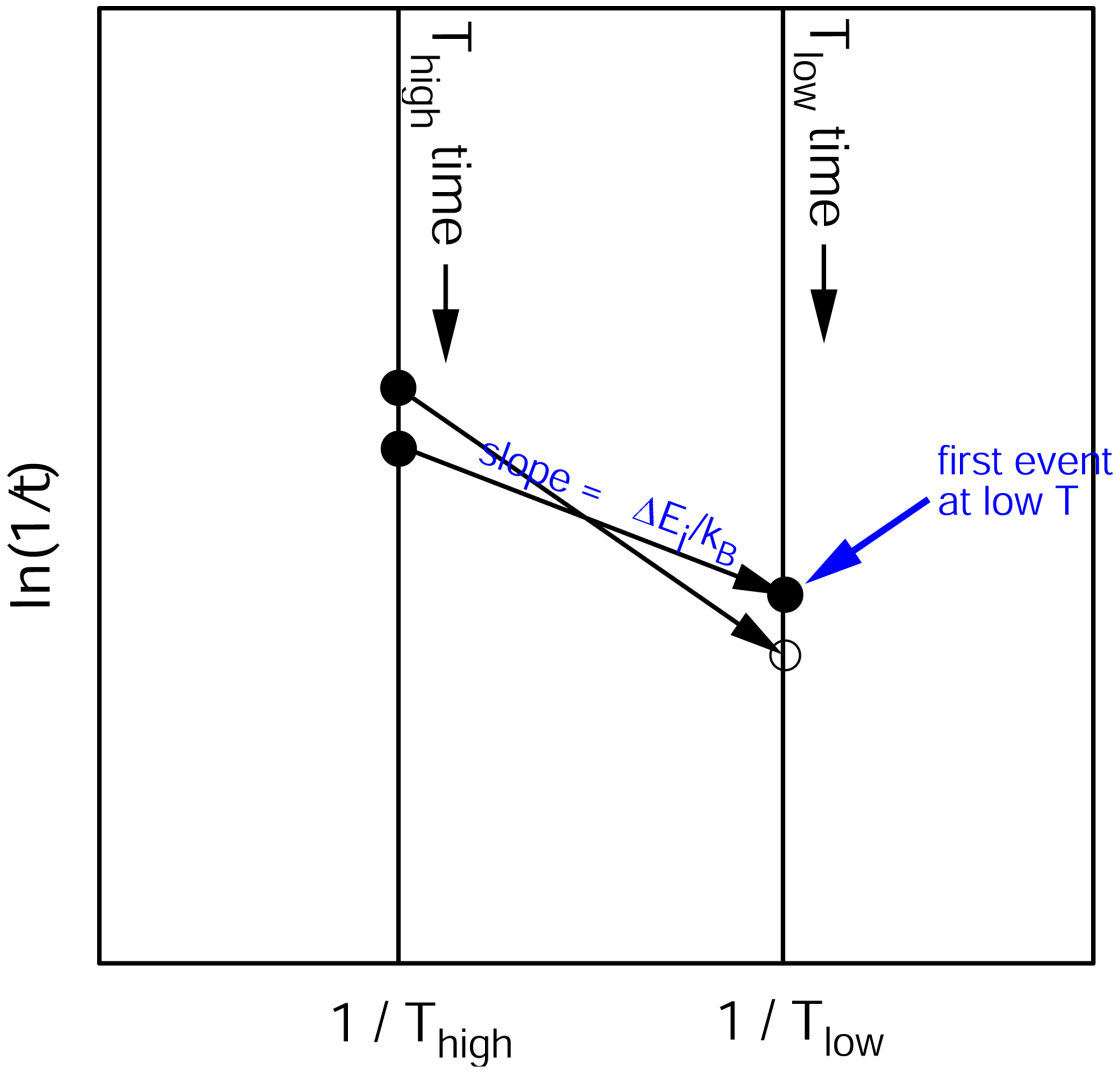

when can we stop?

The confidence line

For a pathway with rate k, the MD time τ required to be certain with confidence $1-\delta$ that at least one escape will occur is

$$\tau = (1/k) \ln(1/\delta)$$

The confidence line

For a pathway with rate k, the MD time τ required to be certain with confidence $1-\delta$ that at least one escape will occur is

$$\tau = (1/k)\ \ln(1/\delta)$$

For an Arrhenius rate, $k = \nu_0 \exp(-E_a/k_B T)$, all but fraction δ of the first escapes will occur above the line with slope E_a and intercept $\ln[\nu_0/\ln(1/\delta)]$

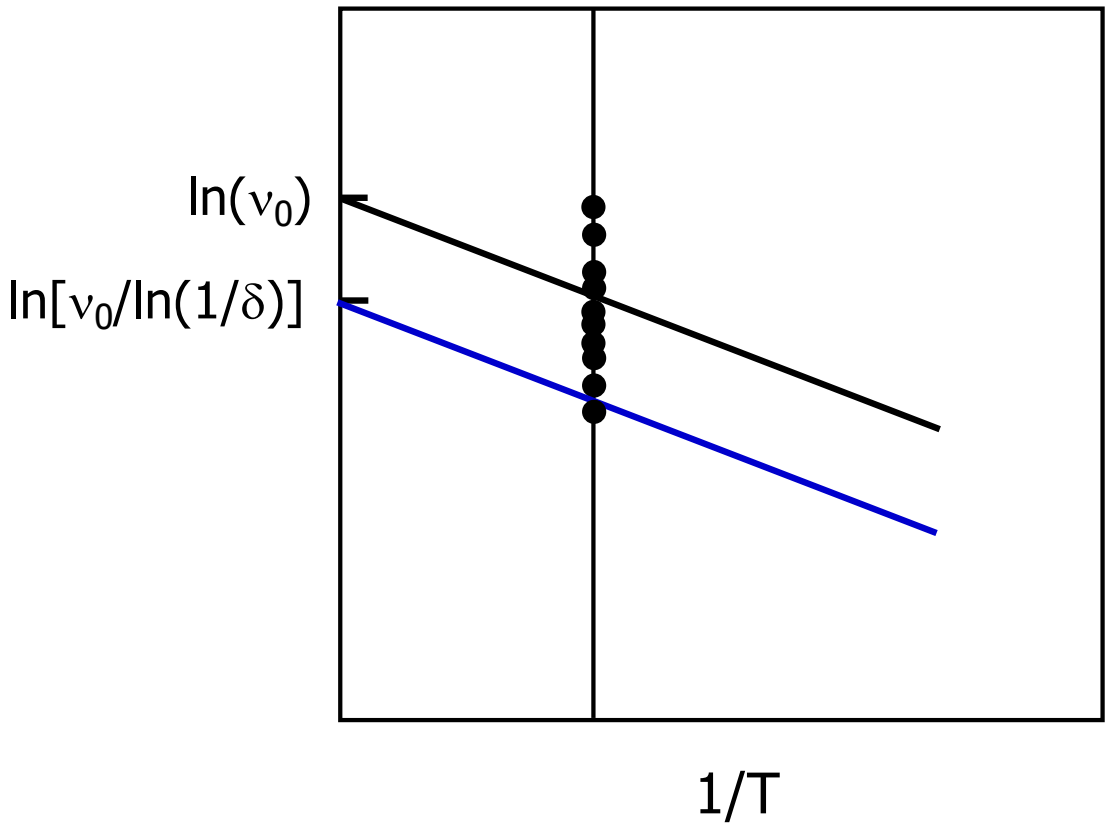

TAD – when can we stop the MD and accept an event?

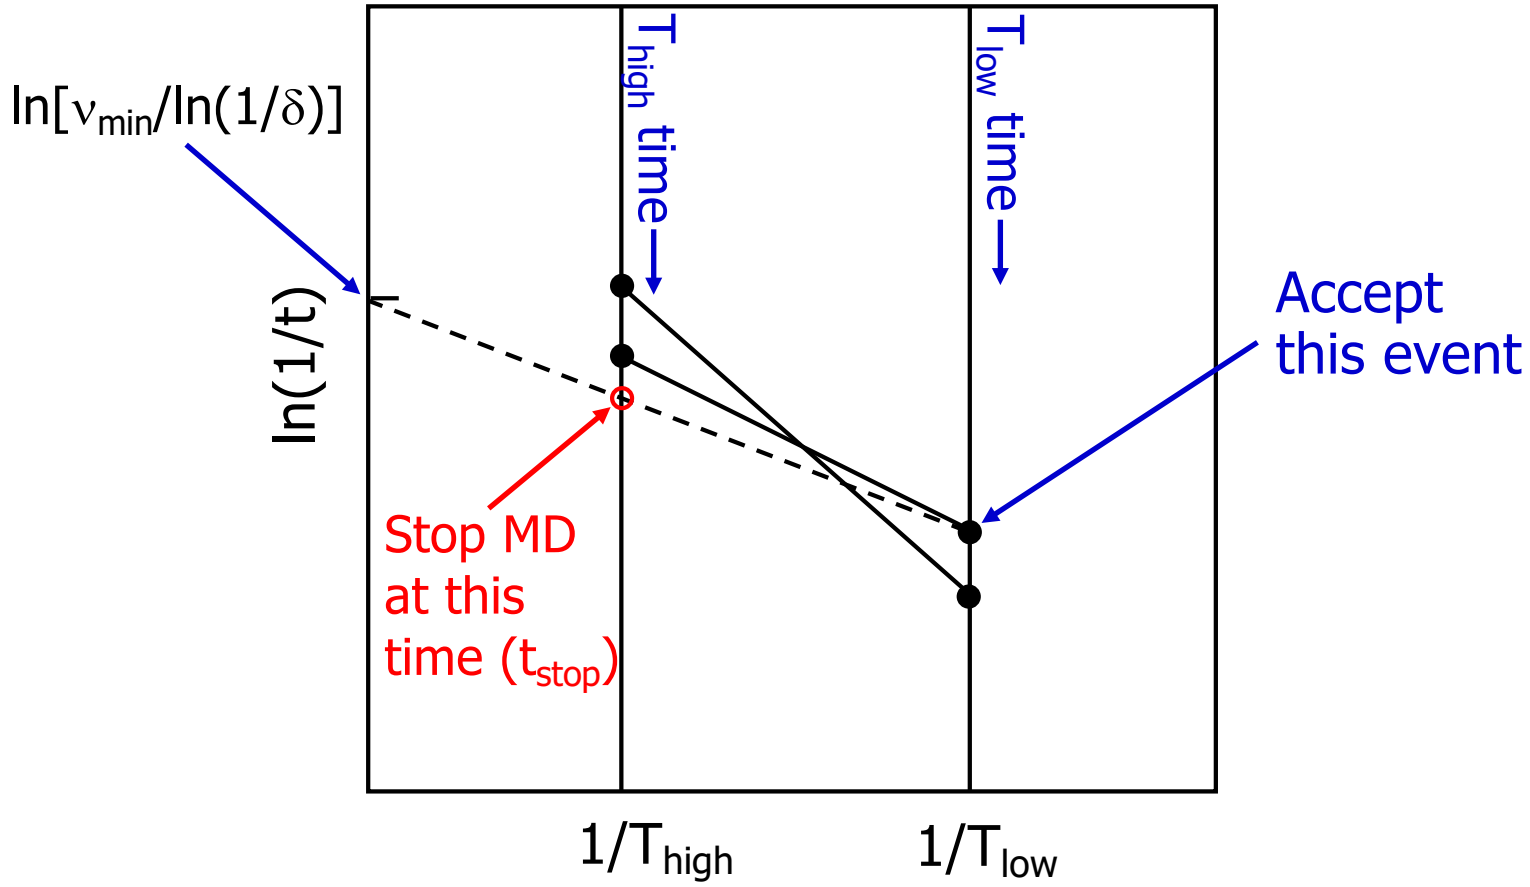

After time t_{stop}, with confidence 1-δ, no event can replace shortest-time event seen at low T.

Move system to this state and start again.

Correct dynamics, assuming harmonic TST, ν_{min}, uncertainty δ.

Effects of Anharmonicity

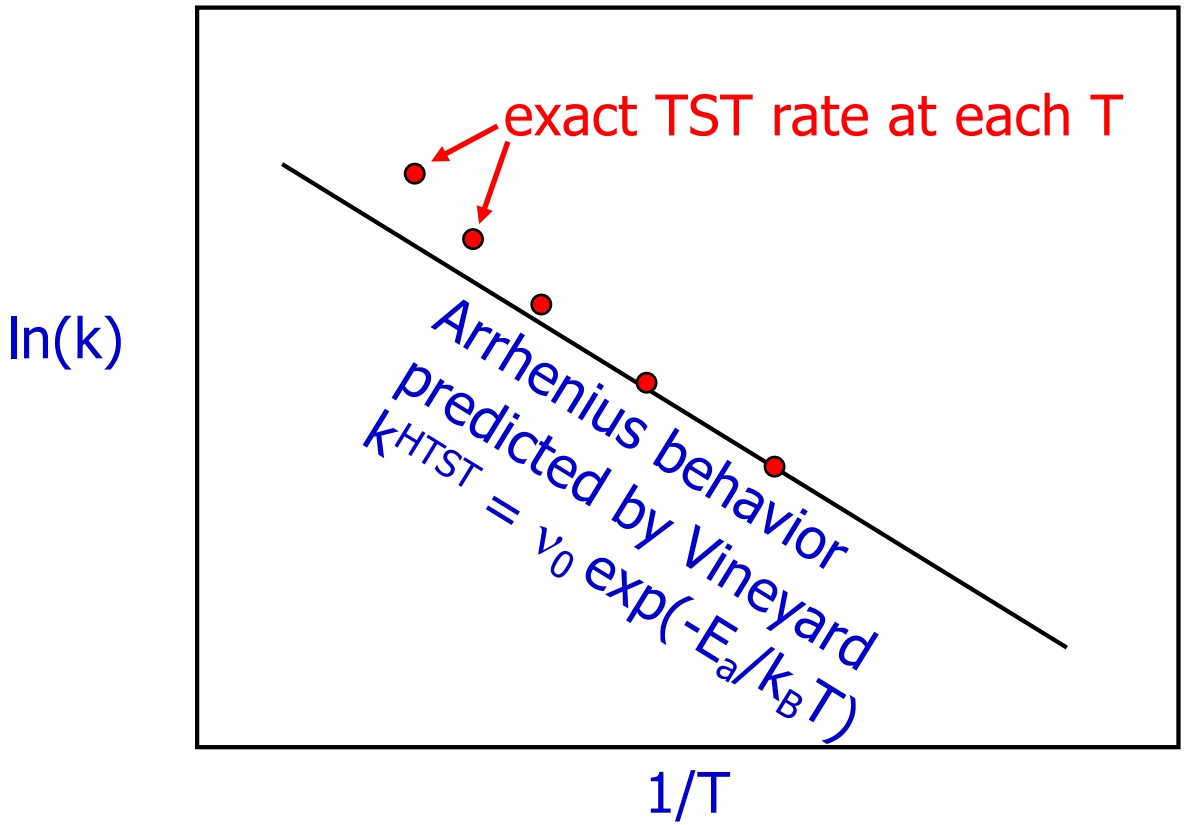

This error will be carried down to the low temperature

TAD procedure

● = current t_{high}

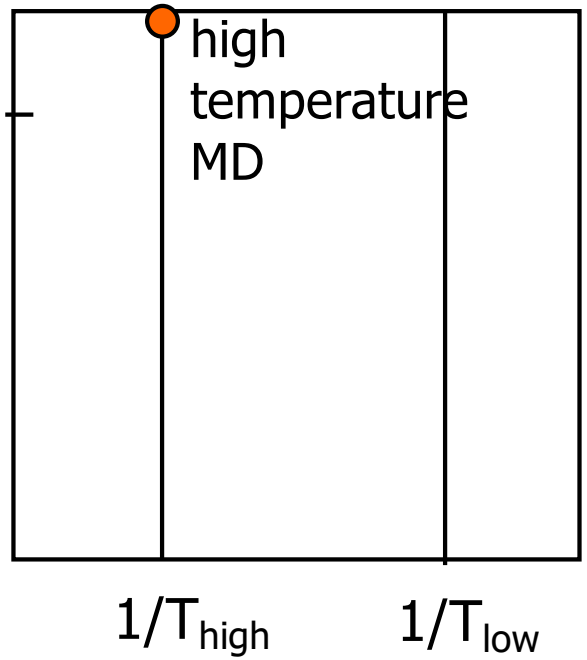

TAD procedure

● = current t_{high}

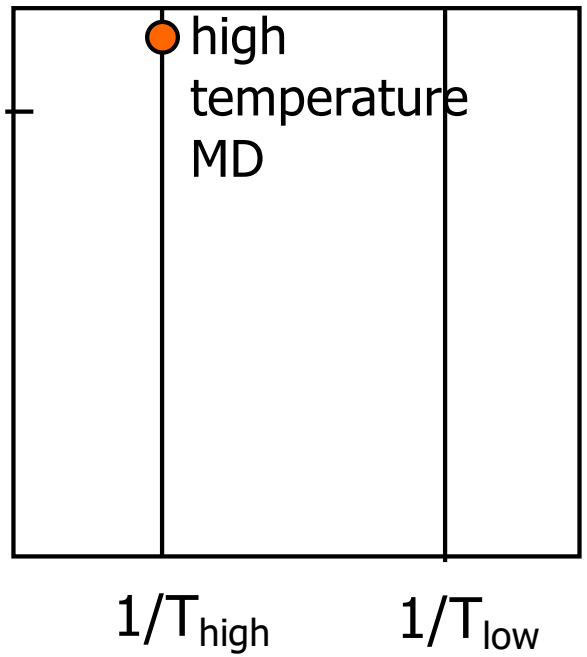

TAD procedure

● = current t_{high}

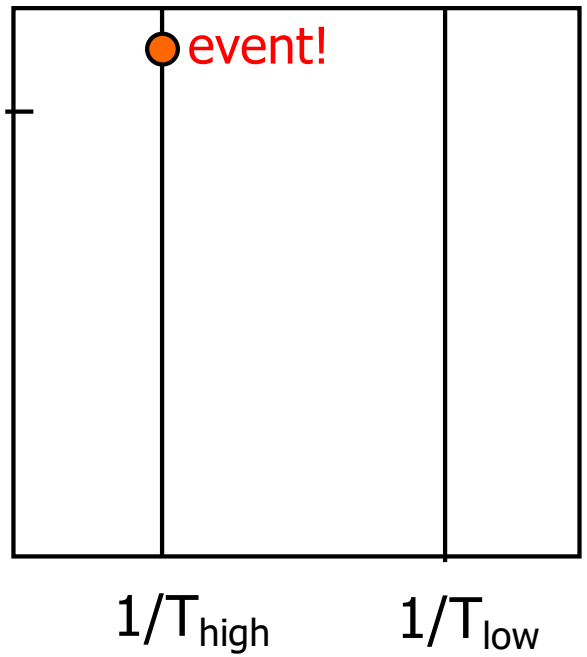

TAD procedure

● = current t_{high}

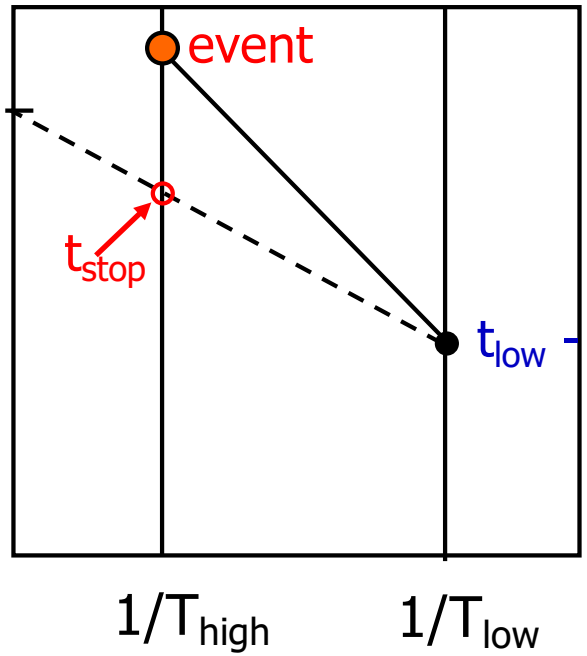

t_{low} - when this event would have occurred at low temperature

TAD procedure

● = current t_{high}

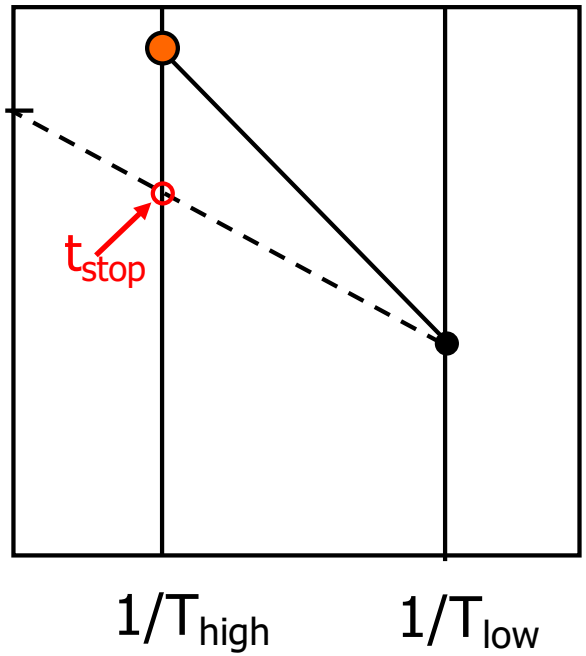

TAD procedure

● = current t_{high}

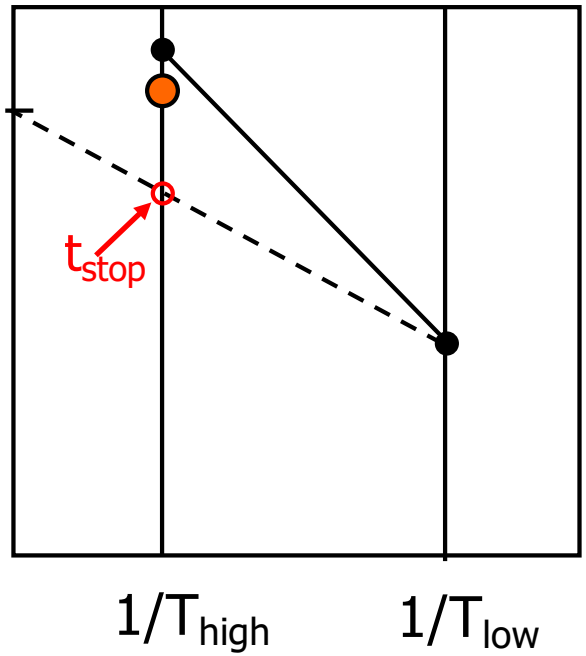

TAD procedure

● = current t_{high}

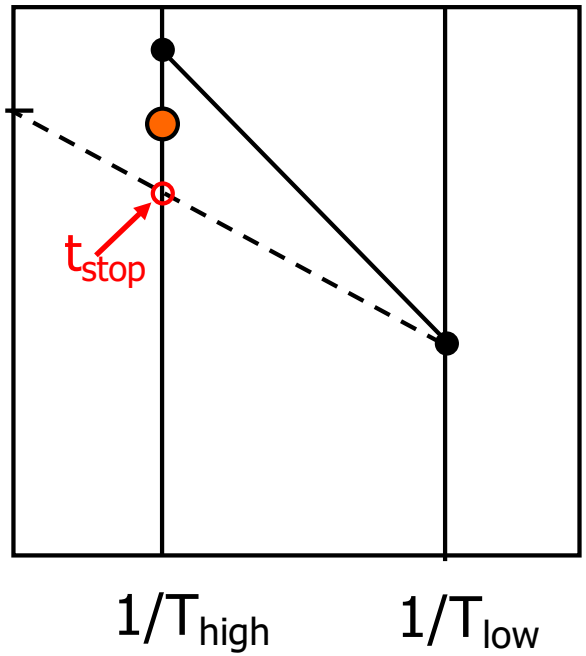

TAD procedure

● = current t_{high}

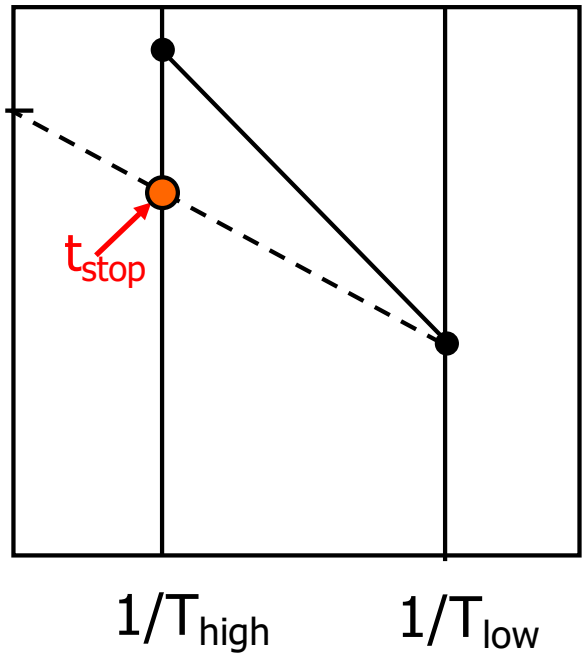

TAD procedure

● = current t_{high}

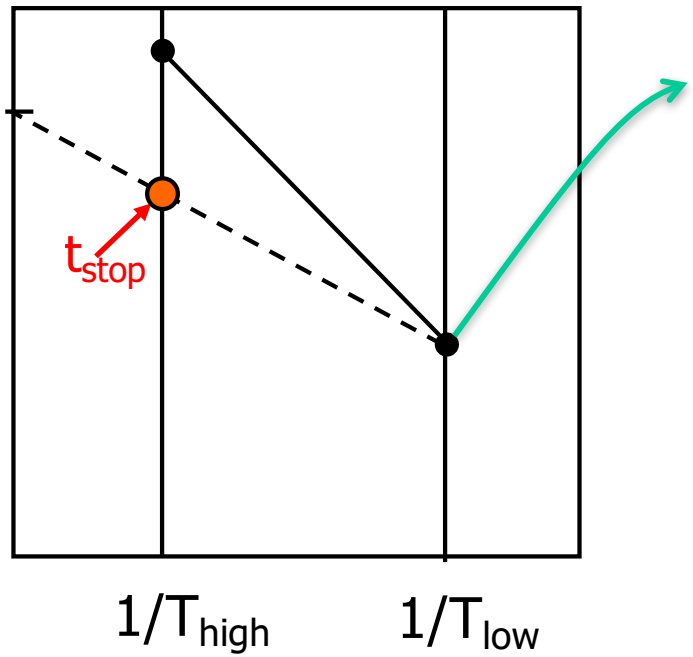

TAD procedure

● = current t_{high}

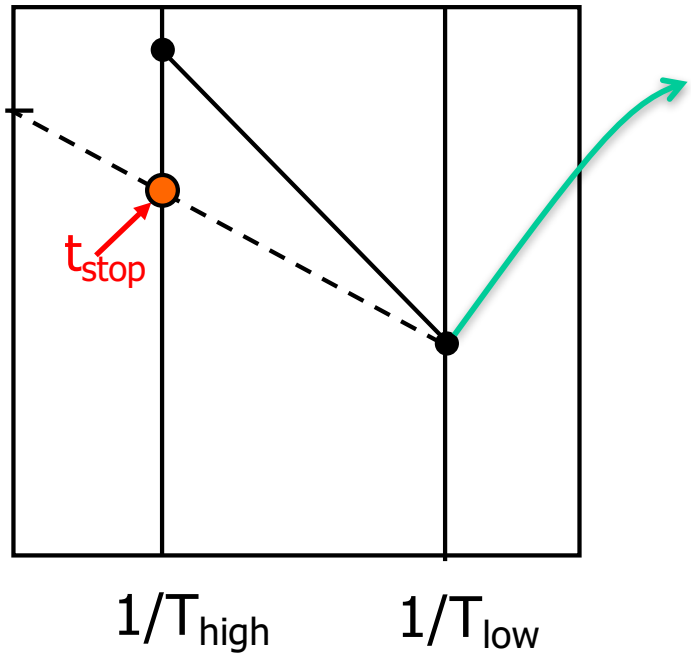

- advance time by t_{low}
- start TAD again in new state

Speculative-transition TAD (SpecTAD)

Speculative-transition TAD (SpecTAD)

A way to parallelize TAD

R.J. Zamora, B.P. Uberuaga, D. Perez, A.F. Voter, Ann. Rev. Chem. Biomol. Eng. (2016);
Zamora, Voter, Perez, Perriot, Uberuaga, Phys. Chem. Chem. Phys. 2016.

Speculative-transition TAD (SpecTAD)

A way to parallelize TAD

- TAD sub-simulation is spawned as soon as each transition is seen.
- Branch continues until/unless it becomes clear that this transition is not the one that will be accepted.

R.J. Zamora, B.P. Uberuaga, D. Perez, A.F. Voter, Ann. Rev. Chem. Biomol. Eng. (2016);
Zamora, Voter, Perez, Perriot, Uberuaga, Phys. Chem. Chem. Phys. 2016.

Speculative Transition Acceptance

● = current t_{high}

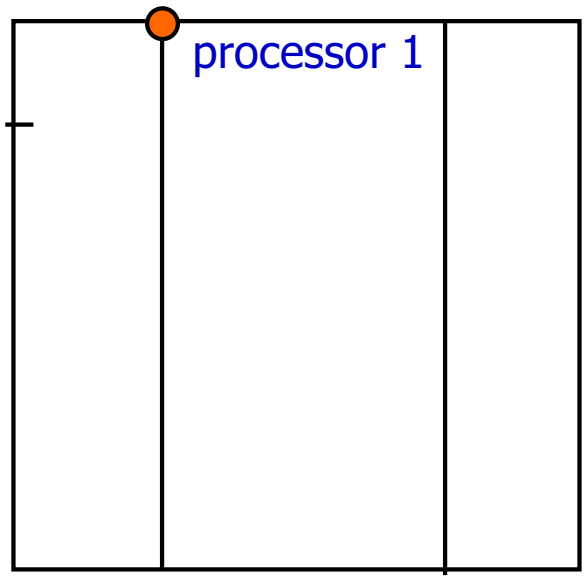

Speculative Transition Acceptance

● = current t_{high}

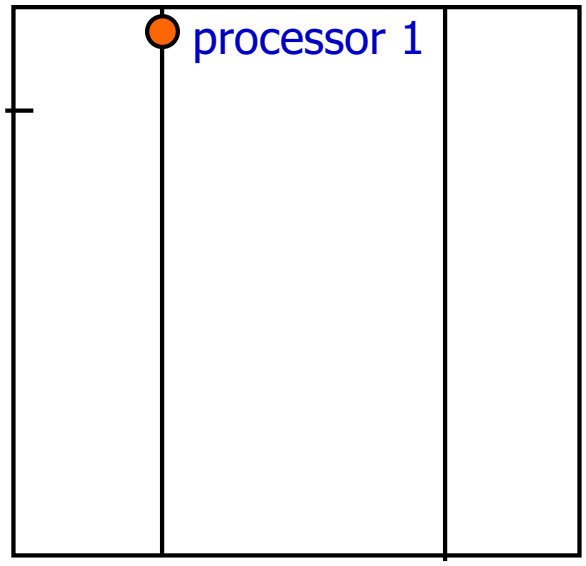

Speculative Transition Acceptance

● = current t_{high}

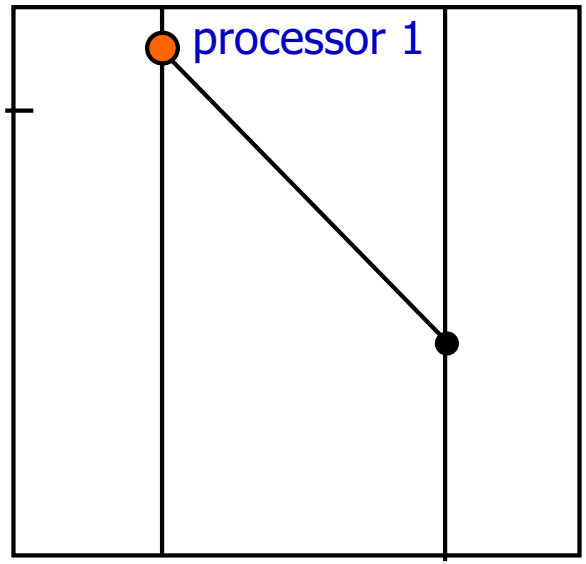

Speculative Transition Acceptance

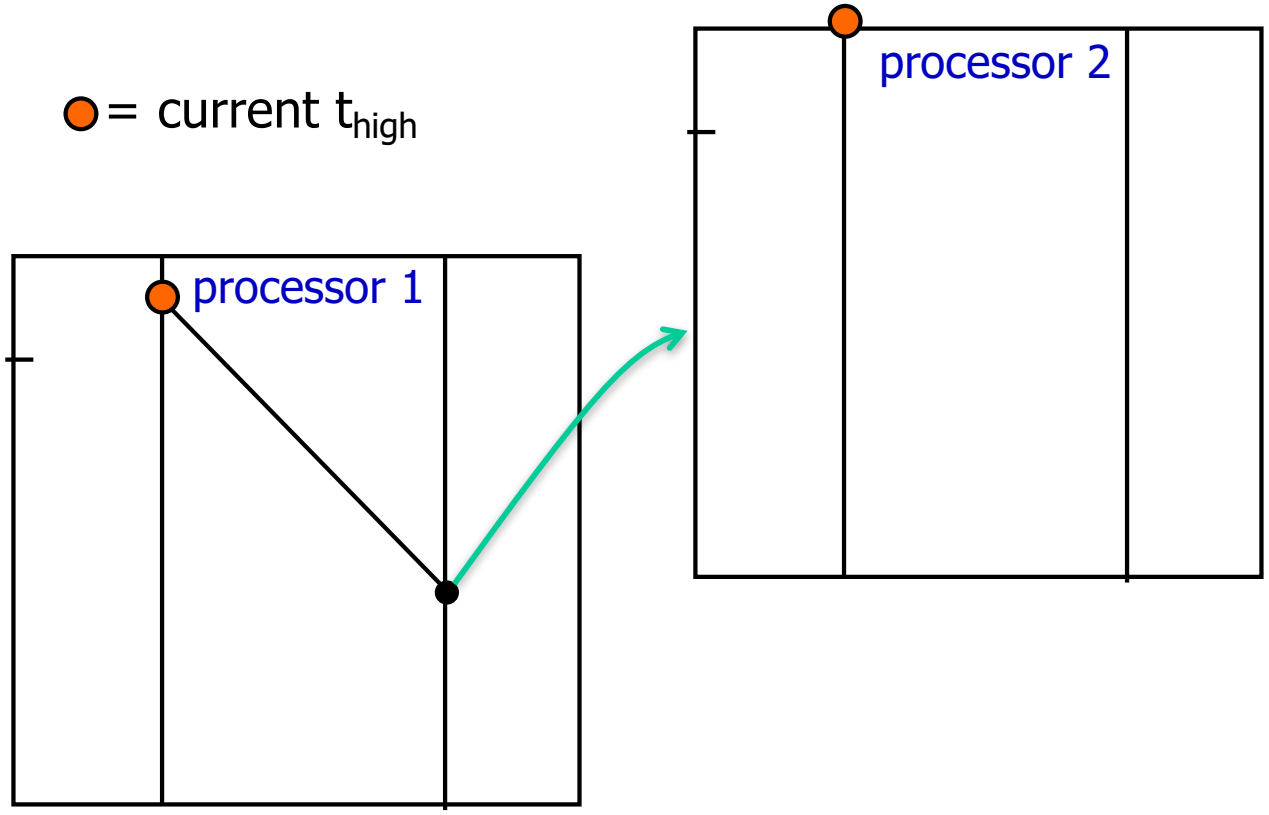

Speculative Transition Acceptance

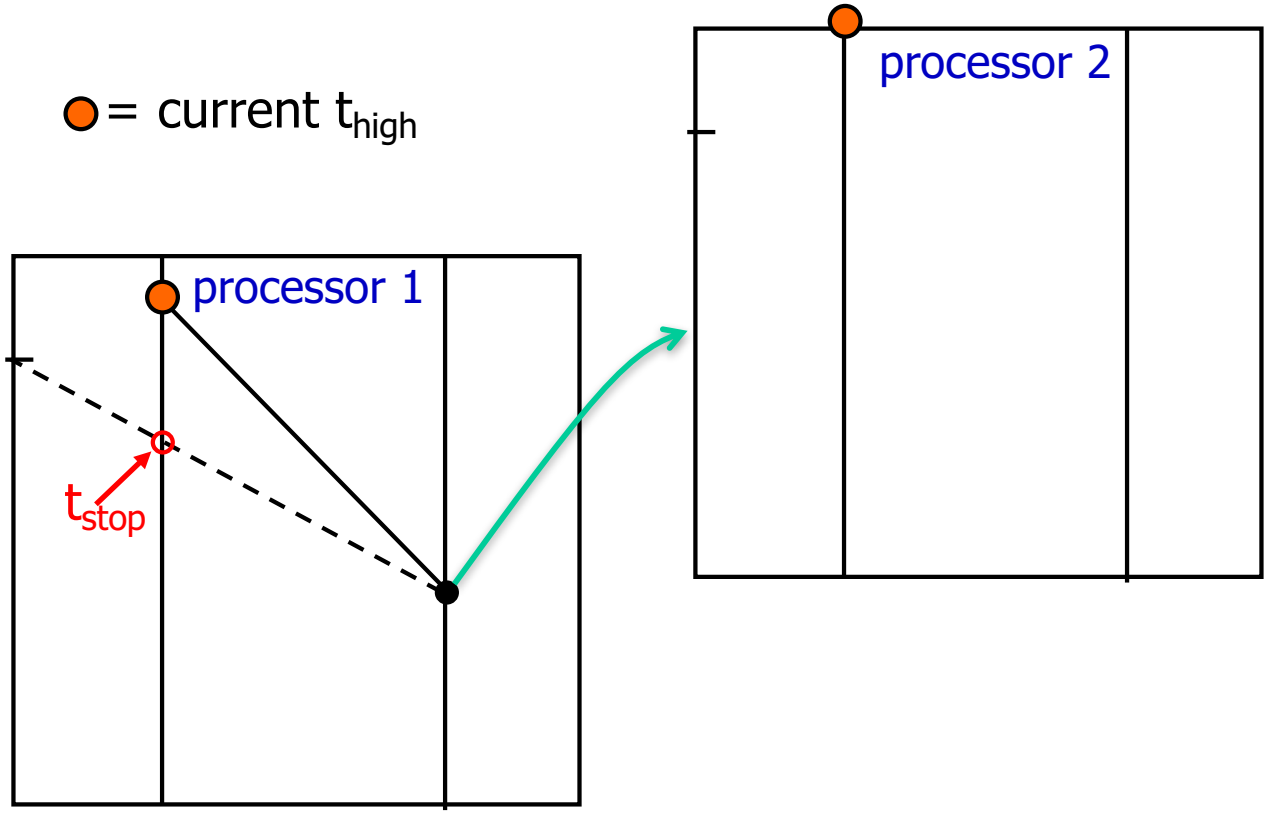

Speculative Transition Acceptance

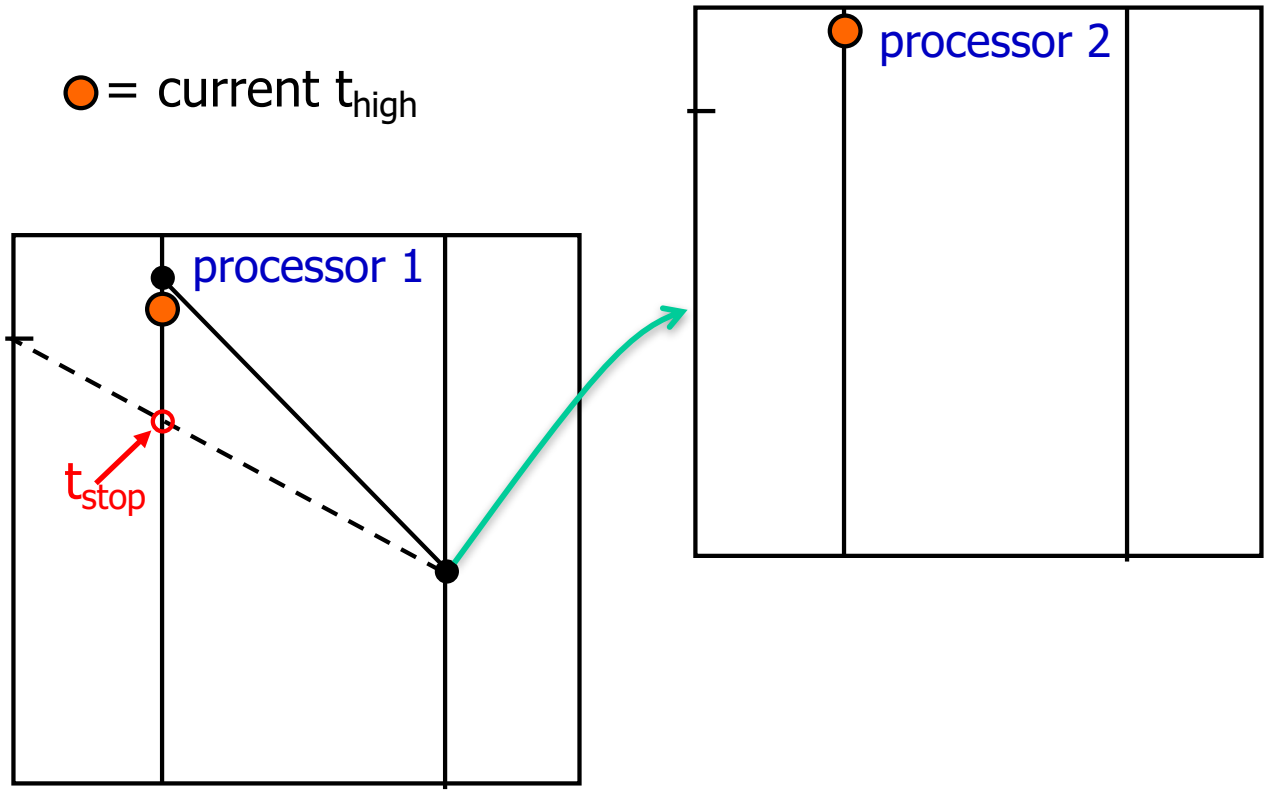

Speculative Transition Acceptance

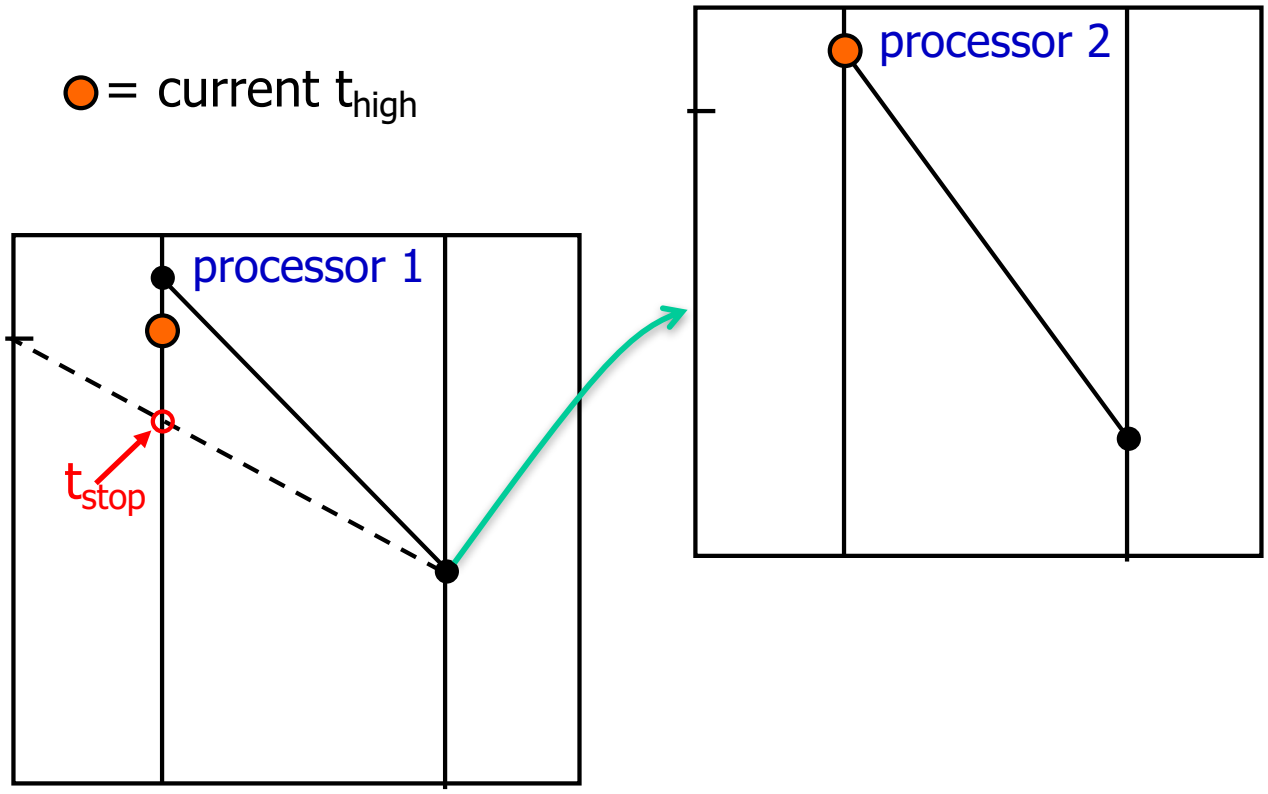

Speculative Transition Acceptance

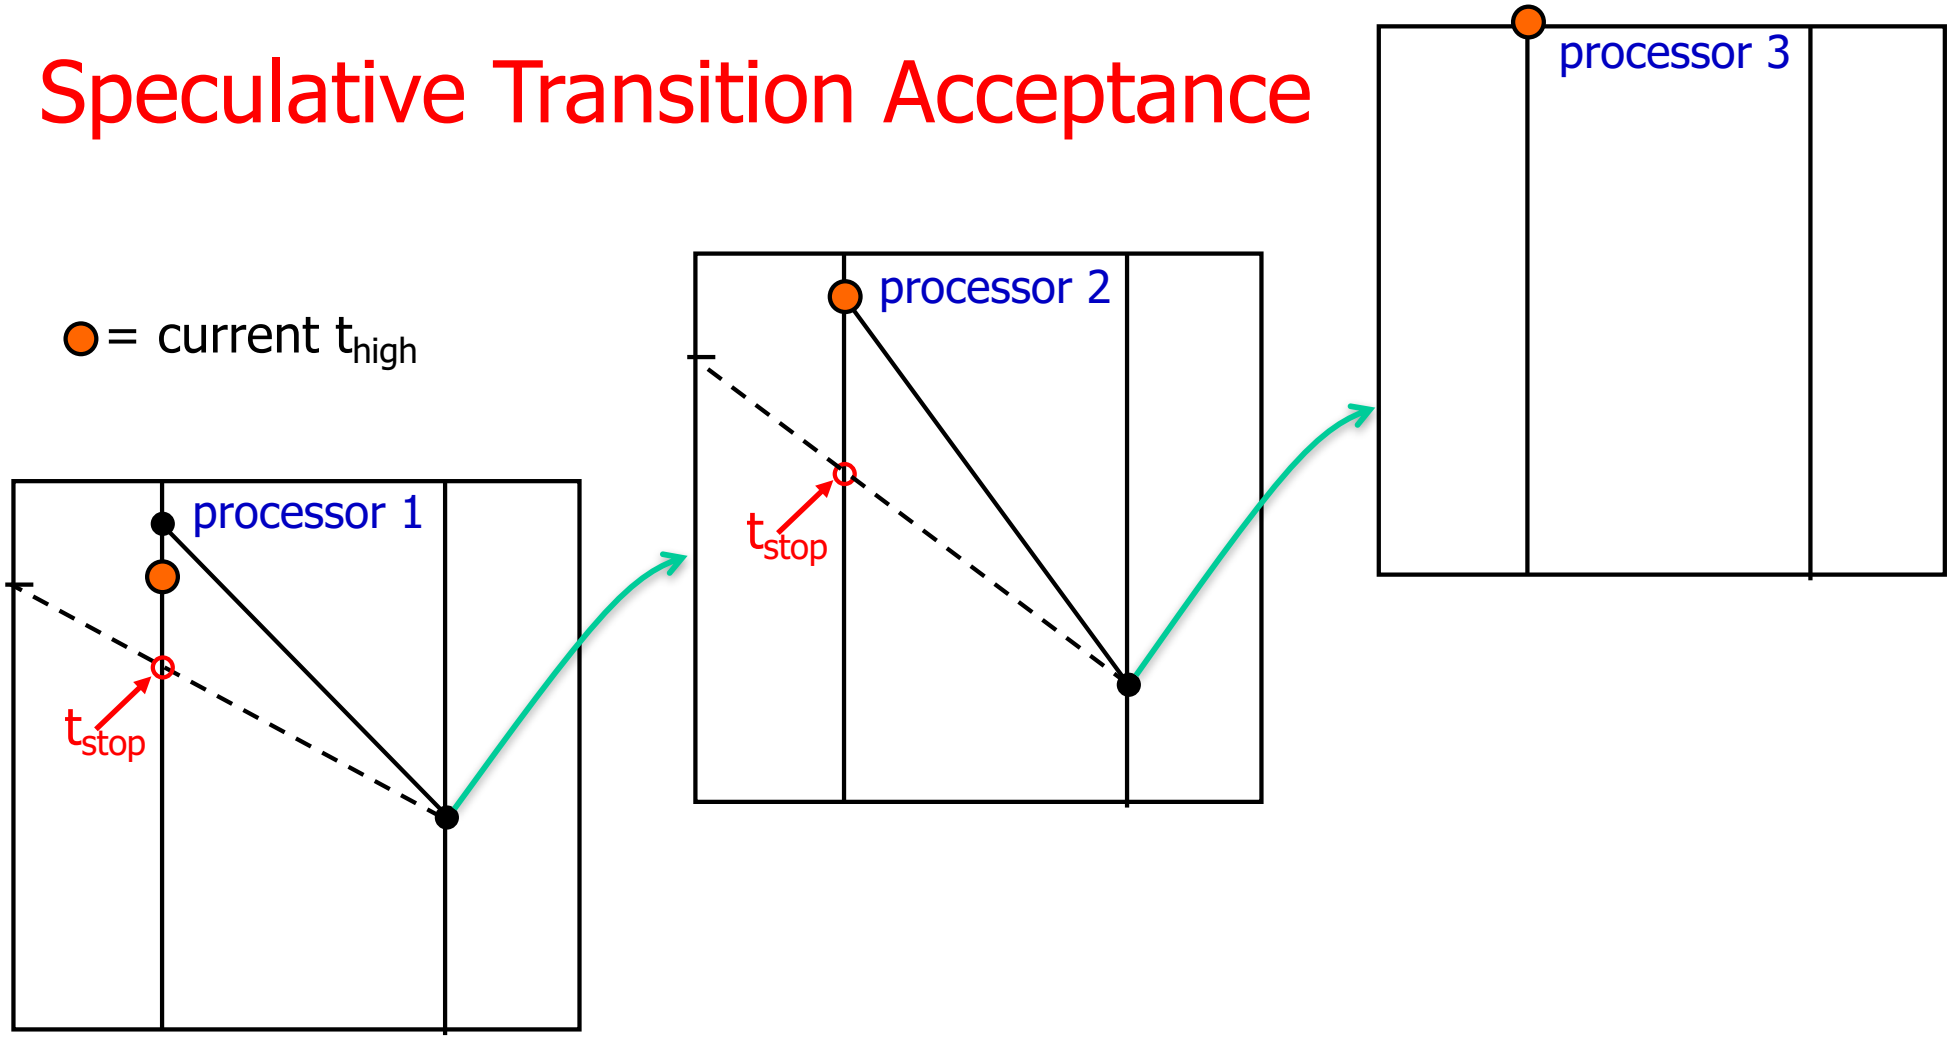

Speculative Transition Acceptance

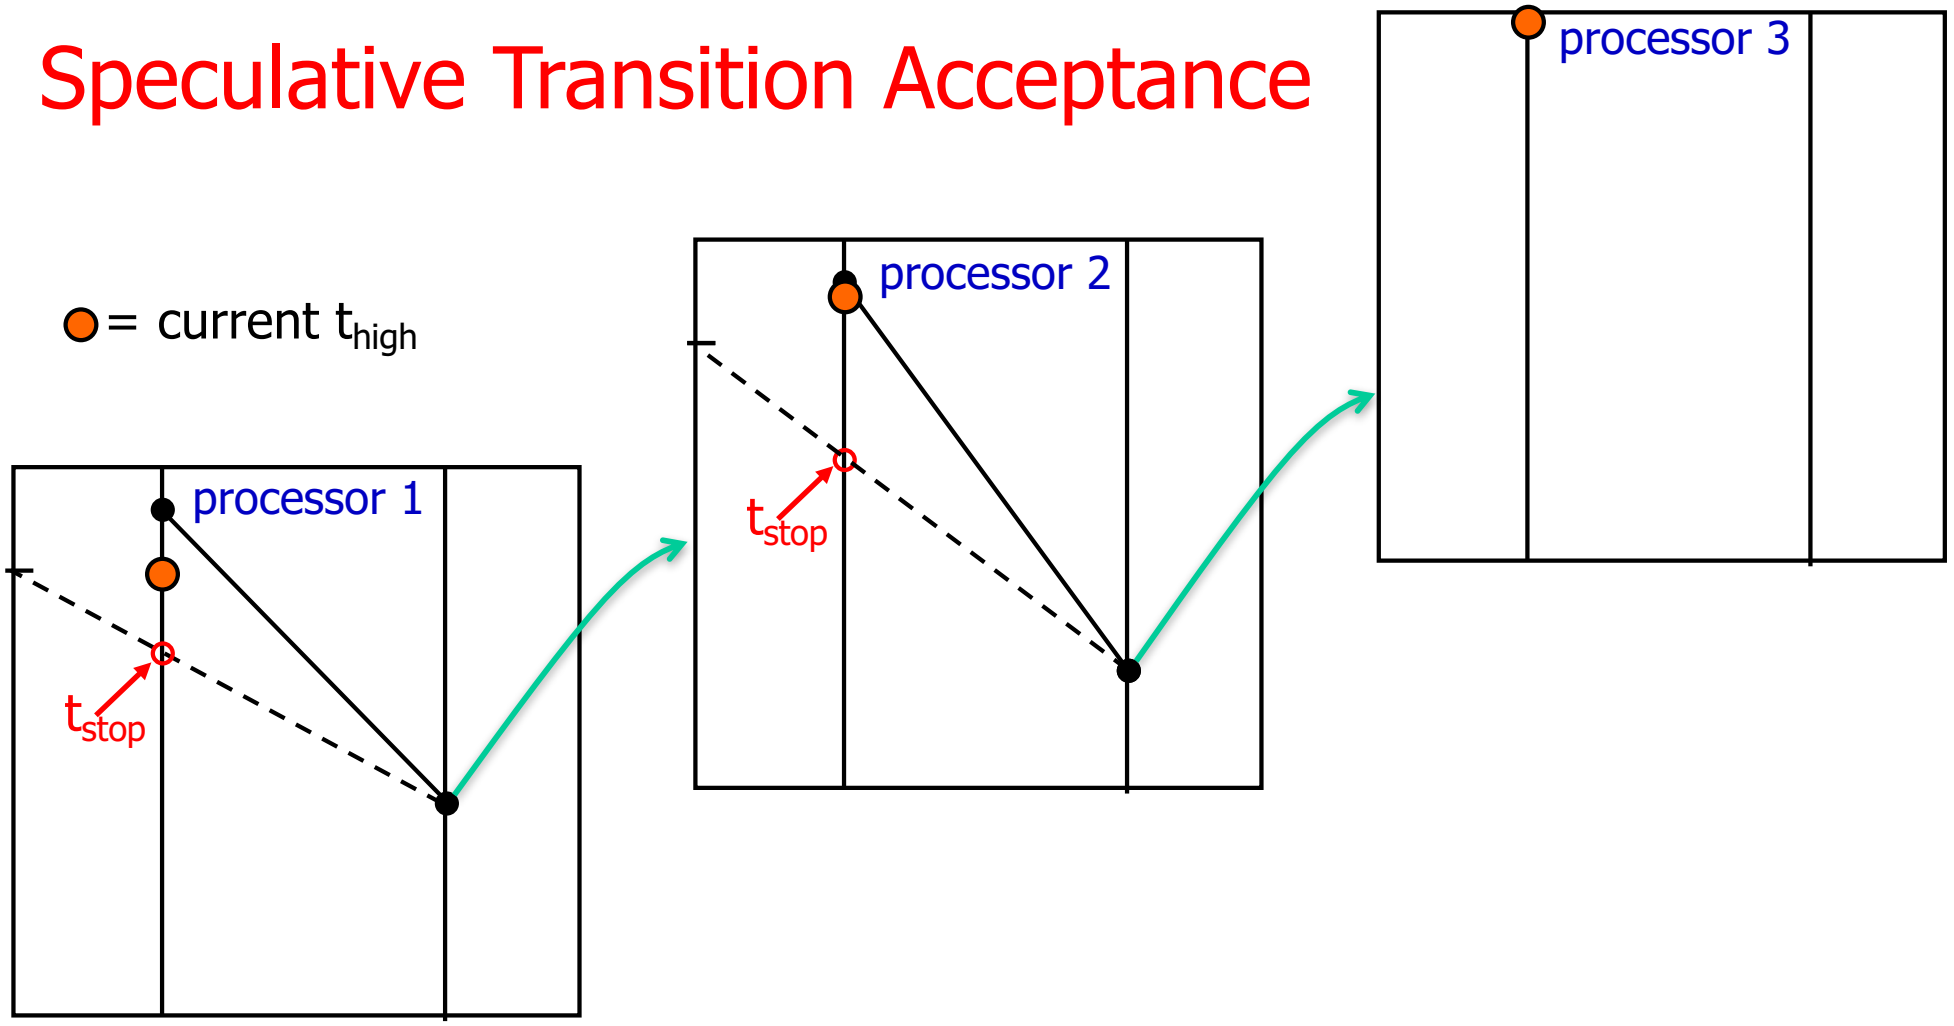

Speculative Transition Acceptance

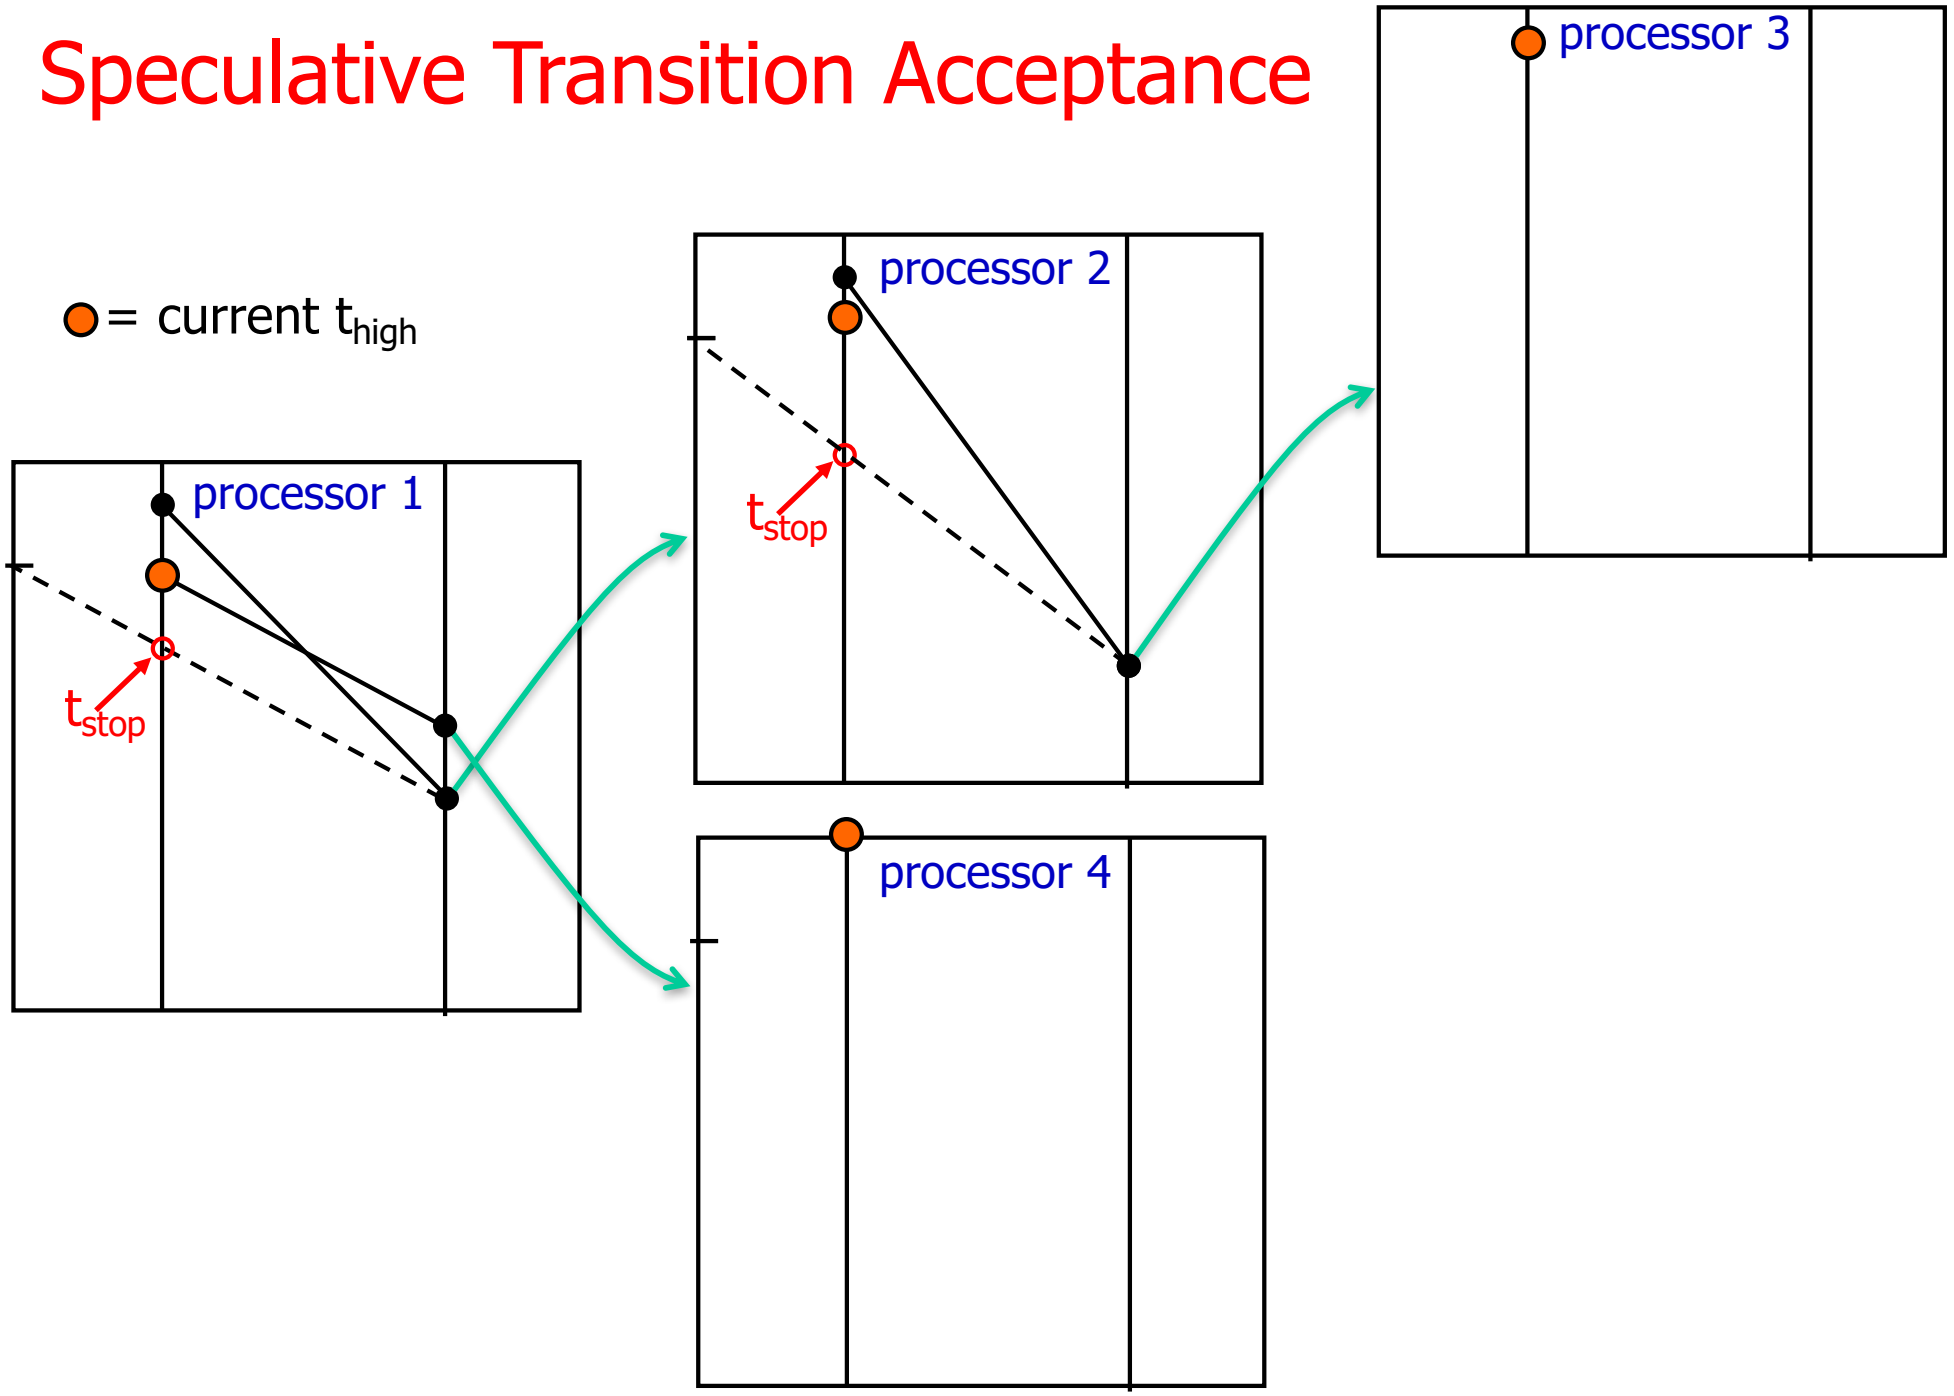

Speculative Transition Acceptance

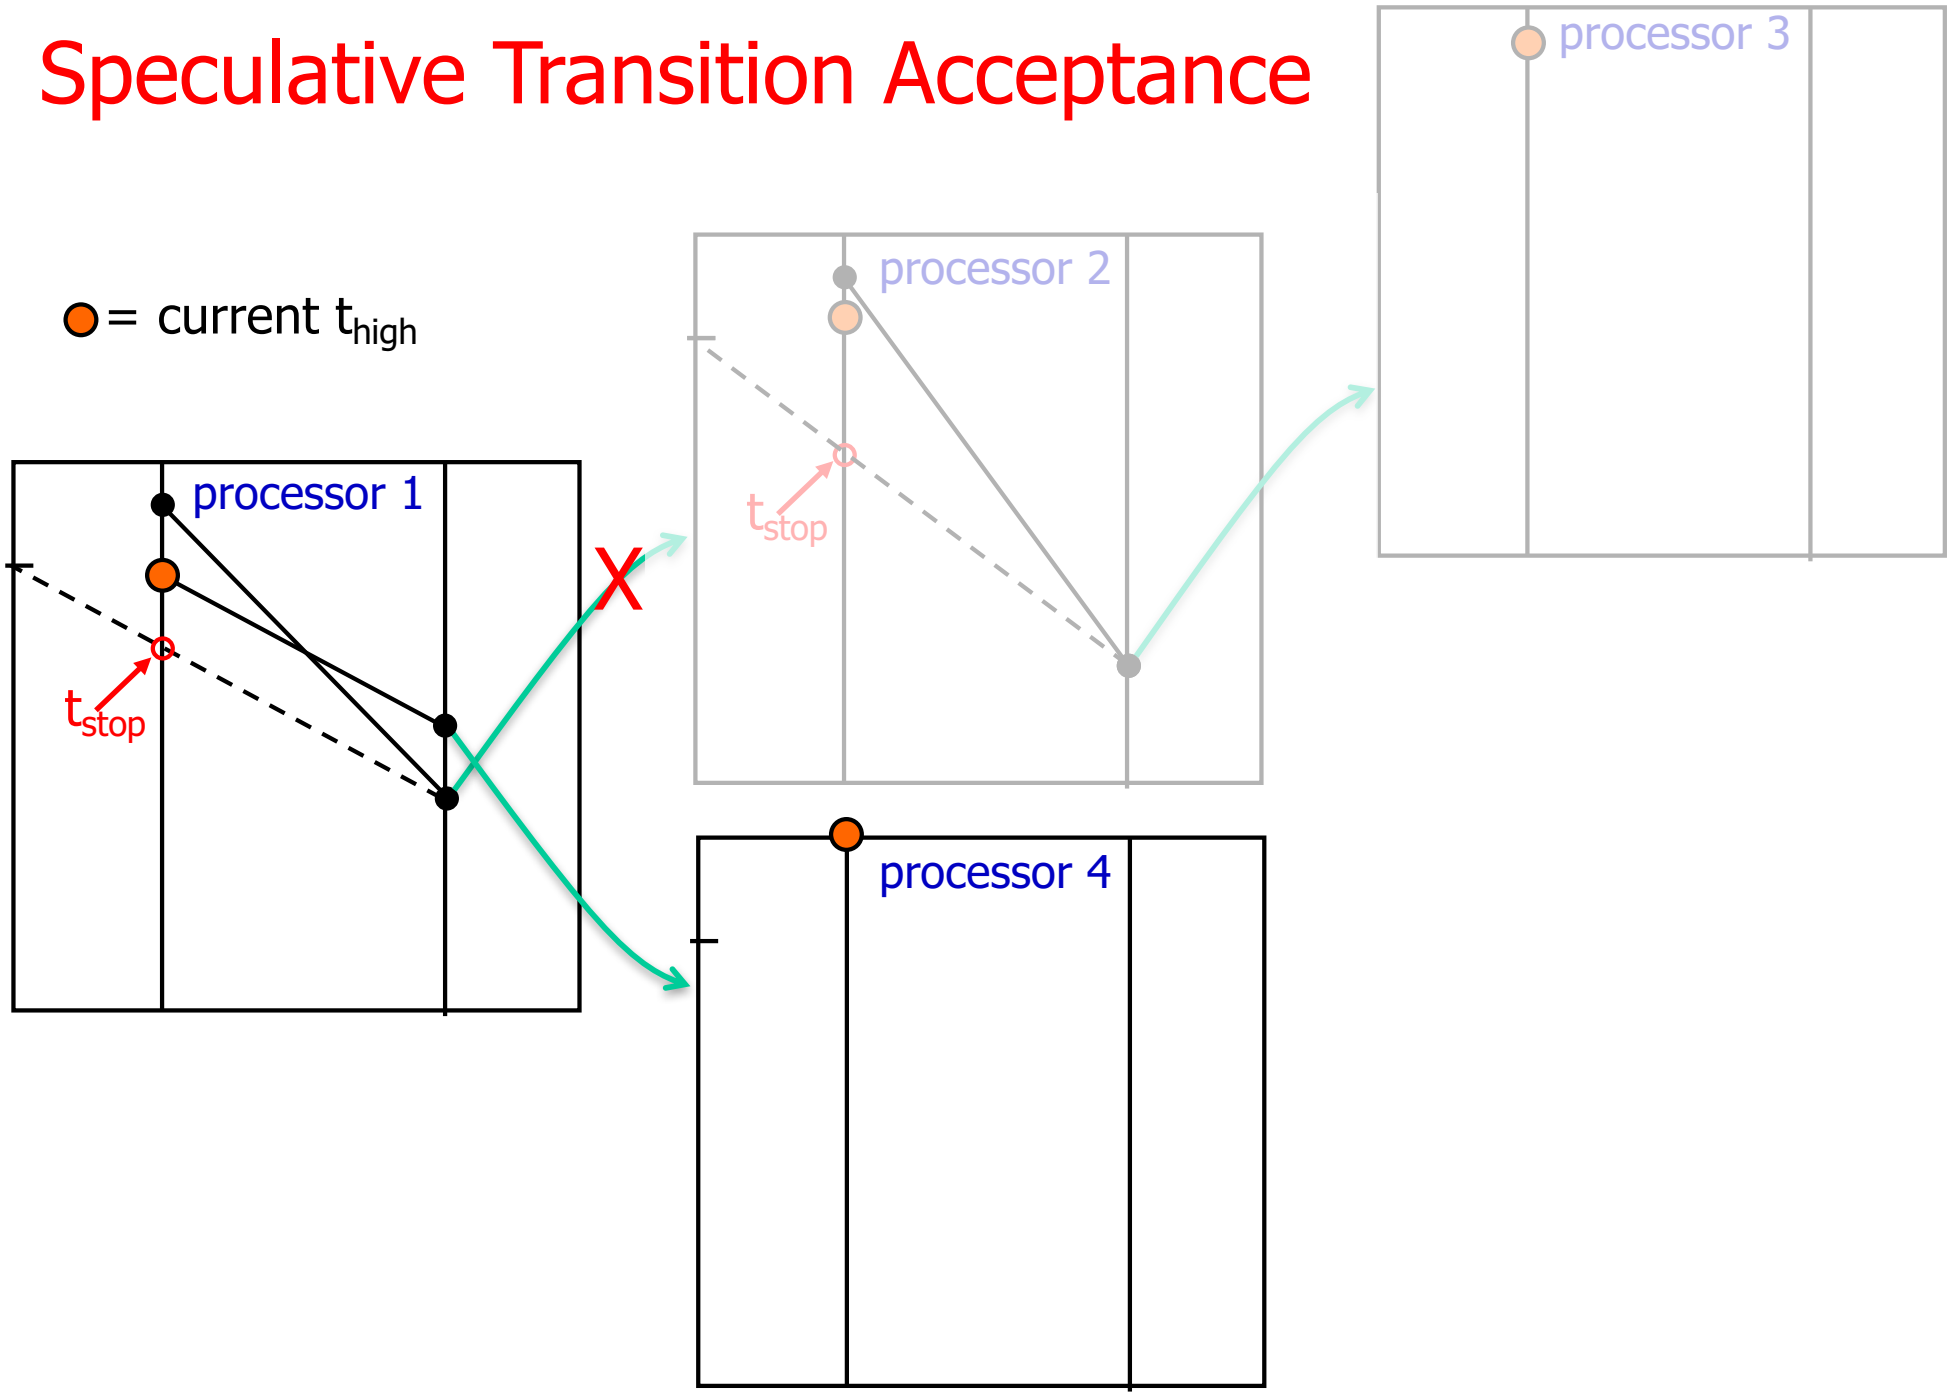

Speculative Transition Acceptance

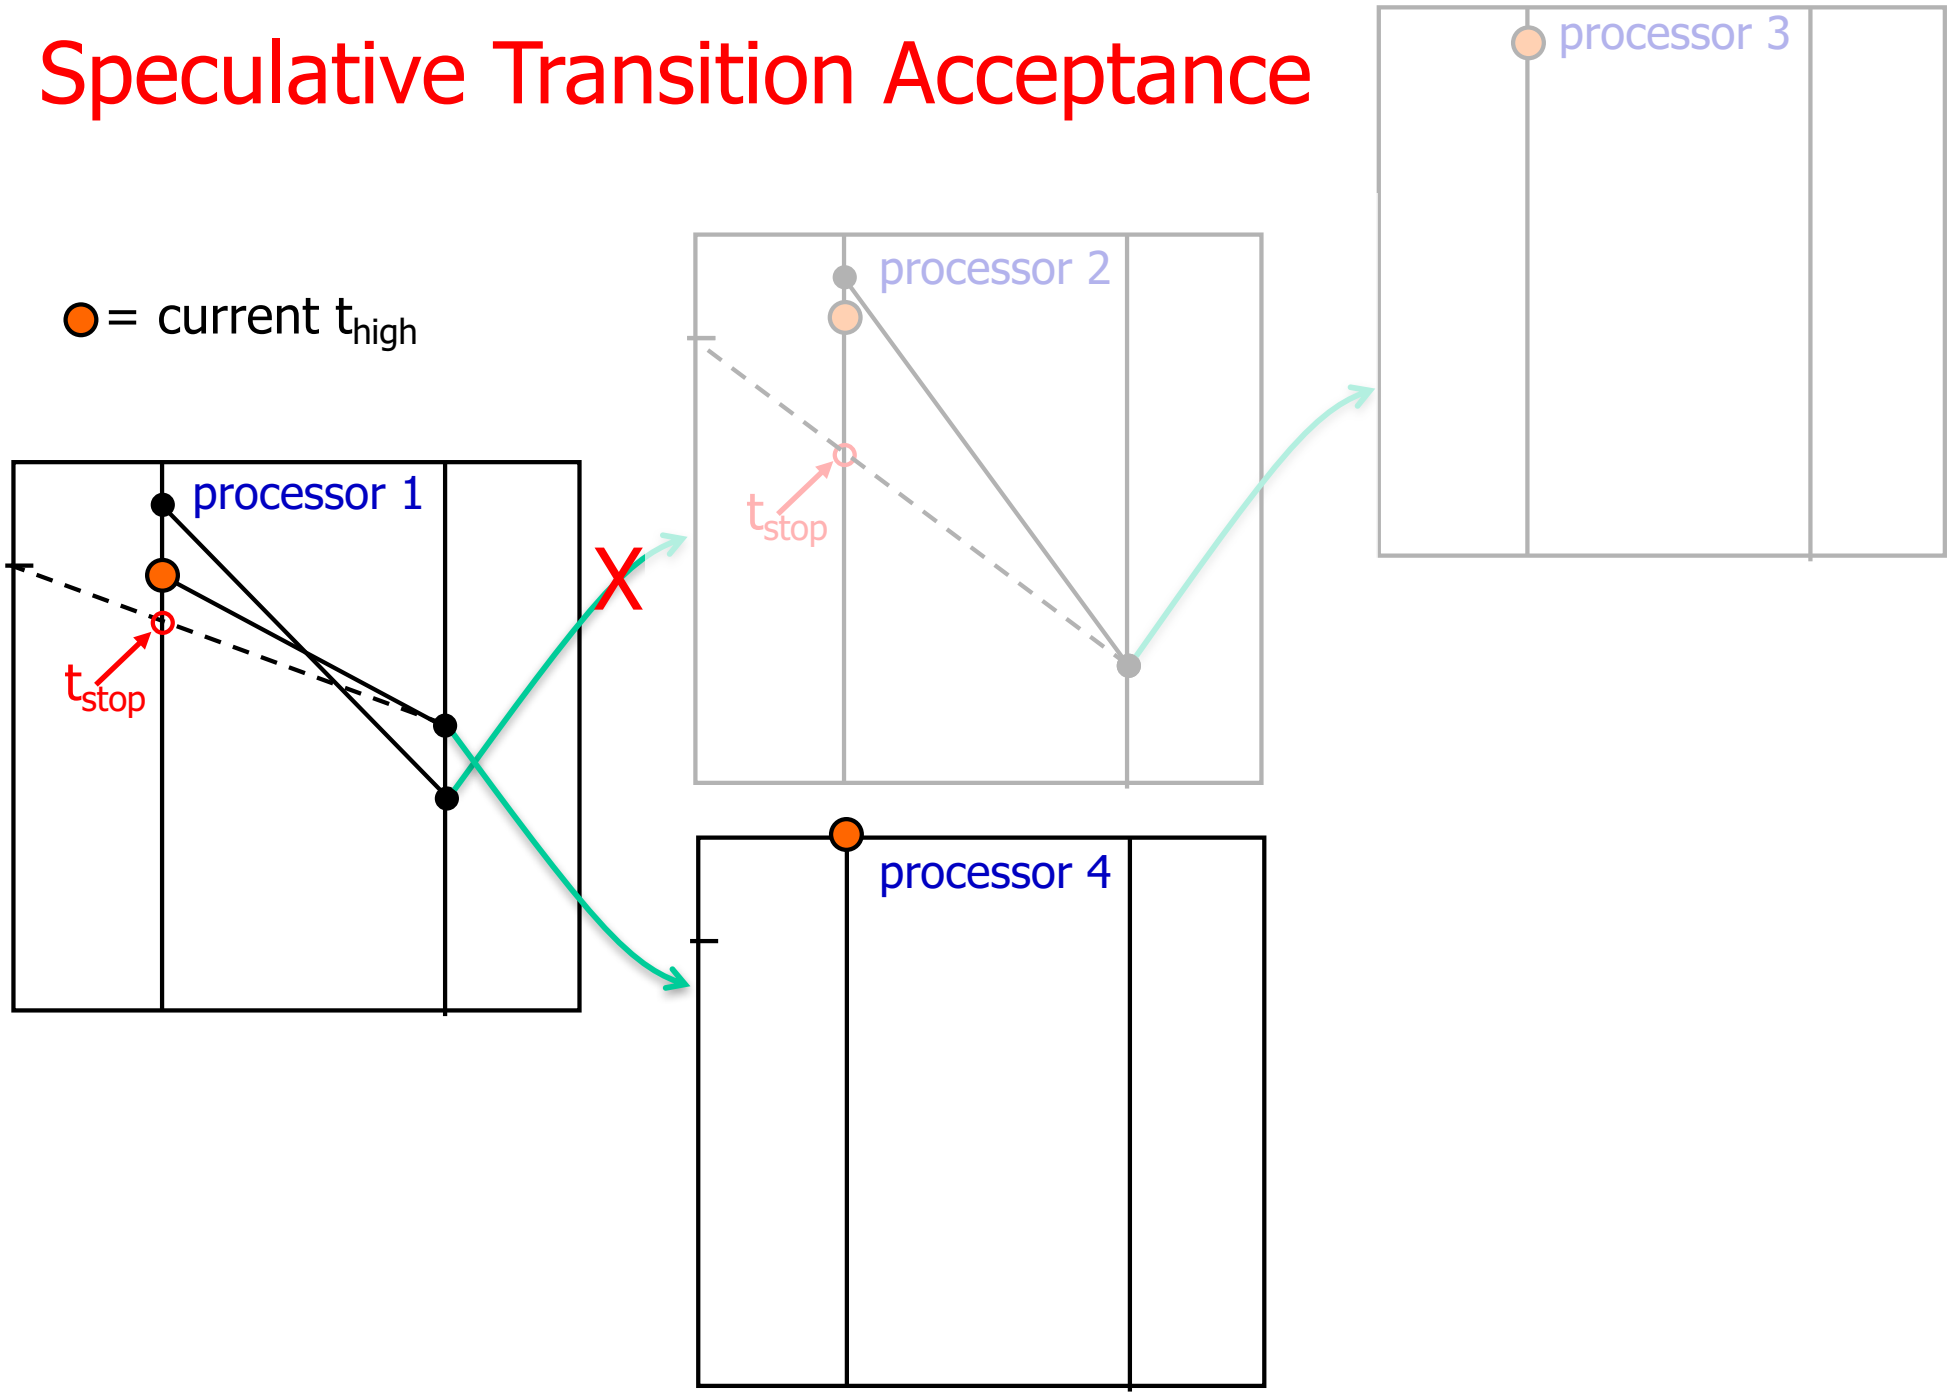

Speculative Transition Acceptance

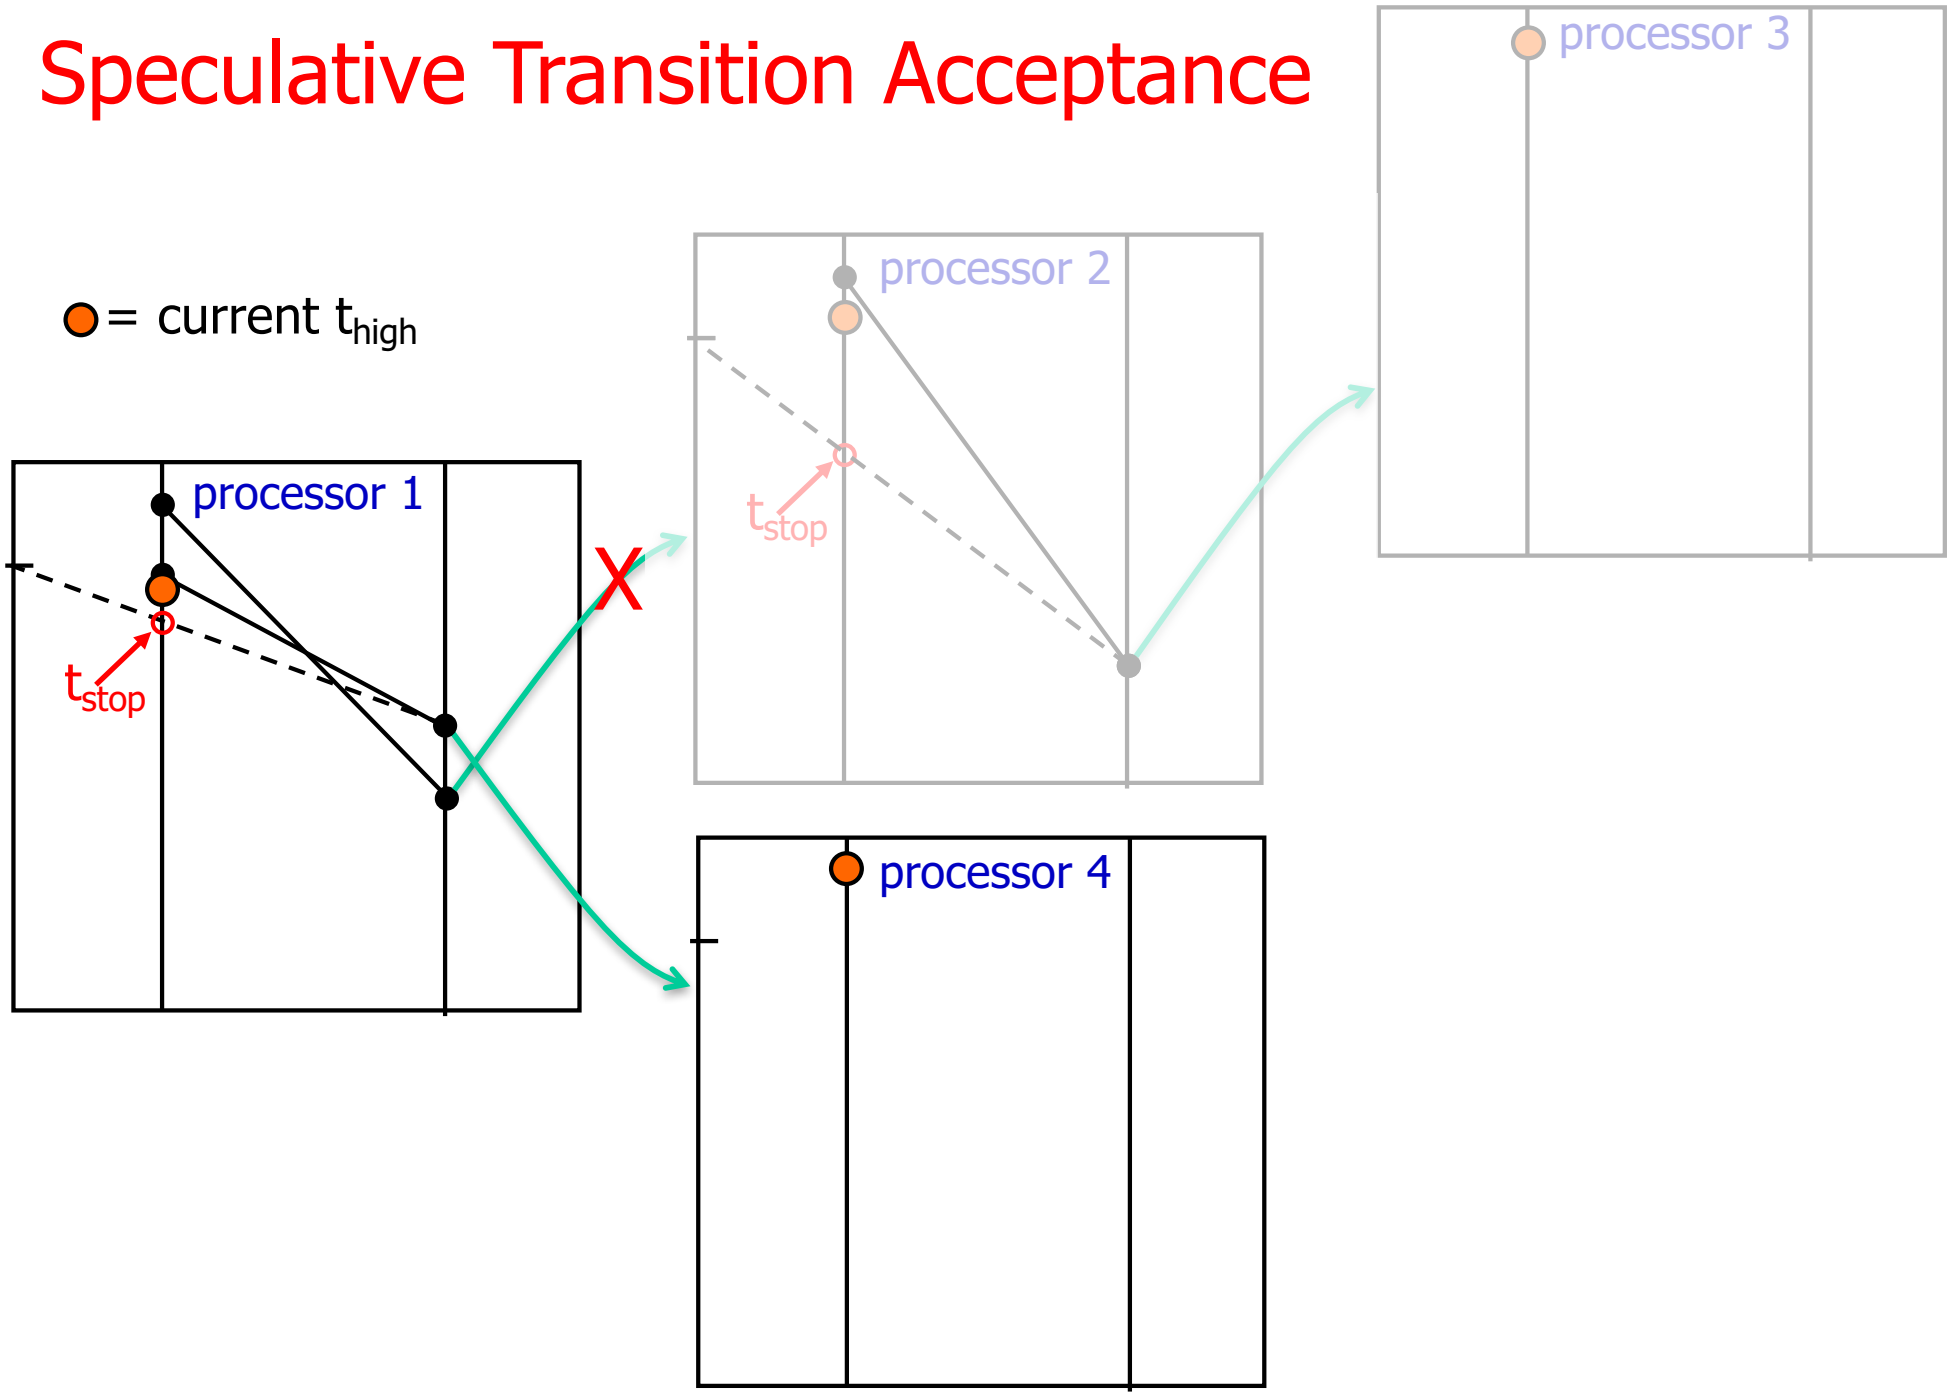

Speculative Transition Acceptance

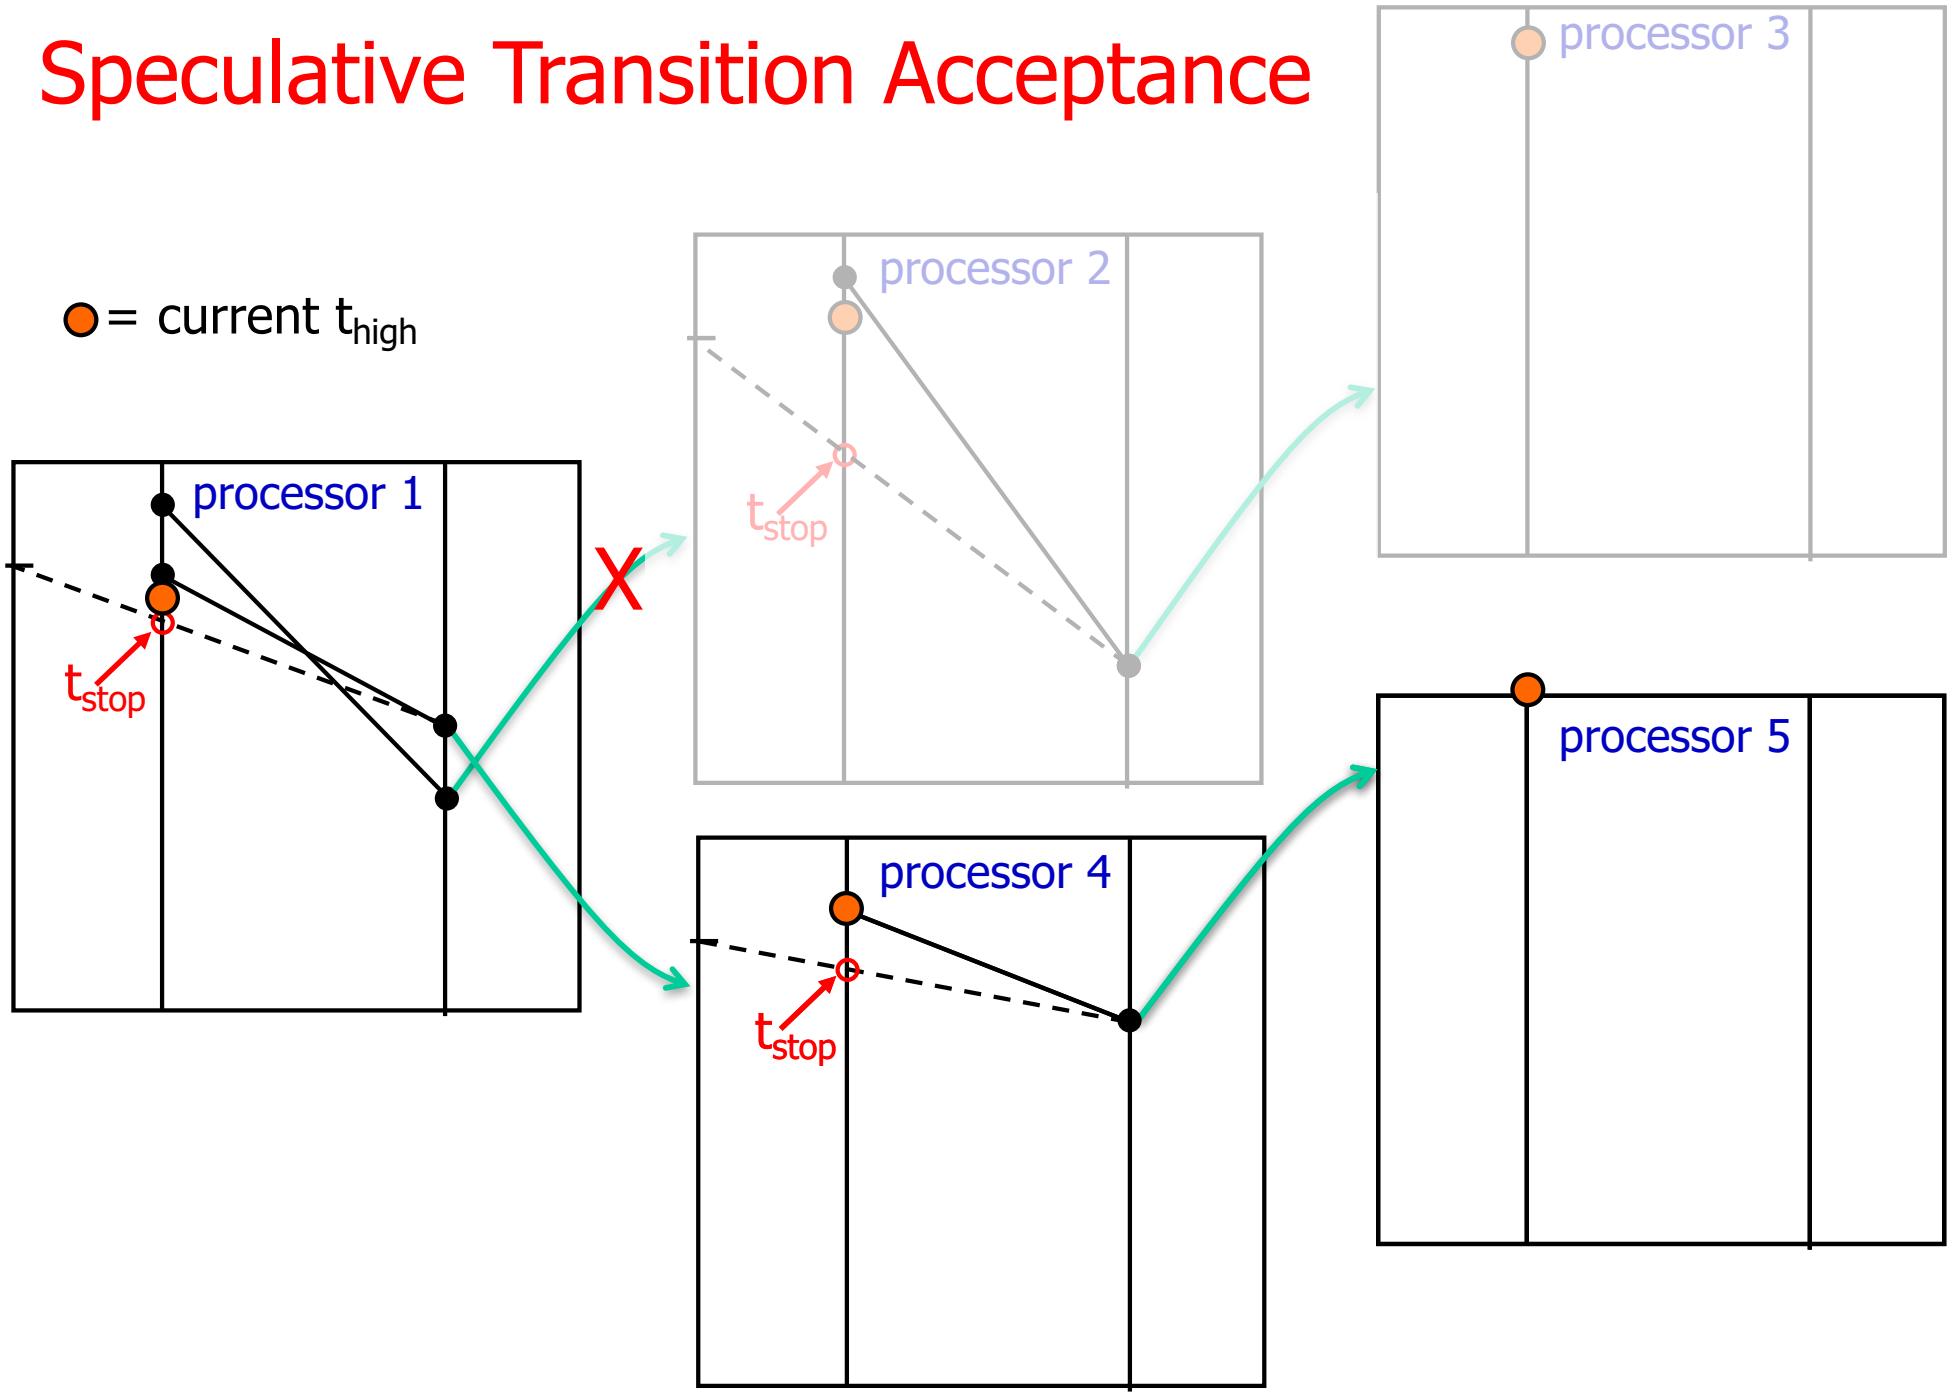

Speculative Transition Acceptance

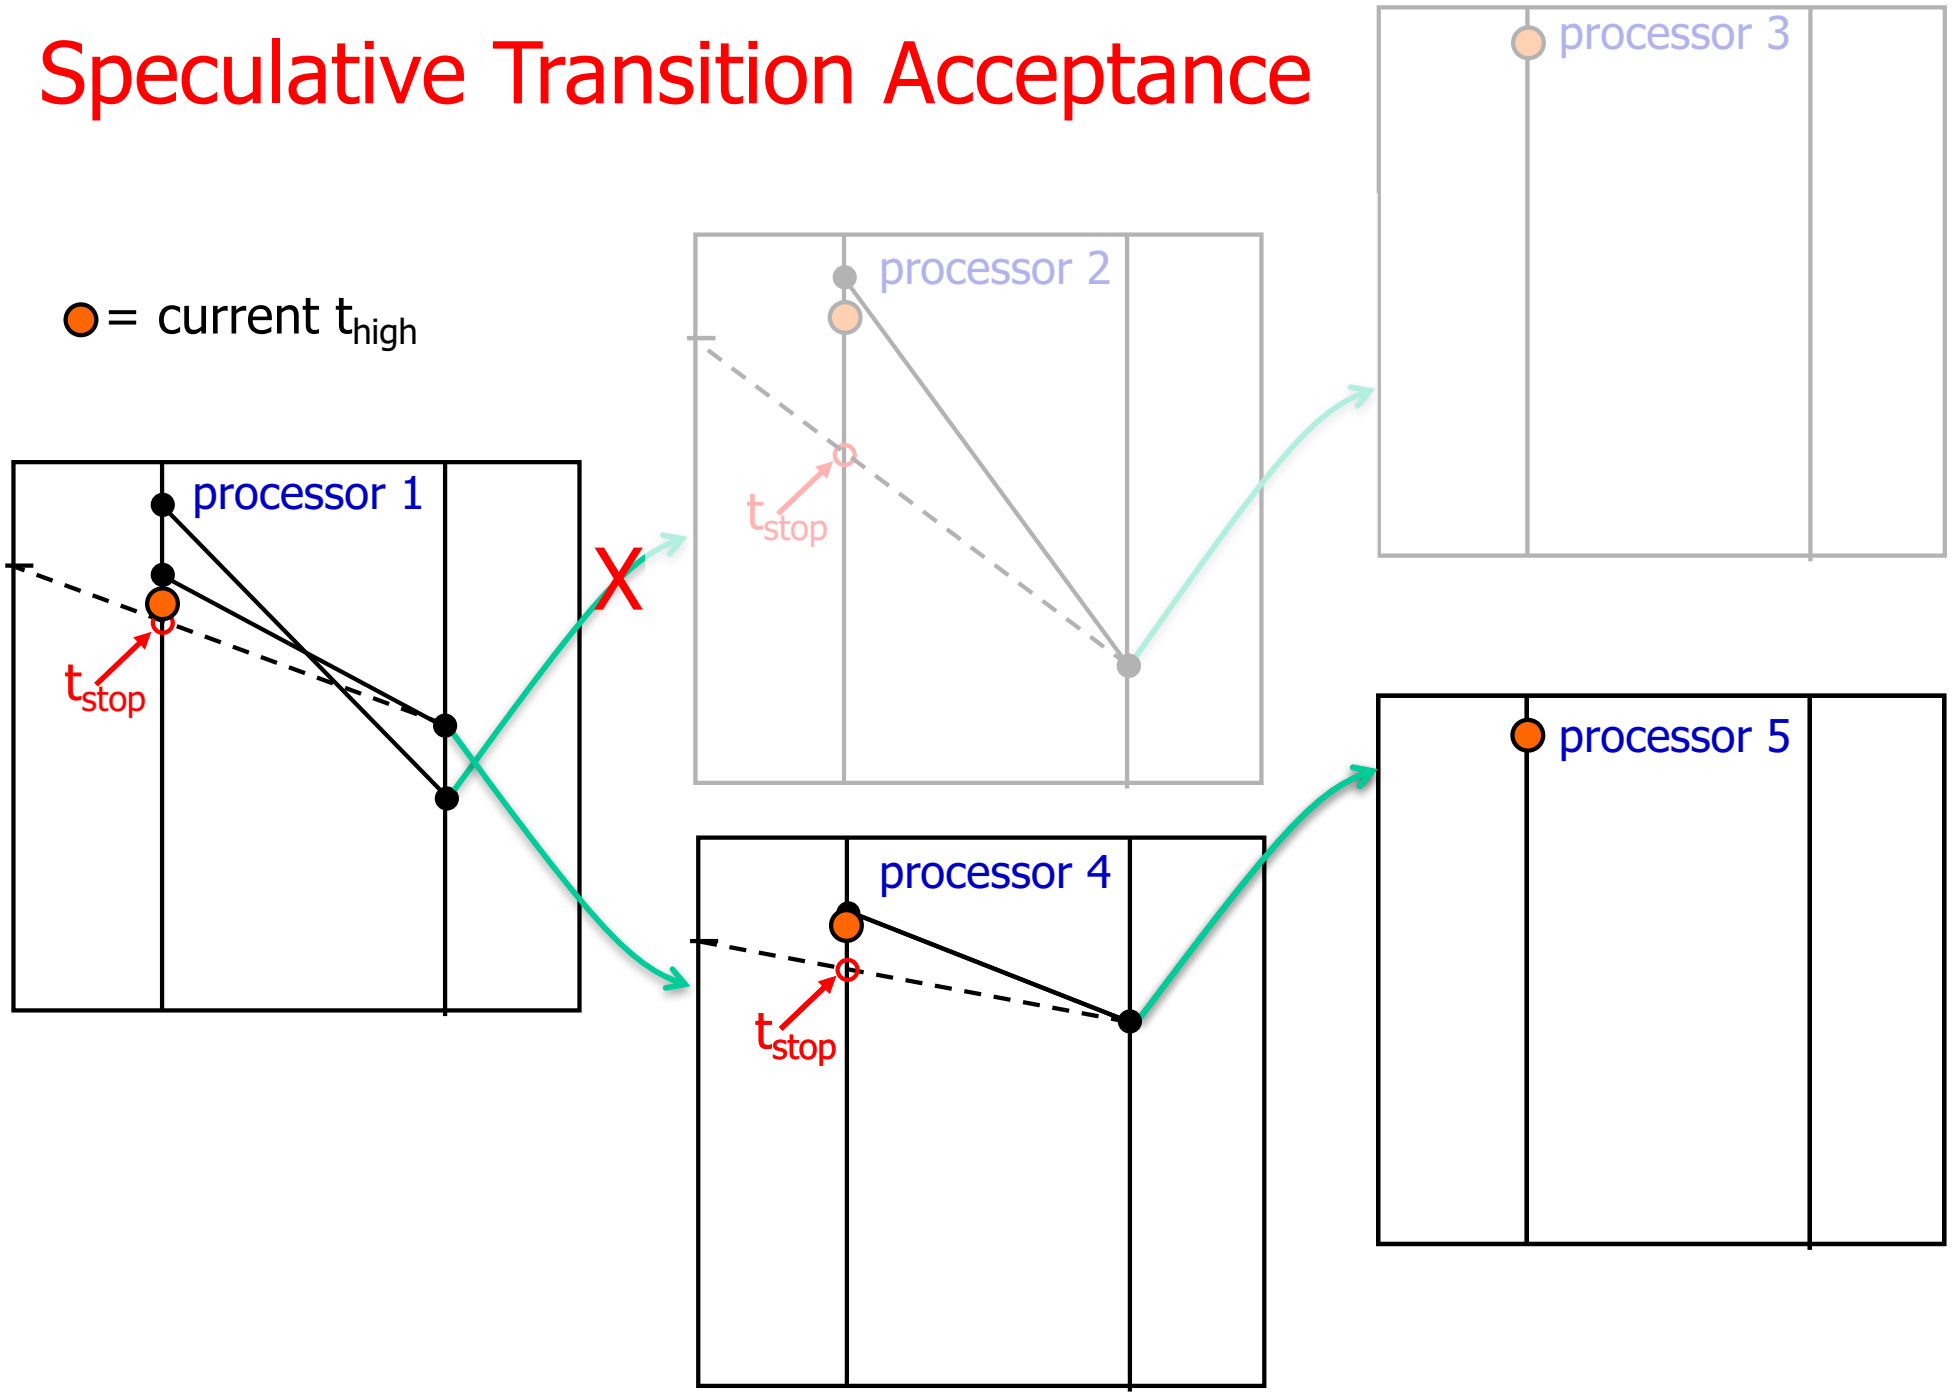

Speculative Transition Acceptance

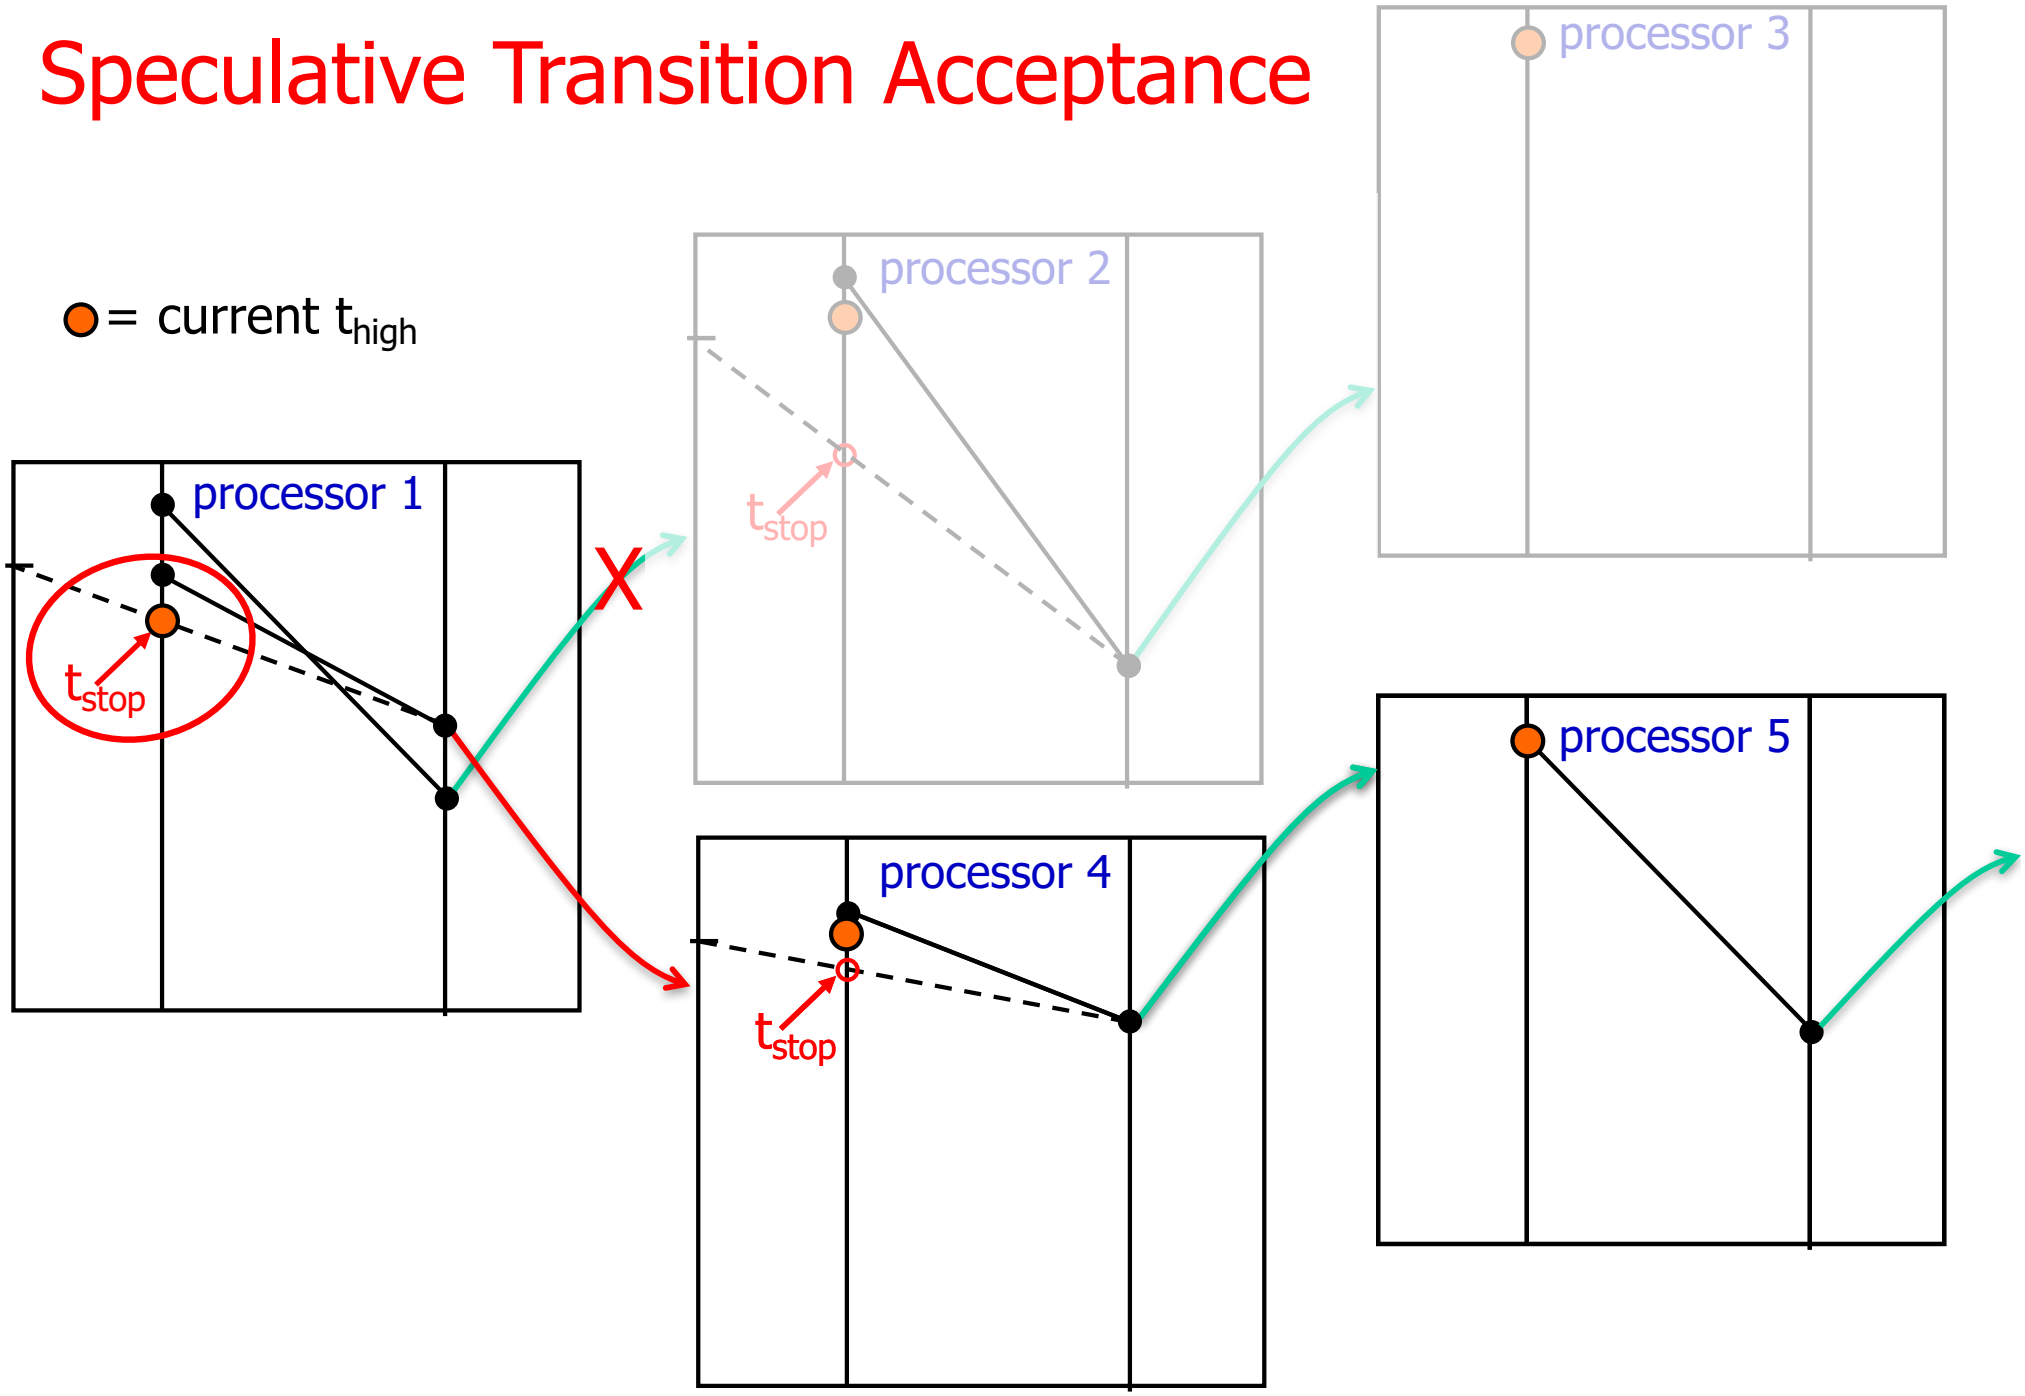

Speculative Transition Acceptance

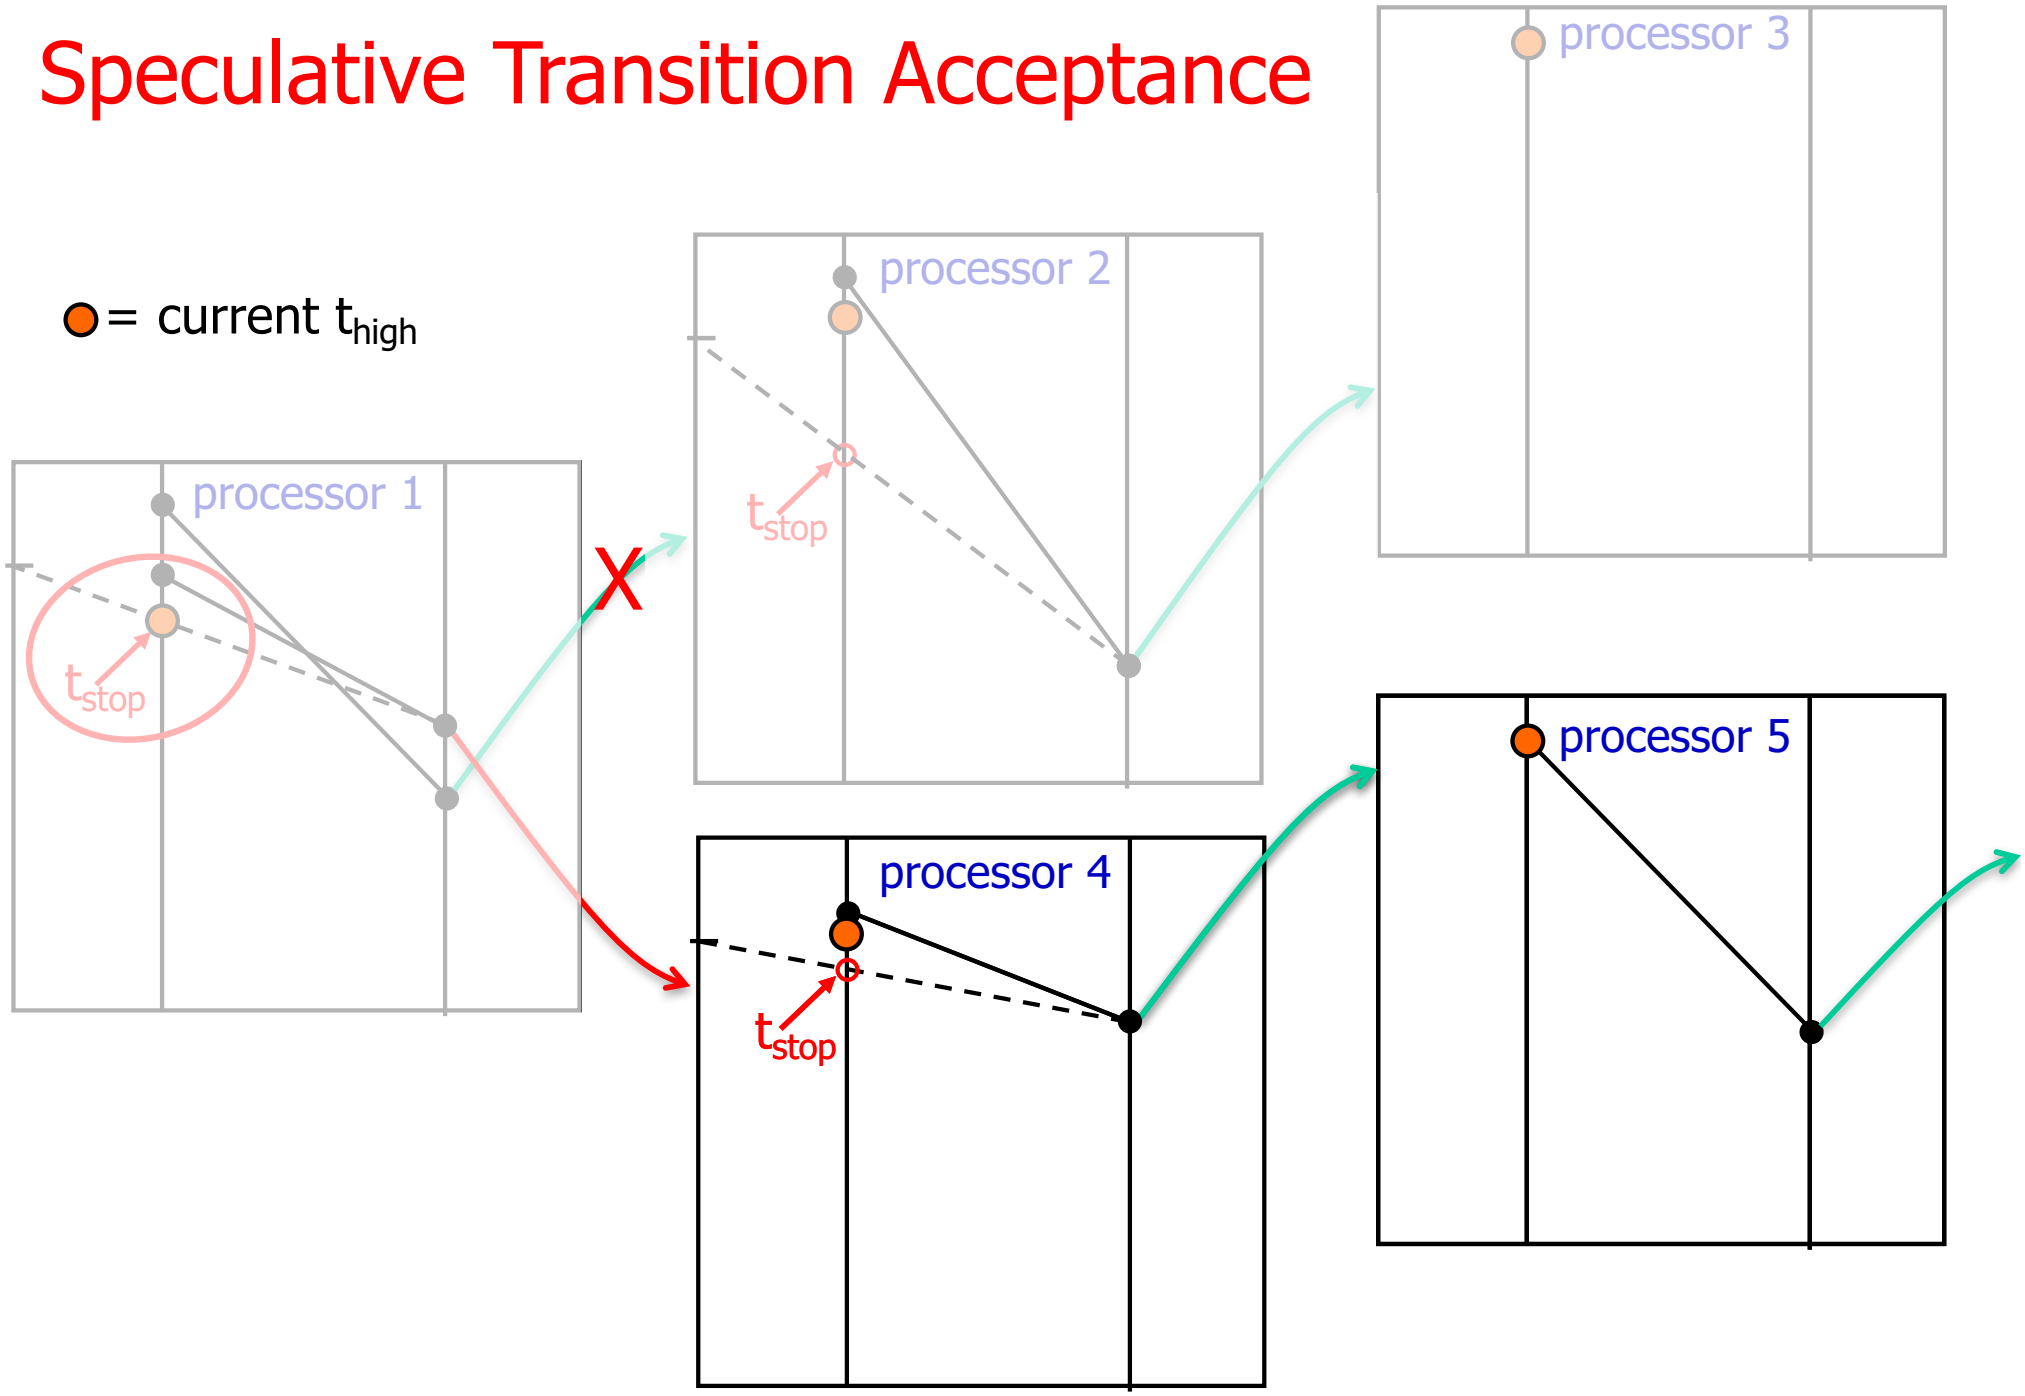

Speculative Transition Acceptance

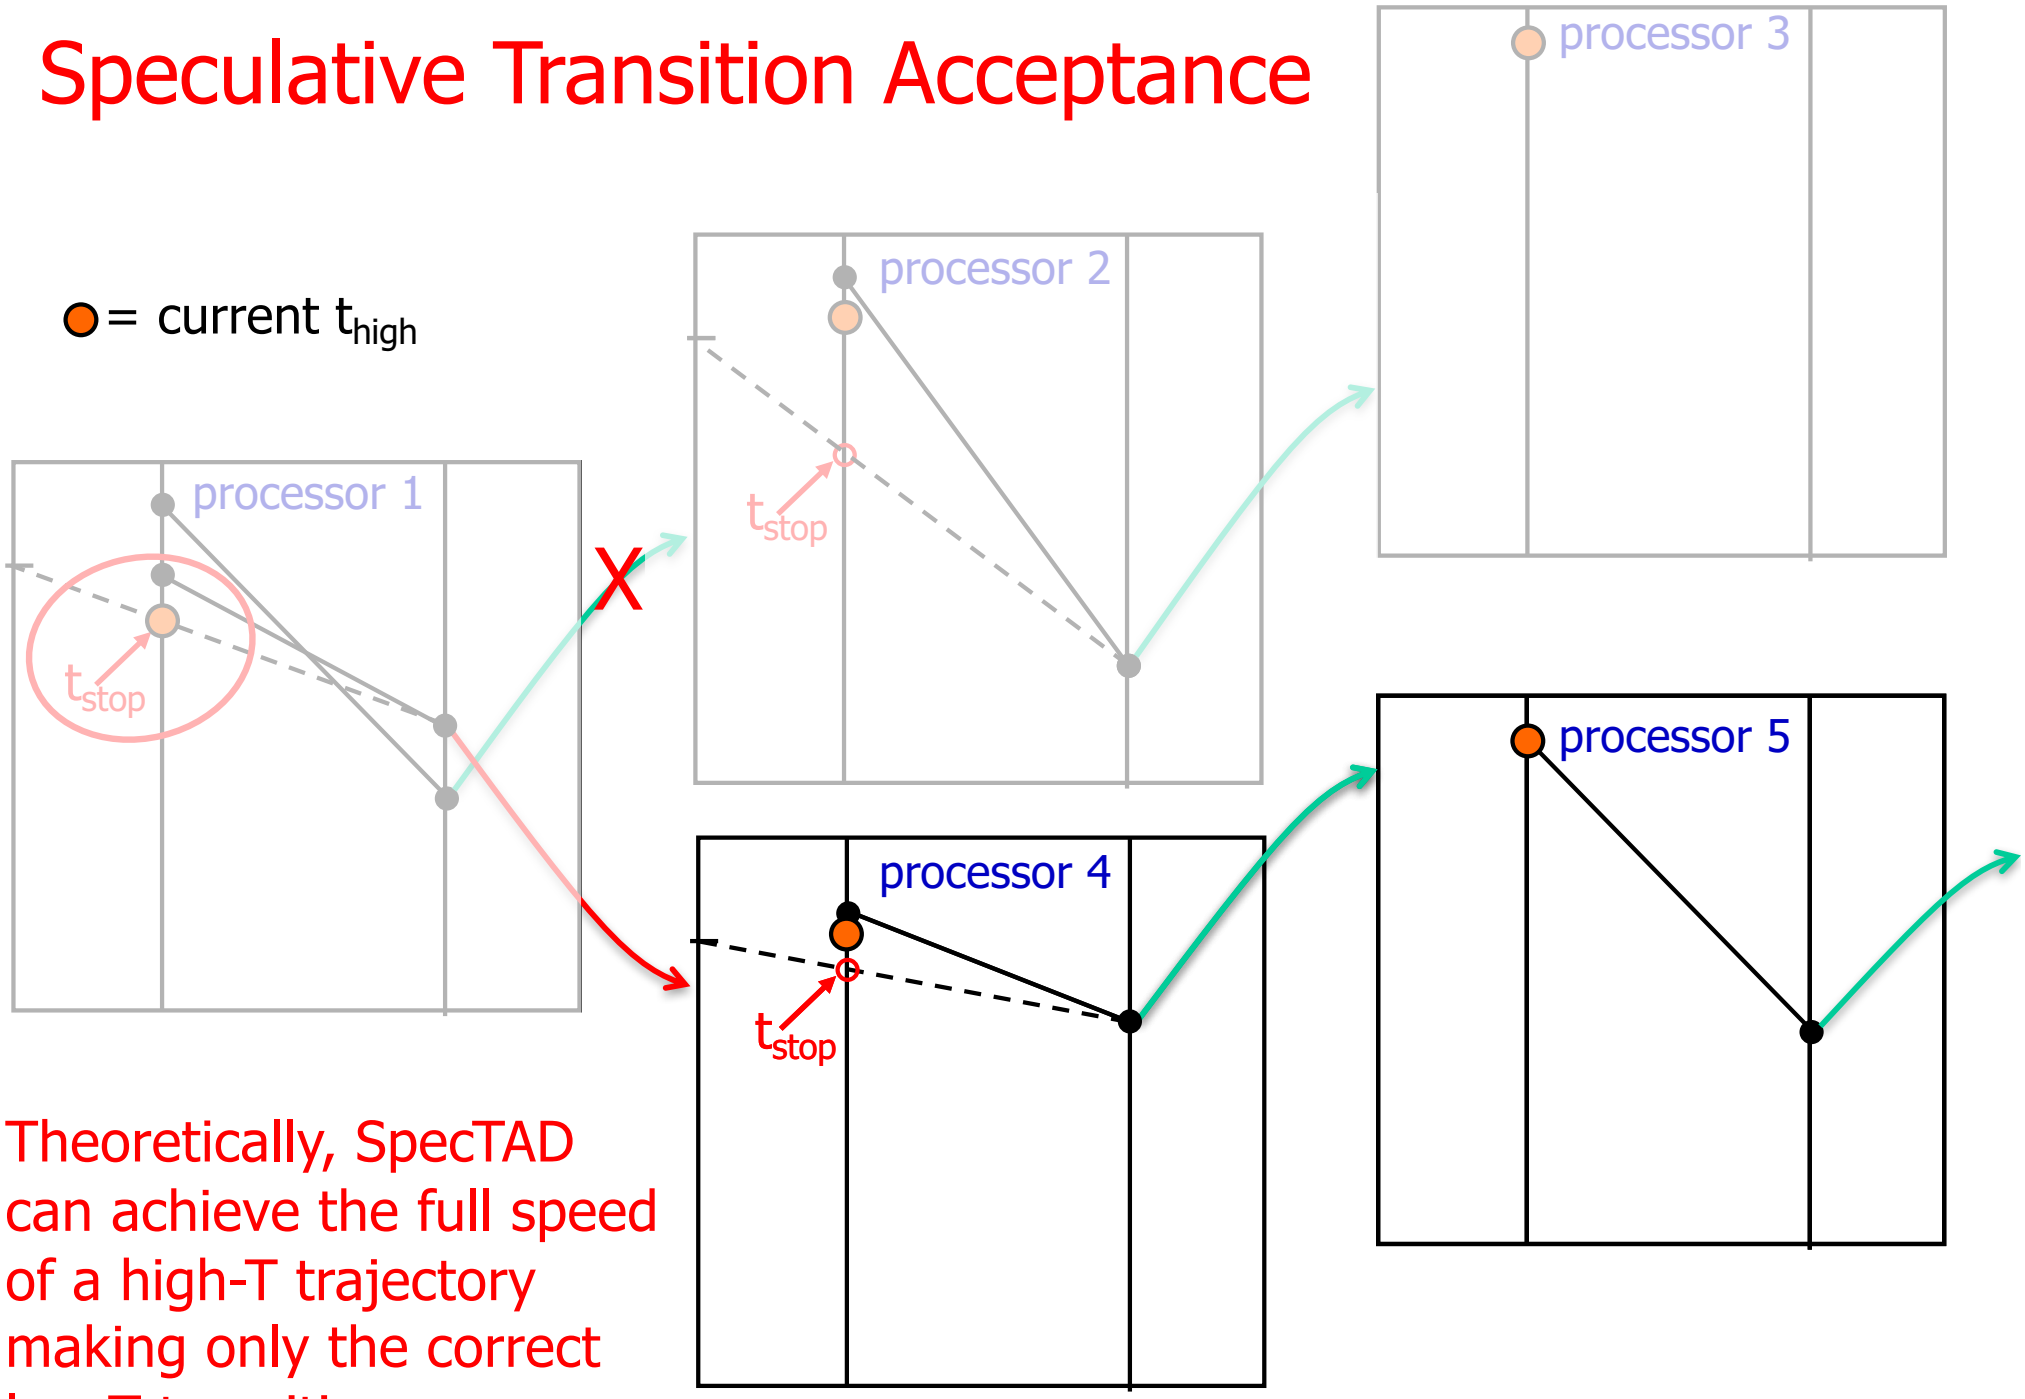

Theoretically, SpecTAD can achieve the full speed of a high-T trajectory making only the correct low-T transitions.

SpecTAD example – vapor deposition

Cu/Ag(100), T=77K, 0.04 ML/s, 128 cores

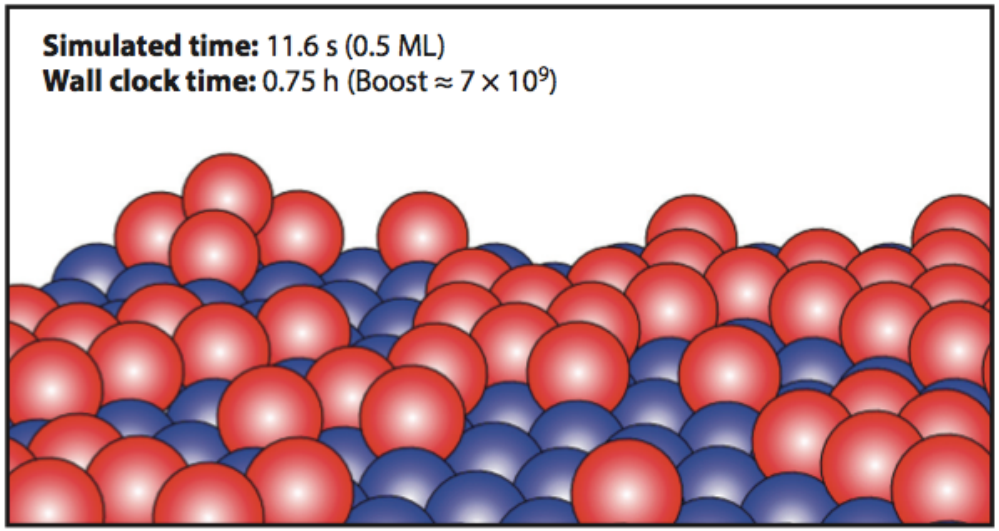

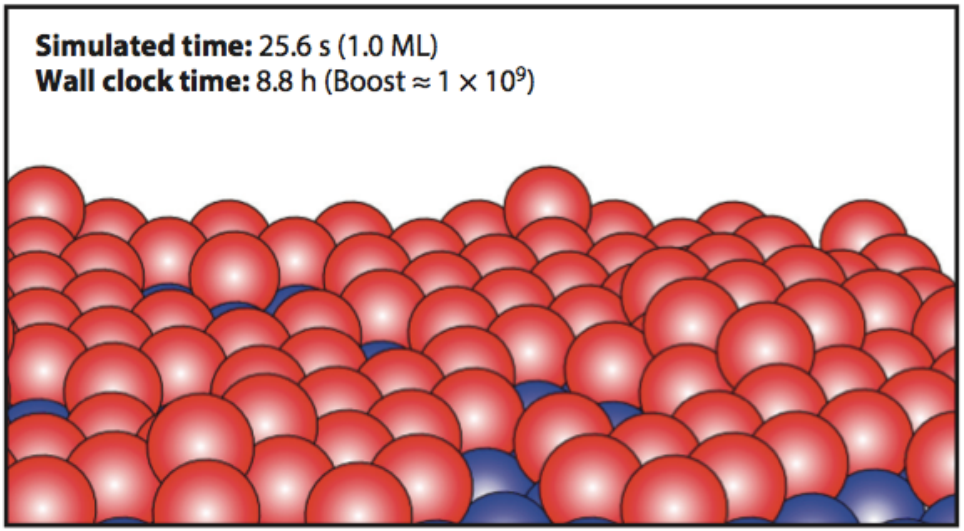

It took ~1 year to grow 1.5 ML with serial TAD
[Sprague et al, Phys. Rev. B 66, 205415 (2002)]

$MgAl_2O_4$ Spinel

diffusion of small defect cluster (O vac + Mg vac)

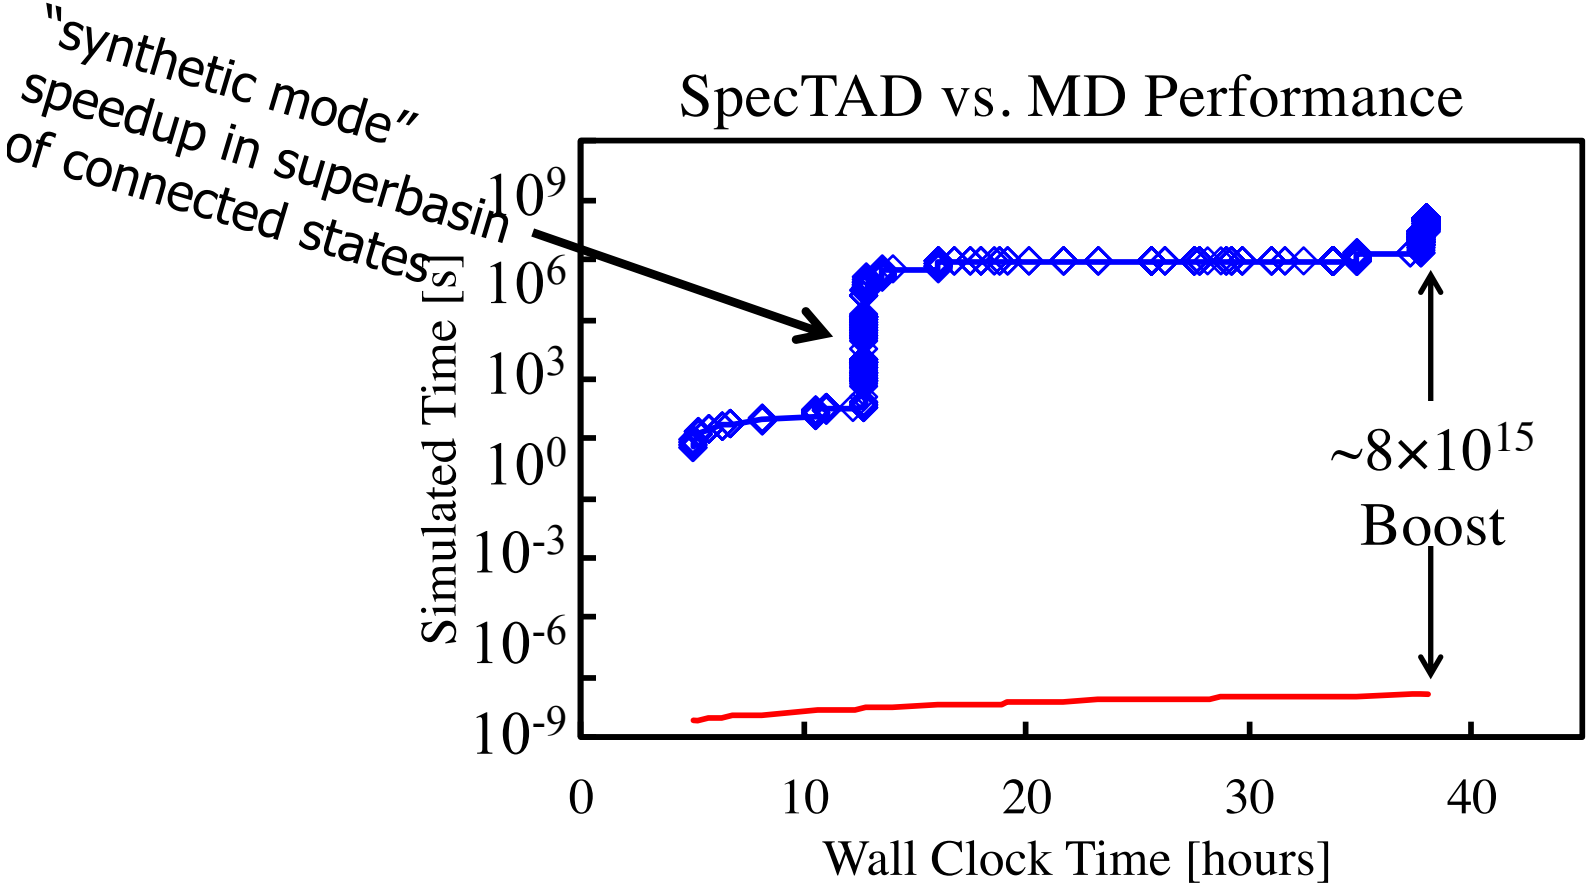

With regular TAD, we probably would not have reached the synthetic-mode jump in time.

Zamora, Voter, Perez, Perriot, Uberuaga, Phys. Chem. Chem. Phys. (2016).

Discrete Event Simulation (DES) of Performance

Simulation of a simulation - runs much faster than a real simulation – can explore complicated parameter space.

Counts up the force calls (the most expensive part of TAD).

Properly accounts for communication lag, serial vs. overlapping parallel occurrences, etc.

Employs a simple hardware model.

Parameterized for Ag/Ag(100) case:

- cpu time for a force call
- rates and barriers for all pathways
- no. of forces to converge NEB
- anharmonicity in rate w/ T_{high}

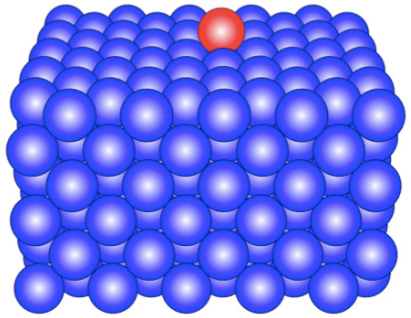

Collaboration with Stephan Eidenbenz, Nandu Santhi, Sunil Thulasidasan, and Sue Mniszewski

Zamora et al, submitted

SpecTADSim Validation

Validation of SpecTADSim using authentic TAD and SpecTAD software

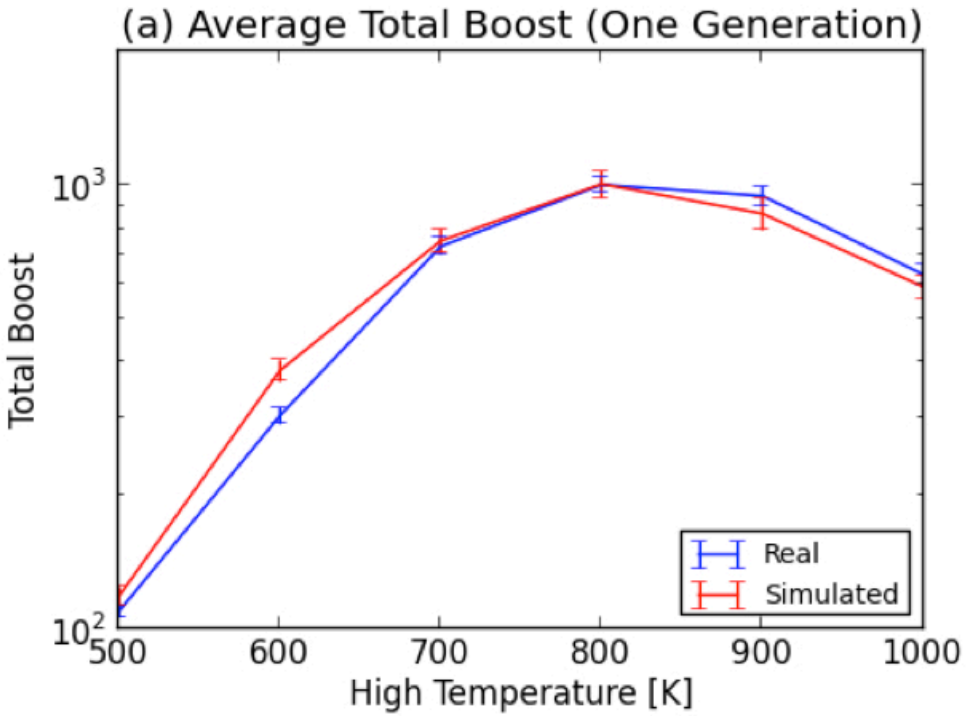

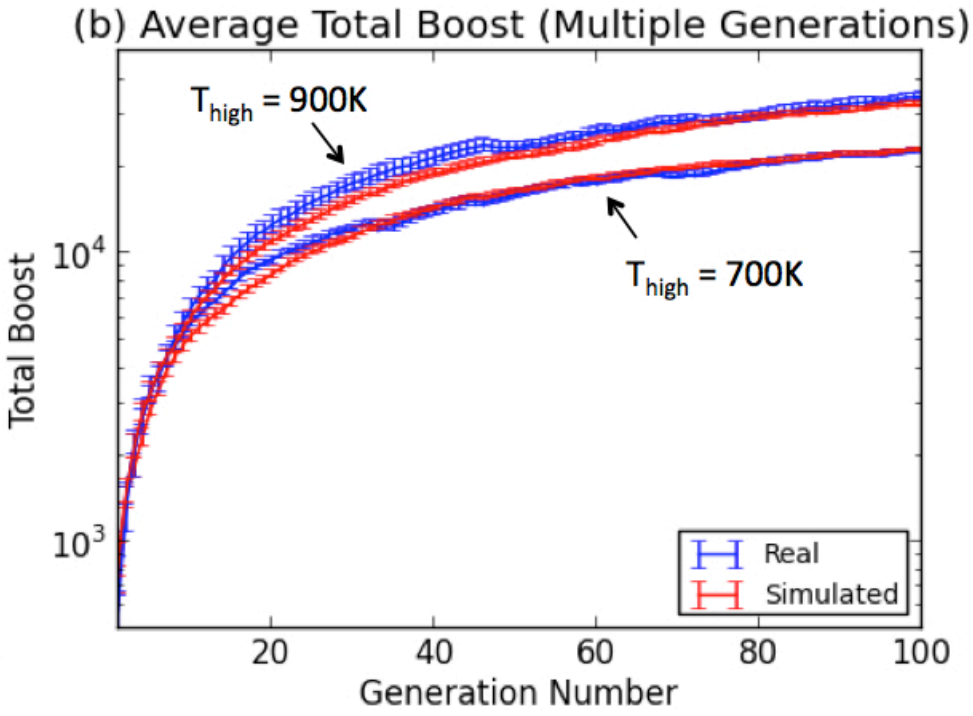

Boost in SpecTAD

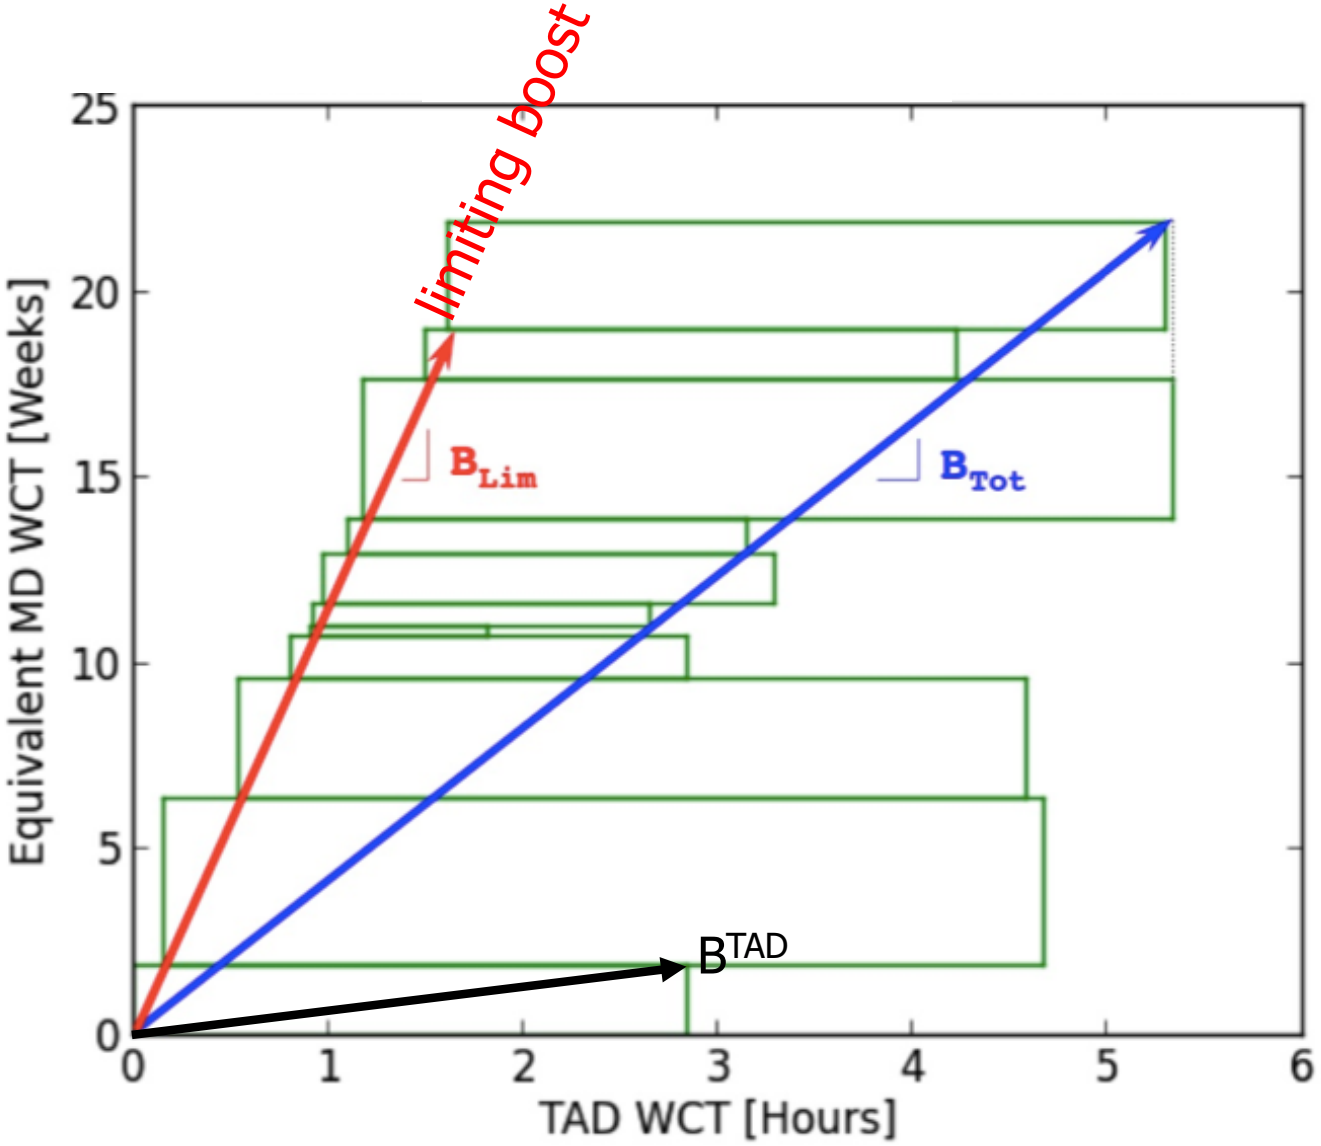

Initial (and final) transient in core count and boost

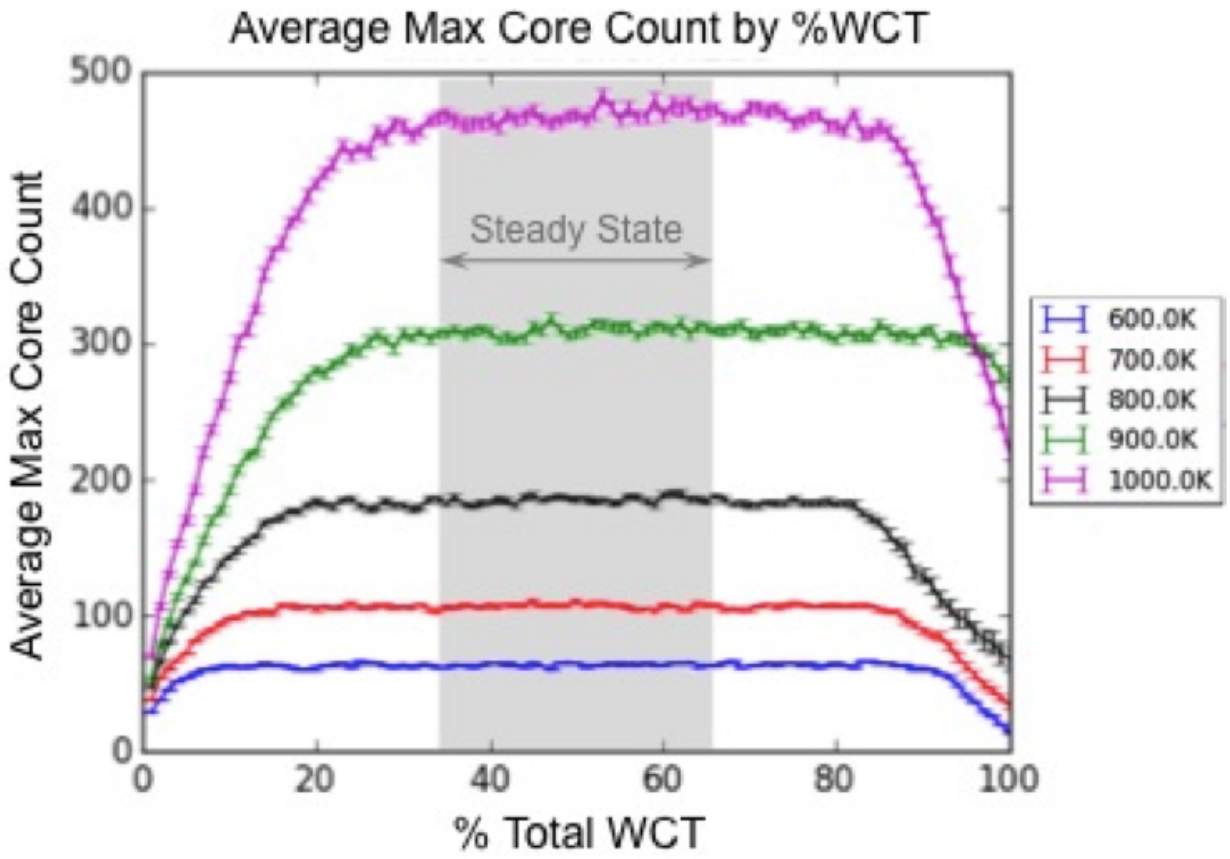

SpecTADSim Limiting Boost Results

Best settings produce >300 times parallel speedup

Worst settings produce ~3-8 times parallel speedup

BN = *BeforeNEB* branching (dashed lines) – good speedup, but uses many cores

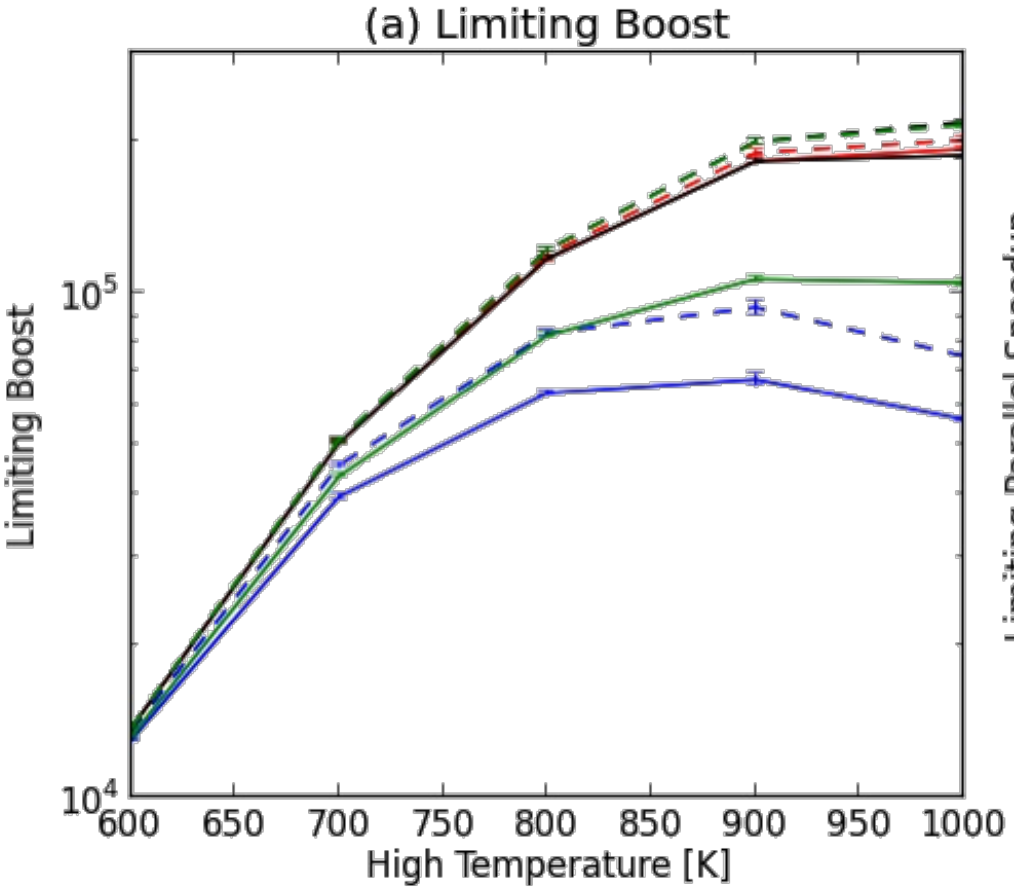

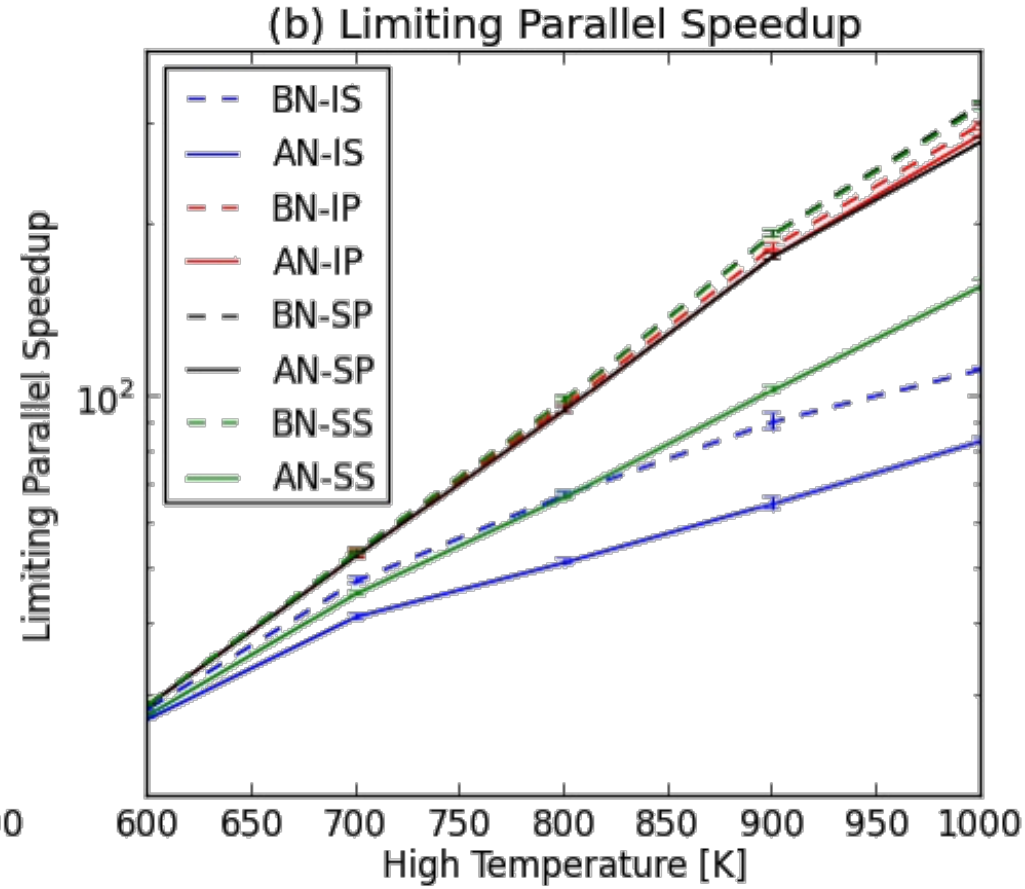

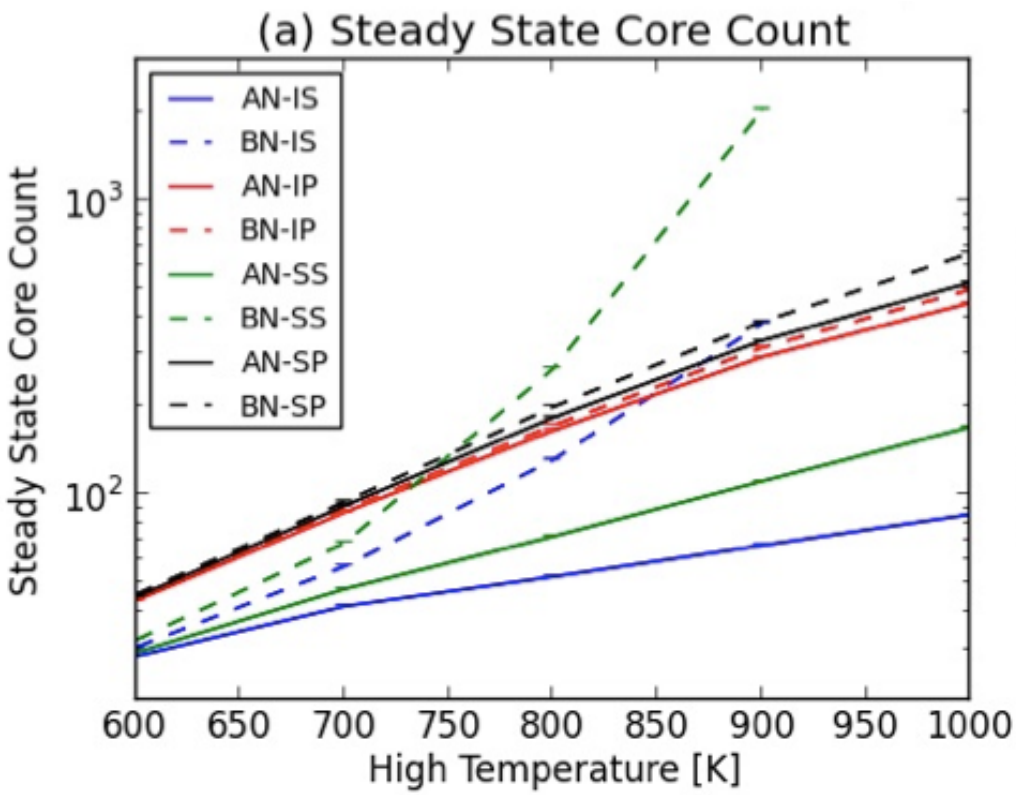
(a) Steady State Core Count
AN-IS
BN-IS
AN-IP
BN-IP
AN-SS
BN-SS
AN-SP
BN-SP
Steady State Core Count
10^3
10^2
600 650 700 750 800 850 900 950 1000
High Temperature [K]

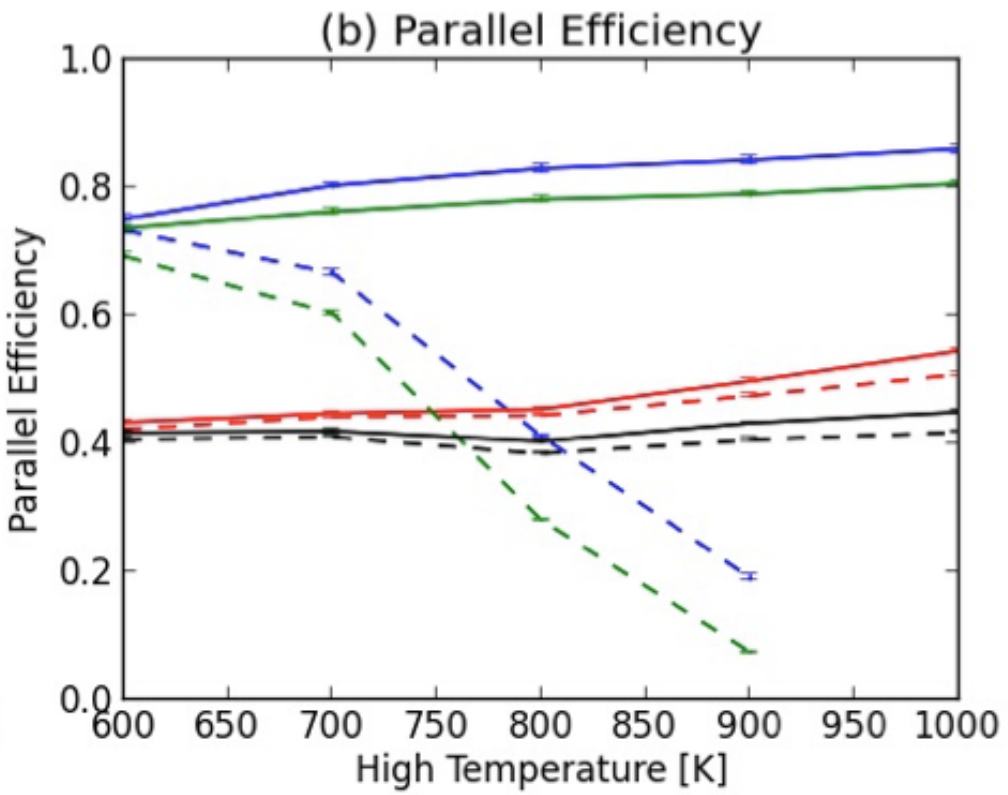
(b) Parallel Efficiency
Parallel Efficiency
1.0
0.8
0.6
0.4
0.2
0.0
600 650 700 750 800 850 900 950 1000
High Temperature [K]

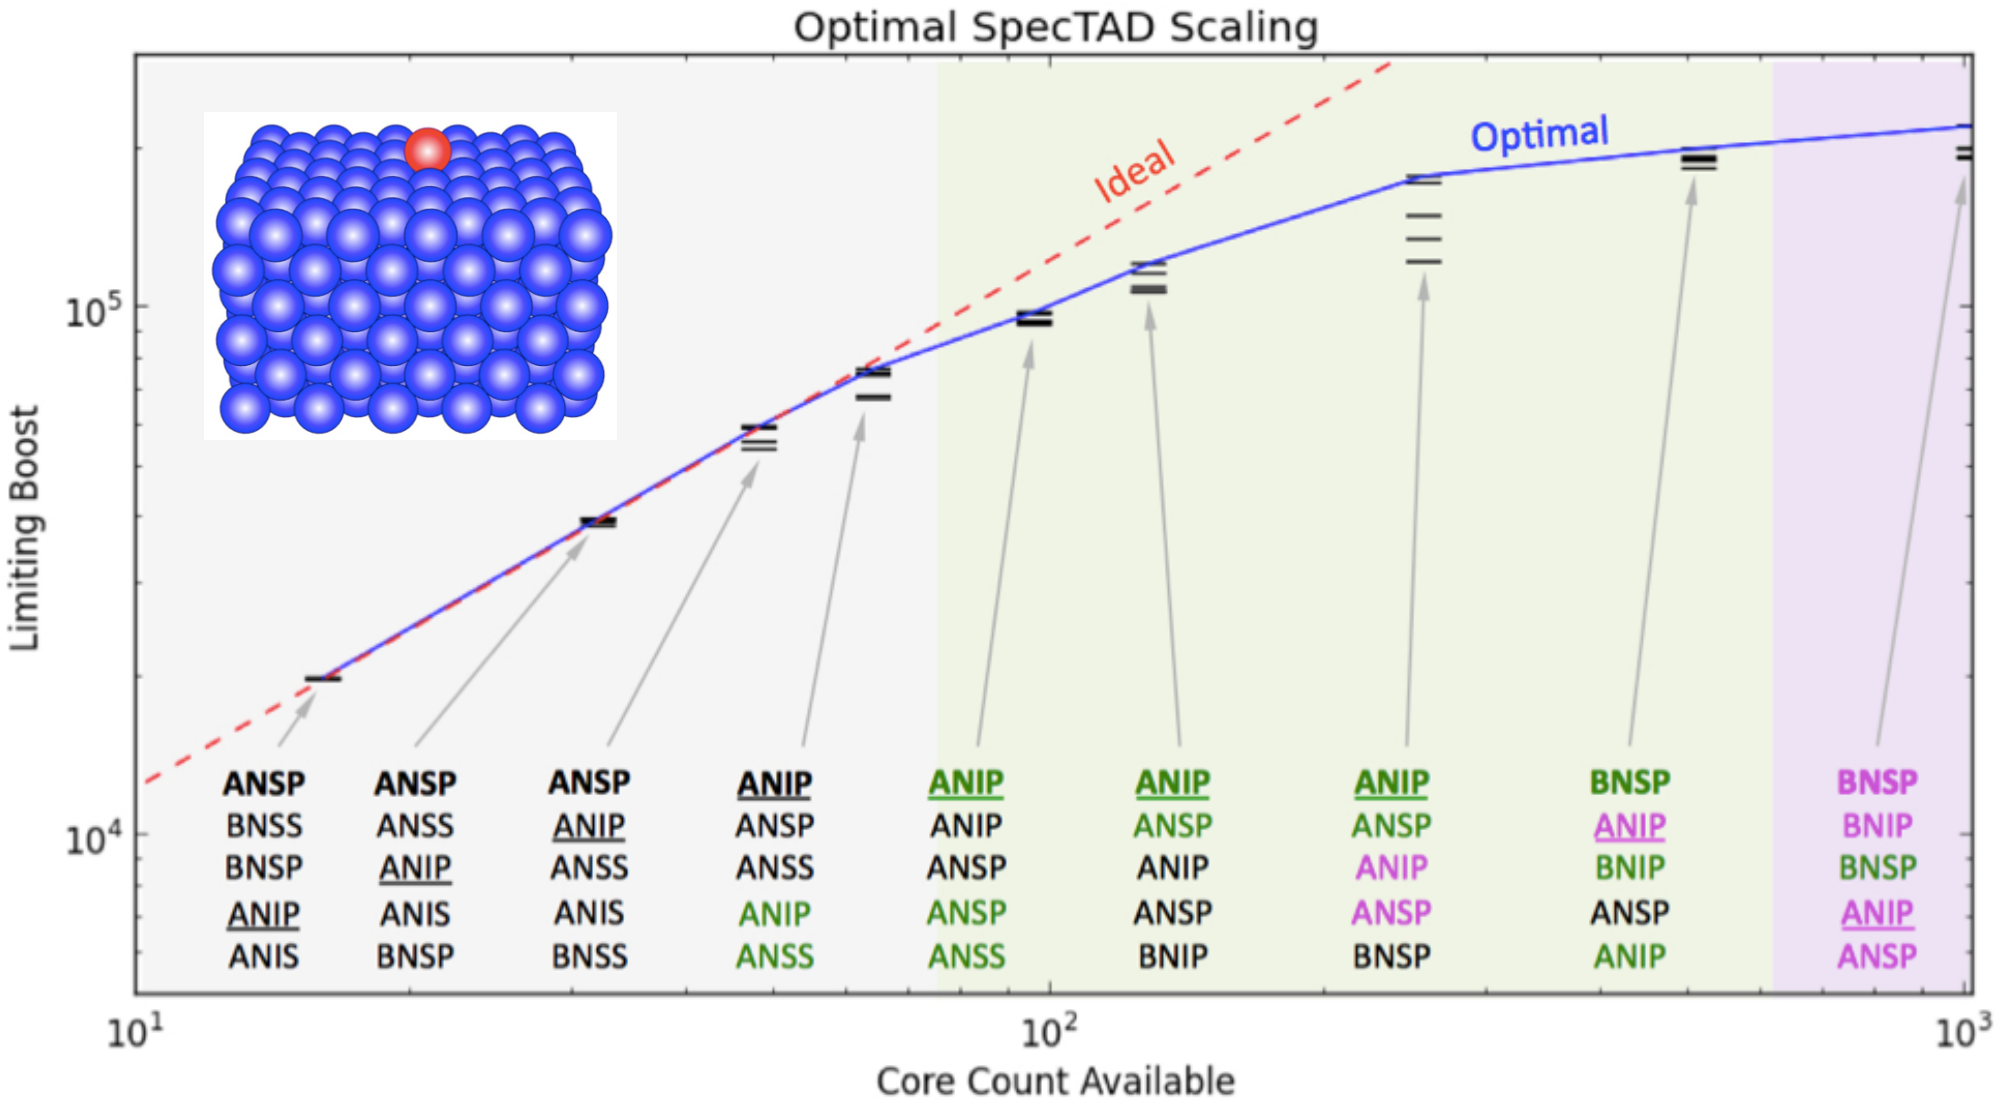
Optimal SpecTAD Scaling
Ideal
Optimal
Limiting Boost
10^5
10^4
10^1
10^2
10^3
Core Count Available
ANSP BNSS BNSP ANIP ANIS
ANSP ANSS ANIP ANIS BNSP
ANSP ANIP ANSS ANIS BNSS
ANIP ANSP ANSS ANIP ANSS
ANIP ANIP ANSP ANSP ANSS
ANIP ANSP ANIP ANSP BNIP
ANIP ANSP ANIP ANSP BNSP
BNSP ANIP BNIP ANSP ANIP
BNSP BNIP BNSP ANIP ANSP

Los Alamos

Vicinal Surface Growth

Cu(223)

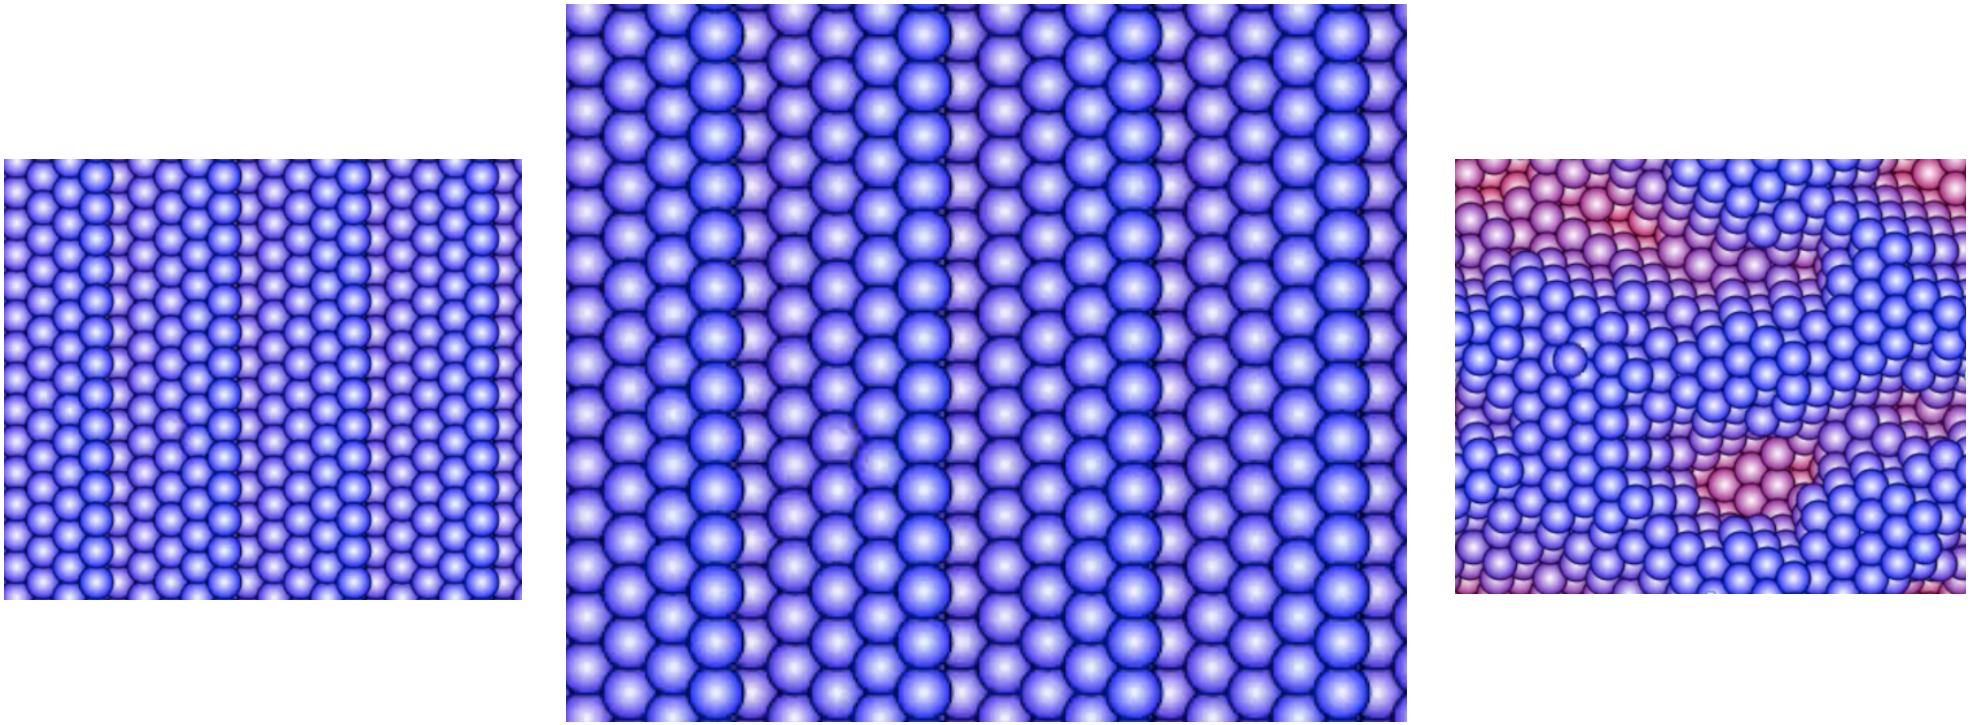

~10.5 min (0.005 ML/s), 77K

0.1 eV incident energy

Los Alamos

A case TAD *cannot* do

Stretching a silver nanowire

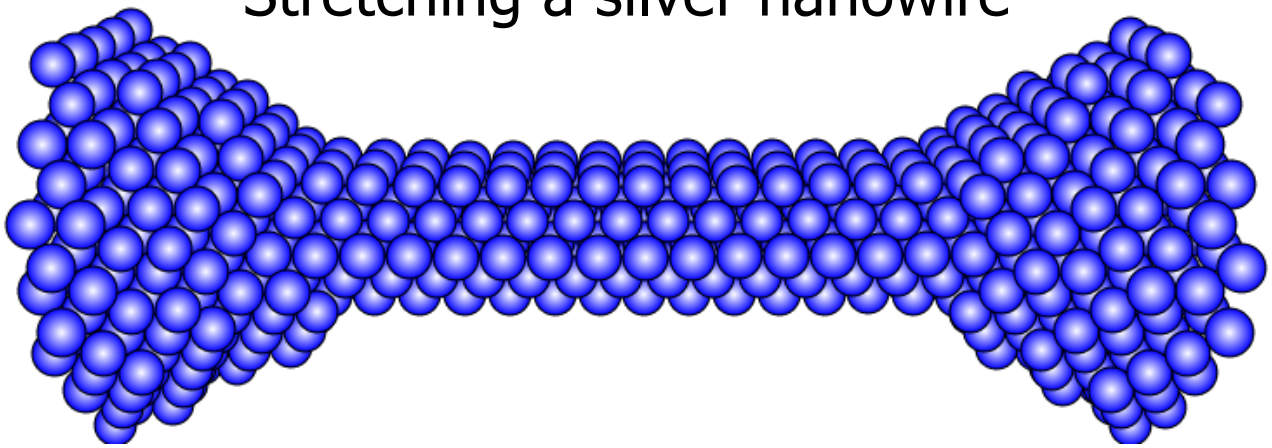

Surprisingly small pre-exponential factor for first slip event:

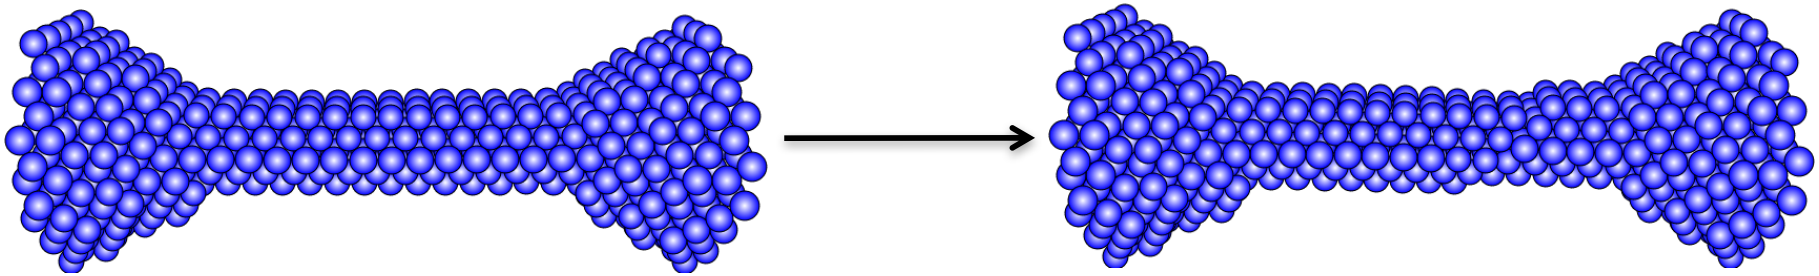

$k = \nu_0 \exp(-E_a/k_bT)$ $\nu_0 \sim 10^9$ Hz! (or much lower!)

We find these very large or very small prefactors for systems under strain.

Small Ag nanopillar

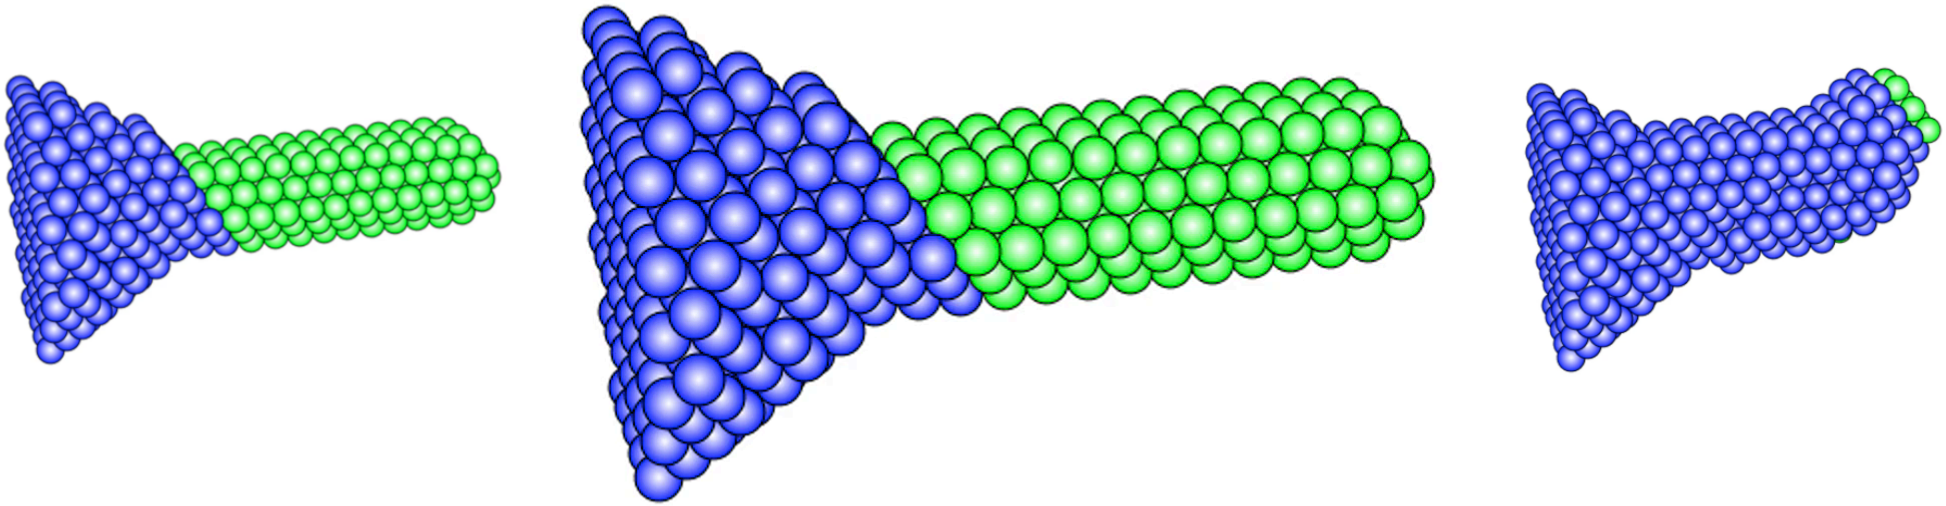

~36 ms, 300K

Los Alamos

Ways to parallelize TAD

Use speculation (SpecTAD) (as just described).

Use a spatially parallel force call (e.g., as in LAMMPS).

Use parallel replica dynamics (ParRep) to parallelize the high-temperature time.

Ways to parallelize TAD

Our latest code combines these.

- Use speculation (SpecTAD) (as just described).
- Use a spatially parallel force call (e.g., as in LAMMPS).
- Use parallel replica dynamics (ParRep) to parallelize the high-temperature time.

Ready for release soon - contact us if you would like to use it.

Ways to parallelize TAD

Our latest code combines these.

- Use speculation (SpecTAD) (as just described).
- Use a spatially parallel force call (e.g., as in LAMMPS).
- Use parallel replica dynamics (ParRep) to parallelize the high-temperature time.

Ready for release soon - contact us if you would like to use it.

Use parallel trajectory splicing (ParSplice) instead of ParRep to parallelize even further (see Danny Perez's talk).

Spatially parallelize using synchronous sub-lattice method [Shim, Amar, Uberuaga, Voter, Phys. Rev. B, (2007)]

Combining SpecTAD and ParRep

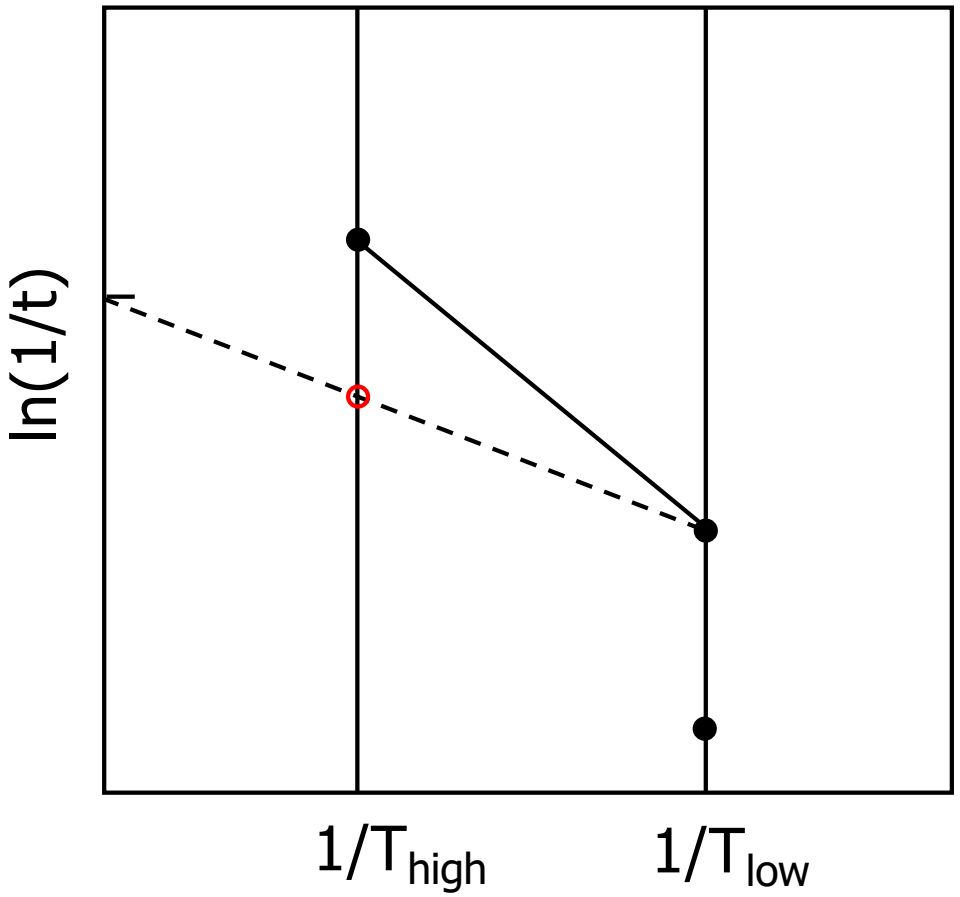

Combining SpecTAD and ParRep

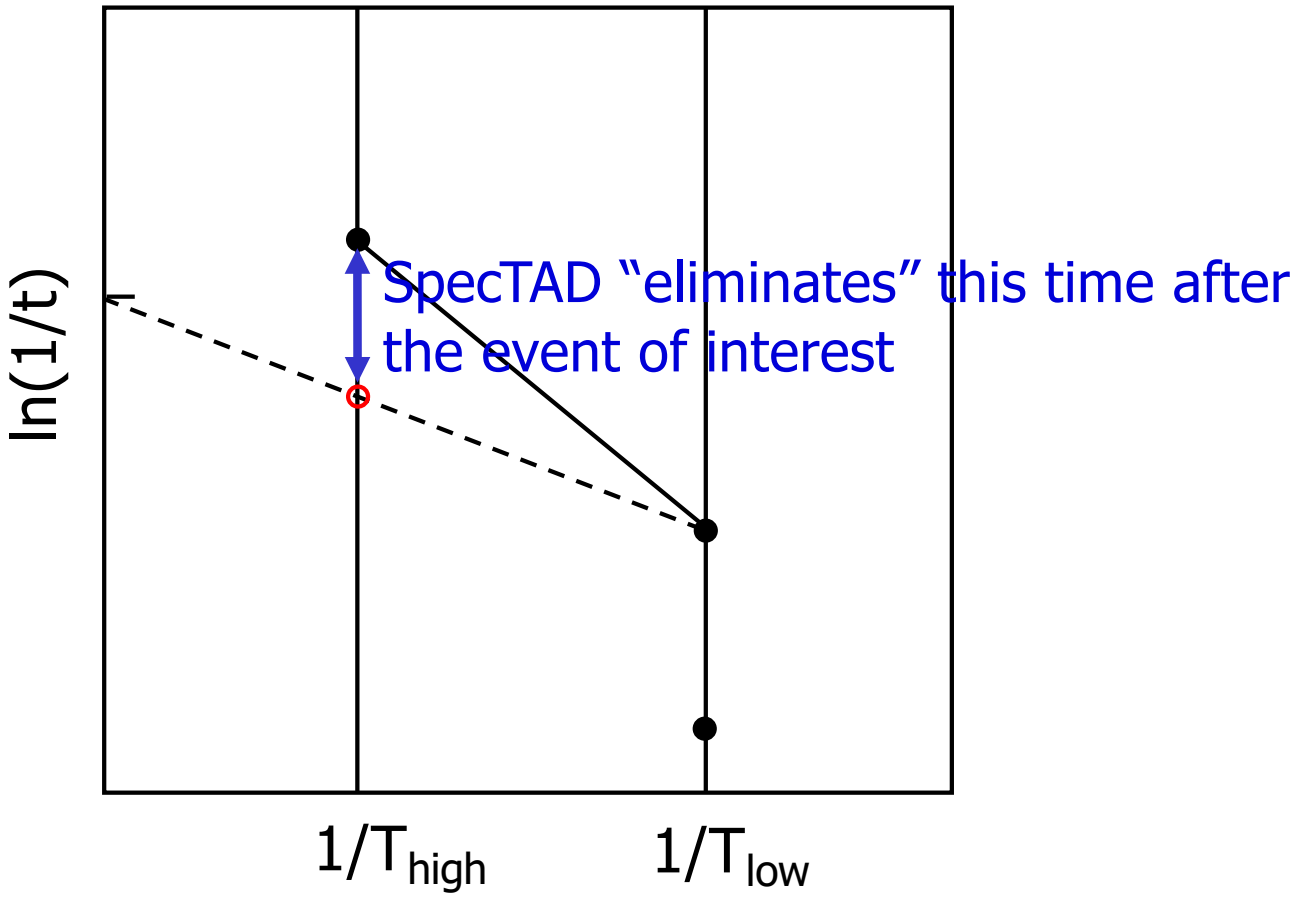

Combining SpecTAD and ParRep

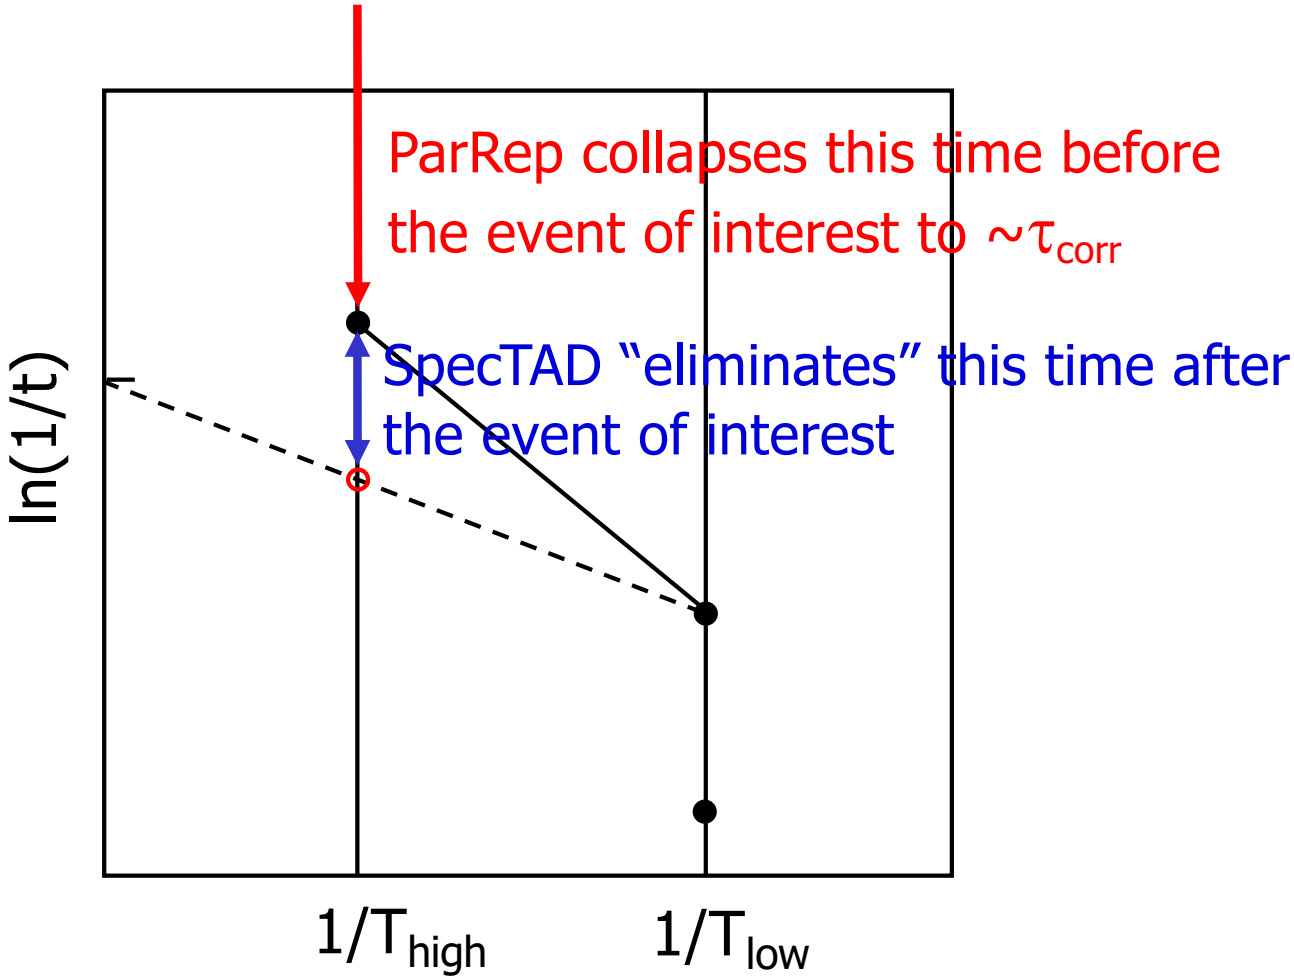

- Escape from any basin in time ~τ_{corr} (e.g., ~1 ps for metal)

Combining SpecTAD and ParRep

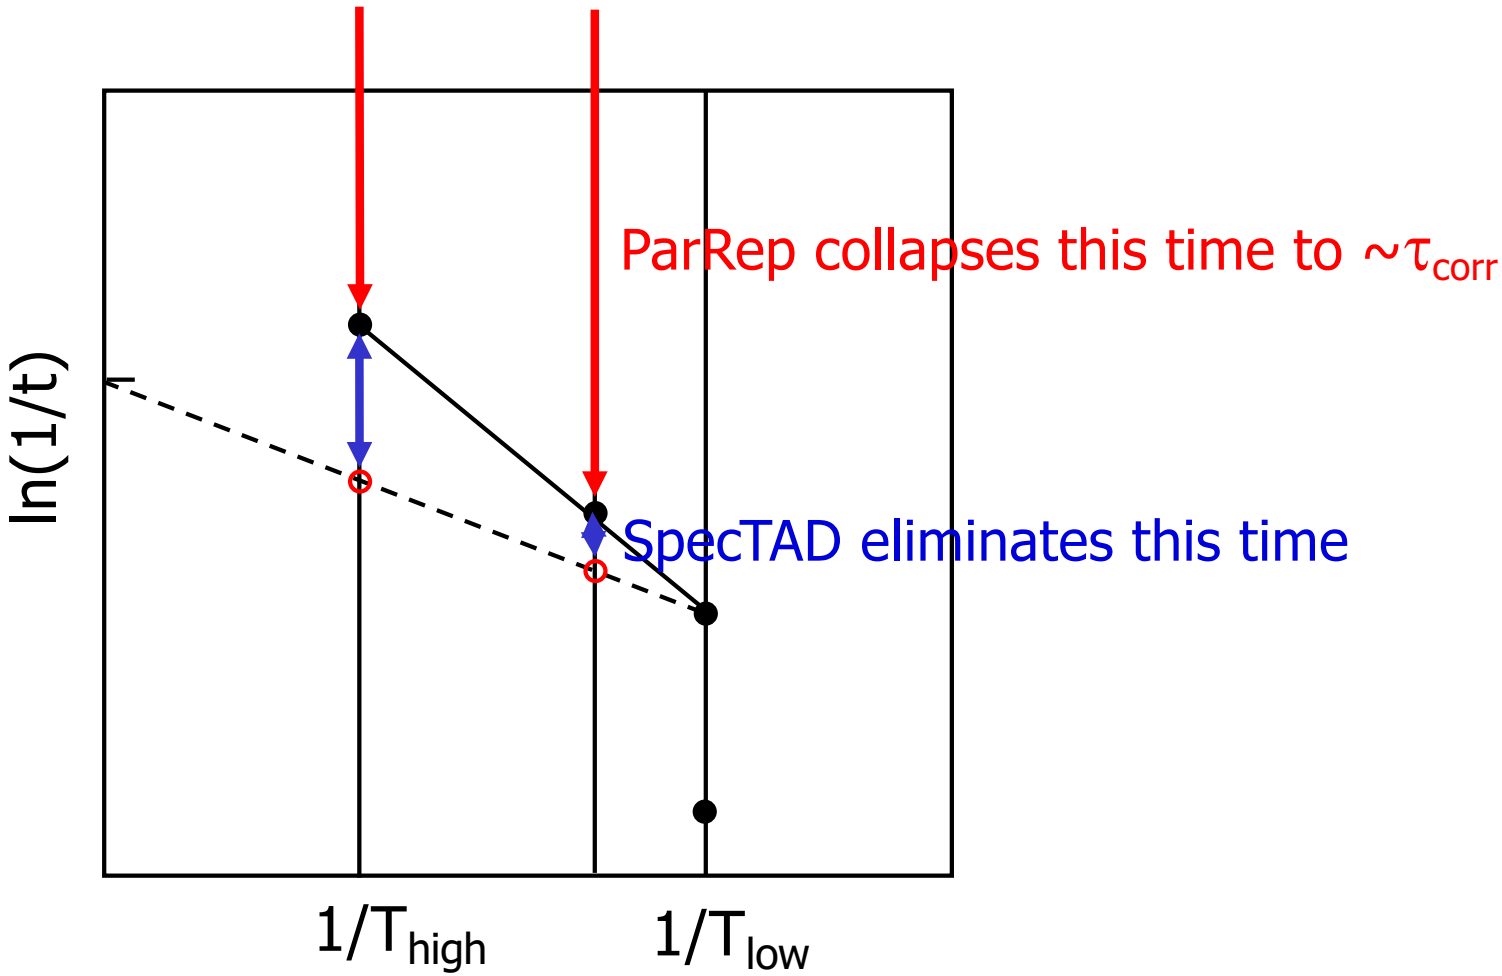

- Escape from any basin in time ~τ_{corr} (e.g., ~1 ps for metal)
- Can also use this for higher TAD accuracy at same wall clock cost
- ParSplice will do even better...

Some final points

- Speculation can be used effectively to increase parallelization.
- Discrete event simulation (DES) is a powerful way to examine the parameter space, and sometimes there are unexpected relationships.
- With unlimited number of processors, the fastest approach is to spawn the new TAD process as soon as the transition is observed at high T and also spawn a NEB to find the saddle.
- With finite number of processors, it can be just as fast (maybe faster) to wait until NEB is finished before spawning. Parallelizing the NEB is then important.
- In a superbasin, SpecTAD in synthetic mode becomes so fast that even one force call slows down the state-to-state rate.