Leakage-Aware Thermal Management for Multi-Core Systems Using Piecewise Linear Model Based Predictive Control

Hai Wang

University of Electronic Science & Technology of China

ASP-DAC 2019, Tokyo, Japan

Outline

- Background
- Problems of leakage-aware DTM
- PWL predictive DTM
- Experimental results
- Conclusion

Two battles lost against leakage

- Leakage power does not scale like dynamic power
  - Power density increases with technology scaling (Dennard scaling lost)
- Power (heat) removal ability remains the same

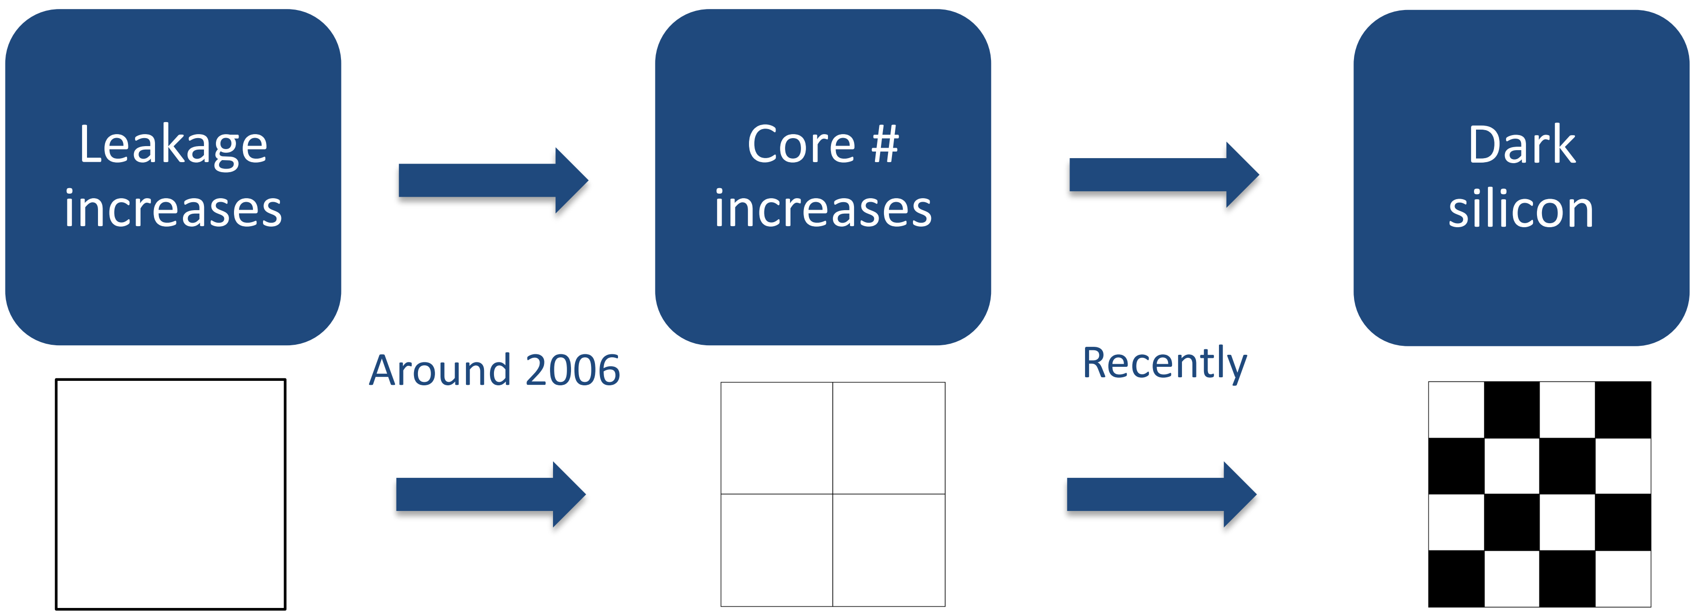

Fix core #
Increase frequency
Best days in performance increase!

Fix frequency
Increase core #

Not all cores operates @ full freq anymore
We lost Dennard scaling
Solutions needed!

Leakage problems in the new era

- How to determine the active core distributions and power budget?
- Our solution: Greedy Dynamic Power (GDP)
  - Locate active core positions at runtime
  - Compute power budget for each core

H. Wang, *et al.*, "GDP: A Greedy based dynamic power budgeting method for multi/many-core systems in dark silicon", *IEEE Trans. on Computers*, 2019

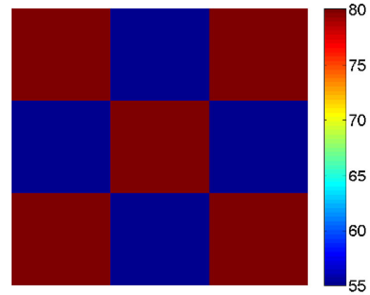

(a) 9-core system with 5 active cores.

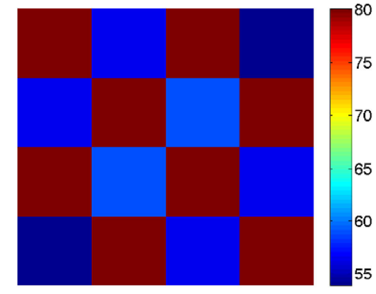

(b) 16-core system with 8 active cores.

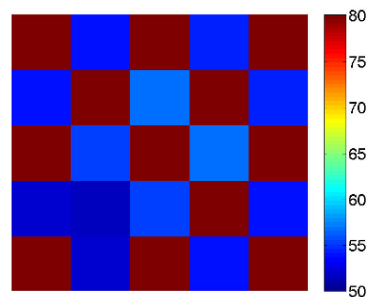

(c) 25-core system with 12 active cores.

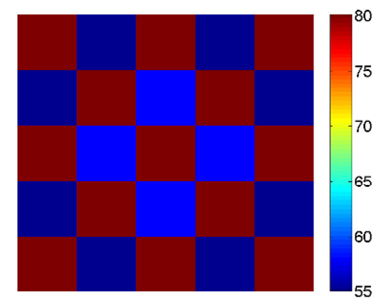

(d) 25-core system with 13 active cores.

Leakage problems in the new era

- How to estimate leakage power distribution at runtime?
- Our solution: Piecewise linear MOR based fast simulation
  - Piecewise linear (PWL) approximation
  - Incremental MOR on local models

H. Wang, *et al.*, "A Fast Leakage-Aware Full-Chip Transient Thermal Estimation Method", *IEEE Trans. on Computers*, 2018

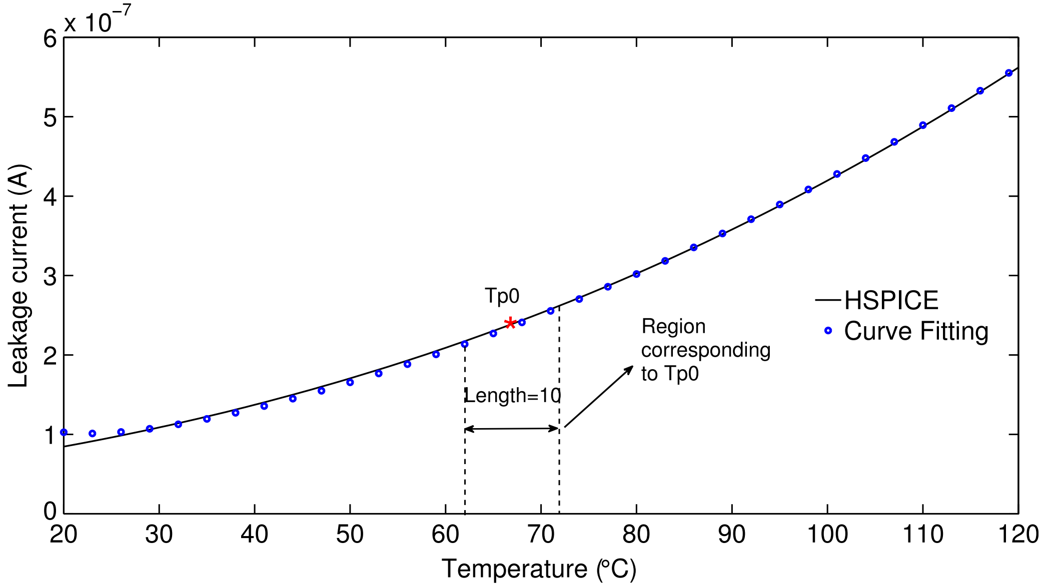

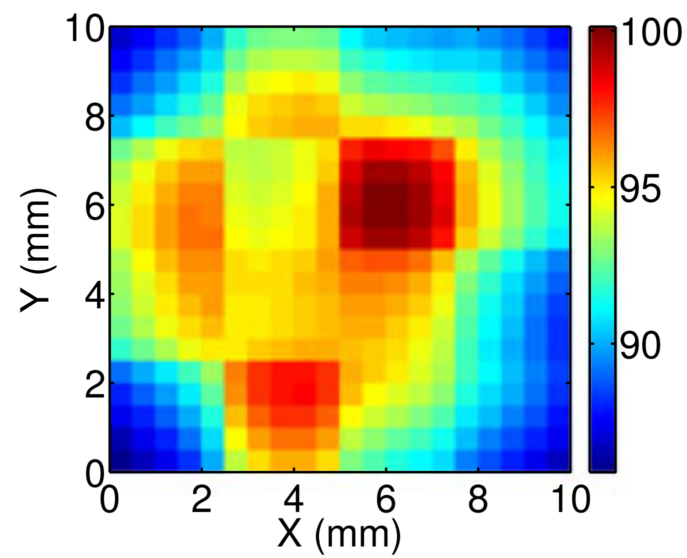

The remaining Leakage problem

- How to control the multi-core system temperature considering leakage?
- In another word, how to compute the dynamic power recommendation in leakage-aware DTM?

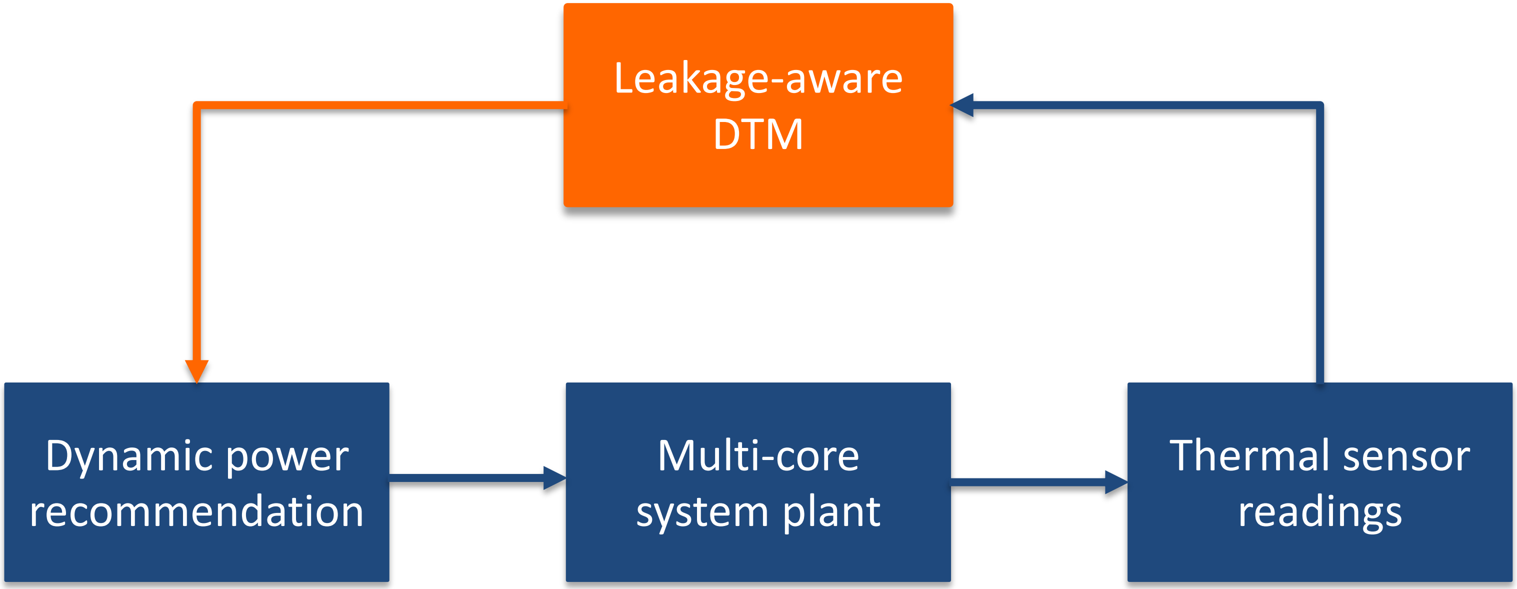

Outline

- Background
- Problems of leakage-aware DTM
- PWL predictive DTM
- Experimental results
- Conclusion

Basic framework of Predictive DTM

- The basic idea of predictive DTM
  - Compute the dynamic power recommendation P_d, which tracks the given target temperature
  - P_d can be solved by optimization using thermal prediction

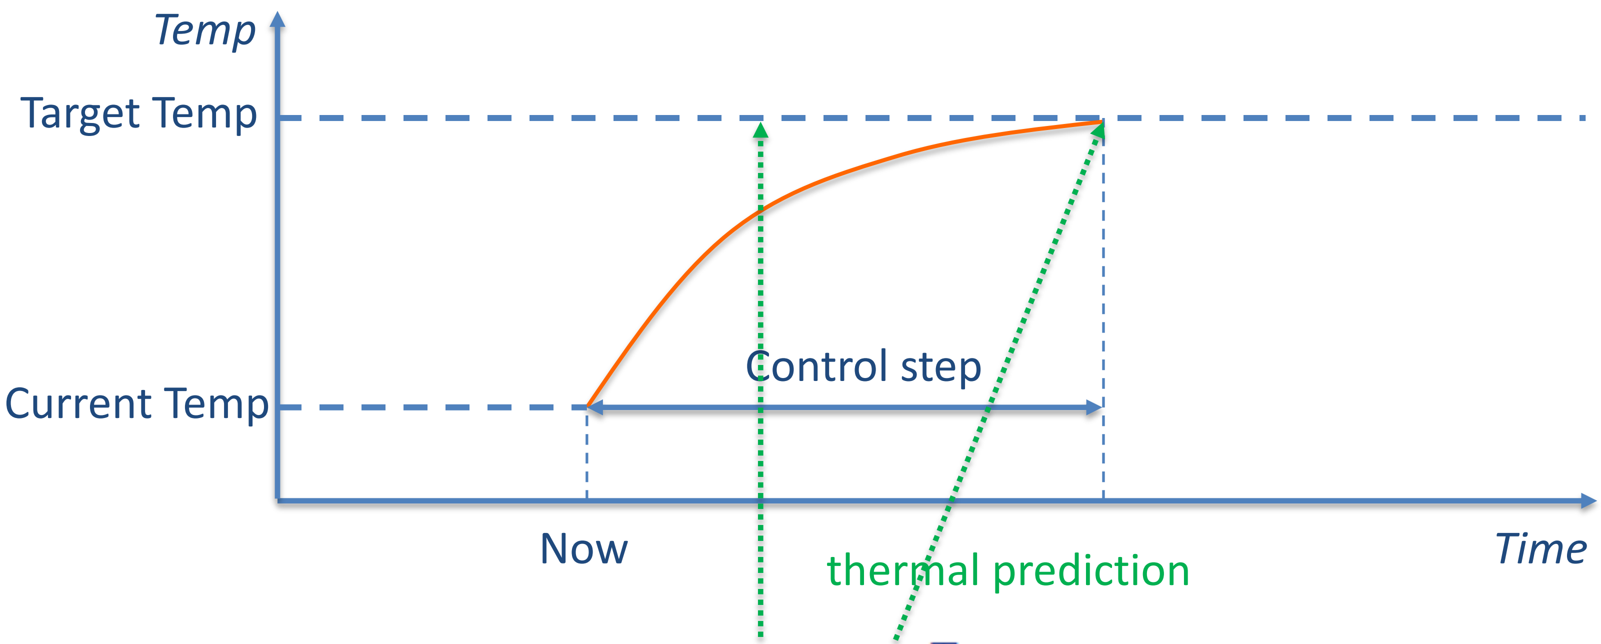

minimize $\mathcal{J} = (\mathcal{Y}_g - \mathcal{Y})^T (\mathcal{Y}_g - \mathcal{Y})$

The root of problem: leakage is nonlinear!

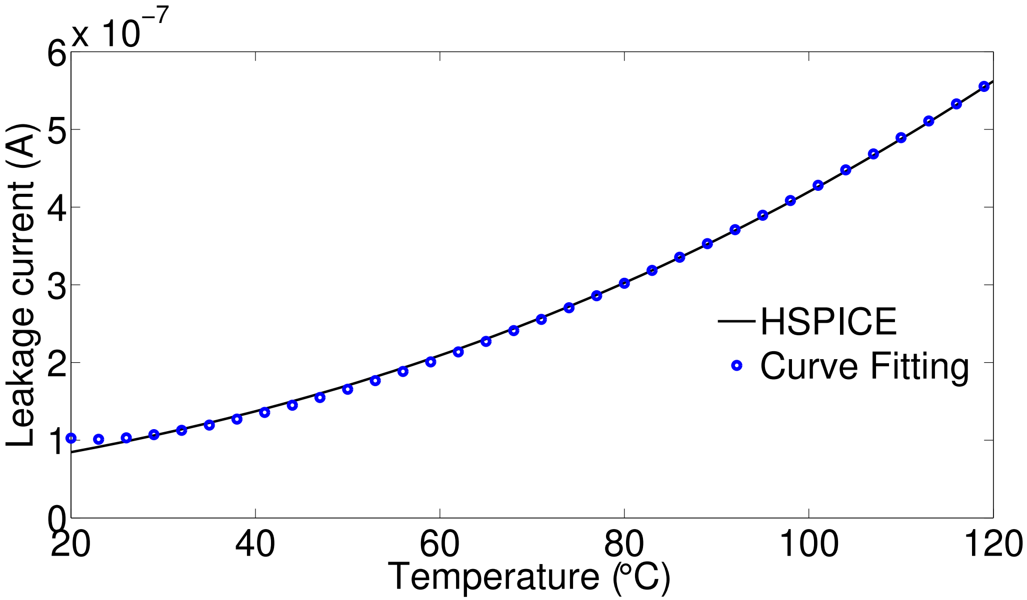

$$p_s = V_{dd} I_{leak} = V_{dd}(I_{sub} + I_{gate}),$$

$$I_{sub} = K {v_T}^2 e^{\frac{V_{GS}-V_{th}}{\eta v_T}} \left(1 - e^{\frac{-V_{DS}}{v_T}}\right)$$

$$\approx K {v_T}^2 e^{\frac{V_{GS}-V_{th}}{\eta v_T}}, \qquad v_T = \frac{k\boxed{T_p}}{q}$$

Temperature

- Leakage power depends on temperature nonlinearly

Problem caused by nonlinearity

$$GT(t) + C\frac{dT(t)}{dt} = B(P_d(t) + \boxed{P_s(T,t)})$$

nonlinear

Discretize

$$T(k+1) = AT(k) + DP_d(k) + \boxed{\int_0^h e^{-(h-\tau)C^{-1}G} C^{-1} B P_s(T,\tau)\,\mathrm{d}\tau,}$$

Plug in as thermal prediction

Compute P_d to minimize $\mathcal{J} = (\mathcal{Y}_g - \mathcal{Y})^T(\mathcal{Y}_g - \mathcal{Y})$

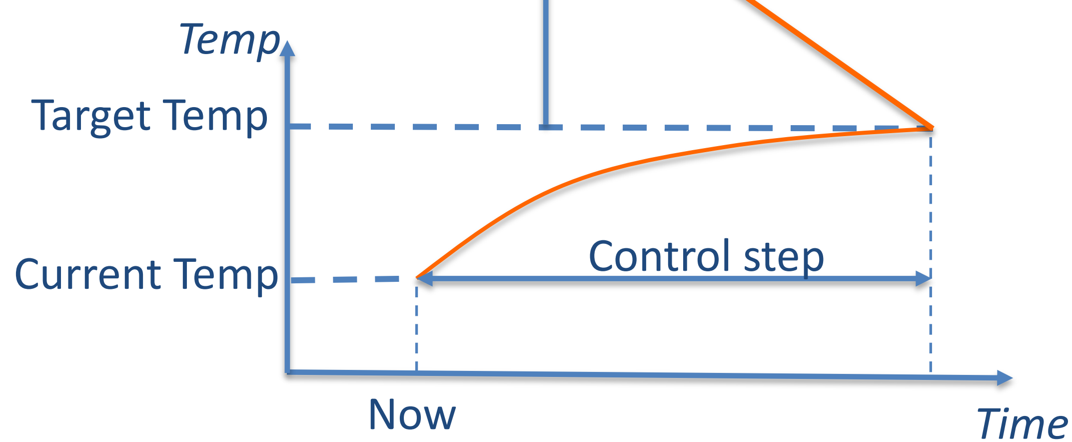

- Power recommendation P_d unsolvable due to nonlinearity

Piecewise linear (PWL) approximation?

- We used piecewise linear (PWL) approximation for leakage-aware thermal estimation before

$$GT(t) + C\frac{dT(t)}{dt} = B(P_d(t) + P_s(T,t))$$

$$\boxed{P_s = \widehat{P} + \widehat{H}T}$$

Taylor expansion

$$\widehat{G}T(t) + C\frac{dT(t)}{dt} = B(P_d(t) + \widehat{P})$$

Local linear thermal model at the local expansion point!

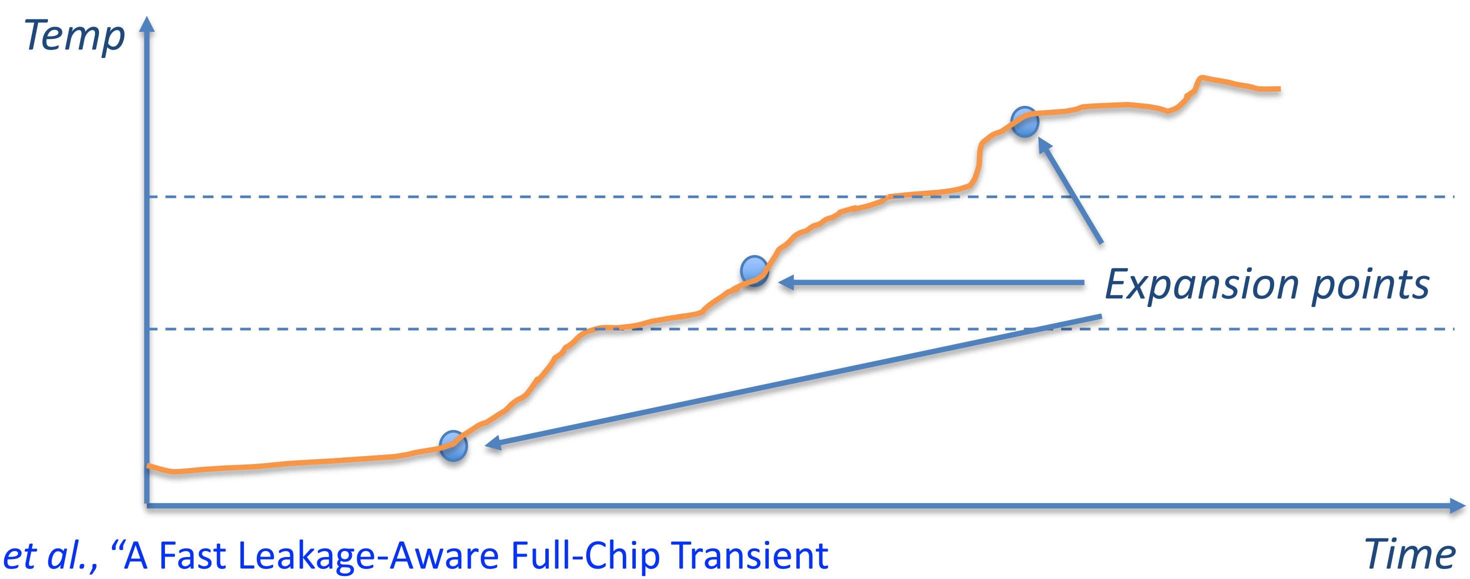

H. Wang, *et al.*, “A Fast Leakage-Aware Full-Chip Transient Thermal Estimation Method”, *IEEE Trans. on Computers*, 2018

Difficulty of PWL thermal prediction in DTM

- PWL cannot be used here directly
  - We do not know the temperature curve yet in DTM!
  - This is because power is the one to be solved (different from temp. estimation problem before)
  - We only know two things: current temp and target temp
  - How can we determine the expansion points in the prediction process?

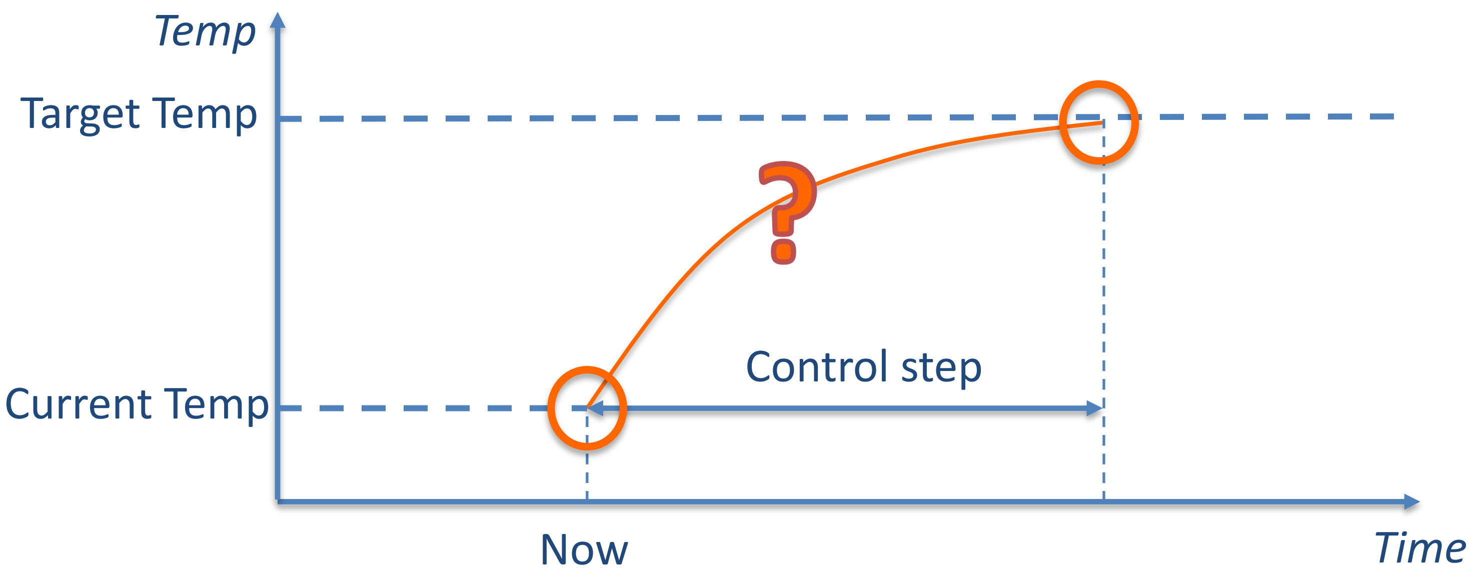

Outline

- Background
- Problems of leakage-aware DTM
- PWL predictive DTM
- Experimental results
- Conclusion

Determine expansion points in DTM

- Simulate the extreme curve: from ambient to the target
- All other curves should be above the extreme curve
- Put expansion points uniformly in Temp axis
- Determine model switching time points using extreme curve
- Expansion points of other curves can be determined

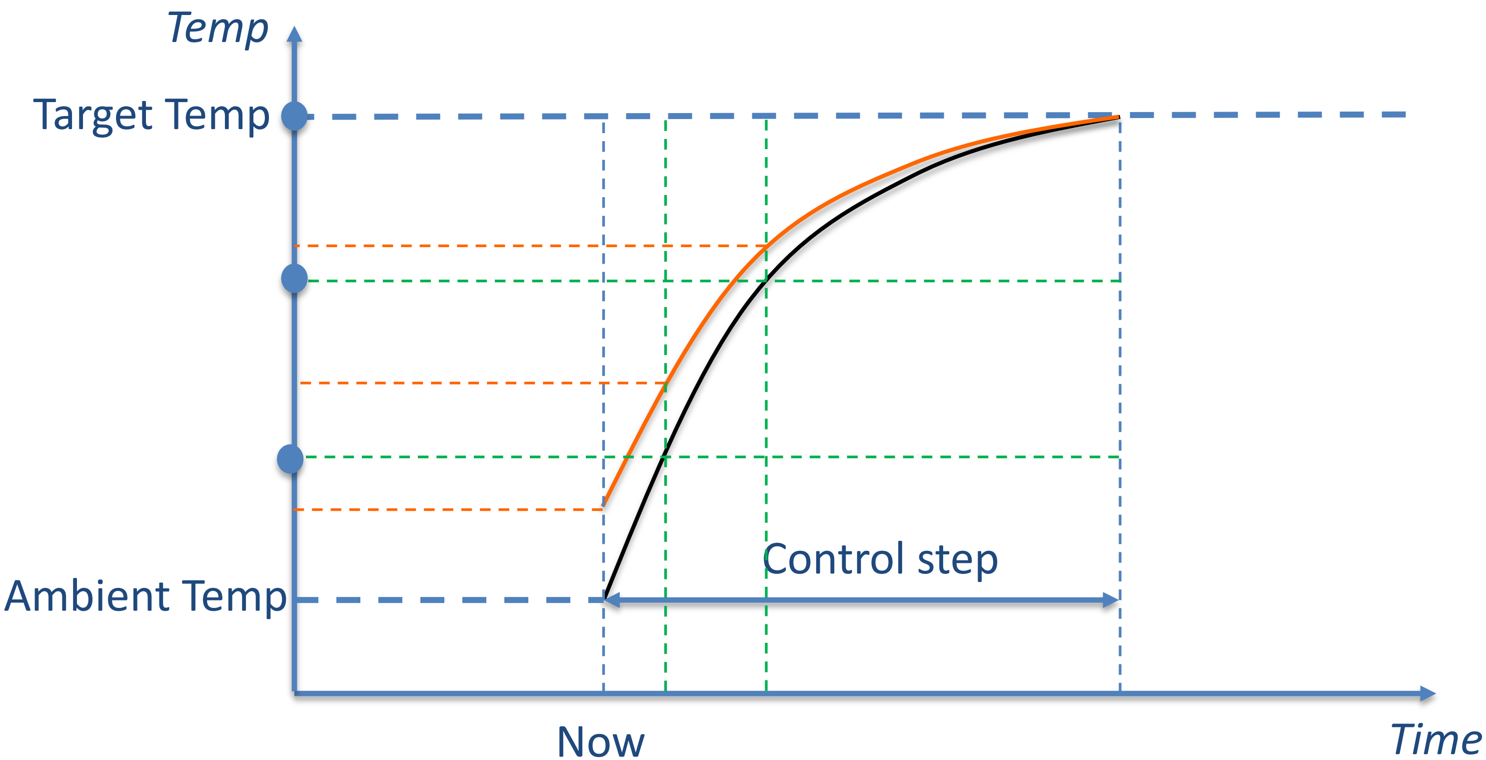

Determine expansion points in DTM

- If current temperature is already close to the target?
- Just use the target as expansion point!

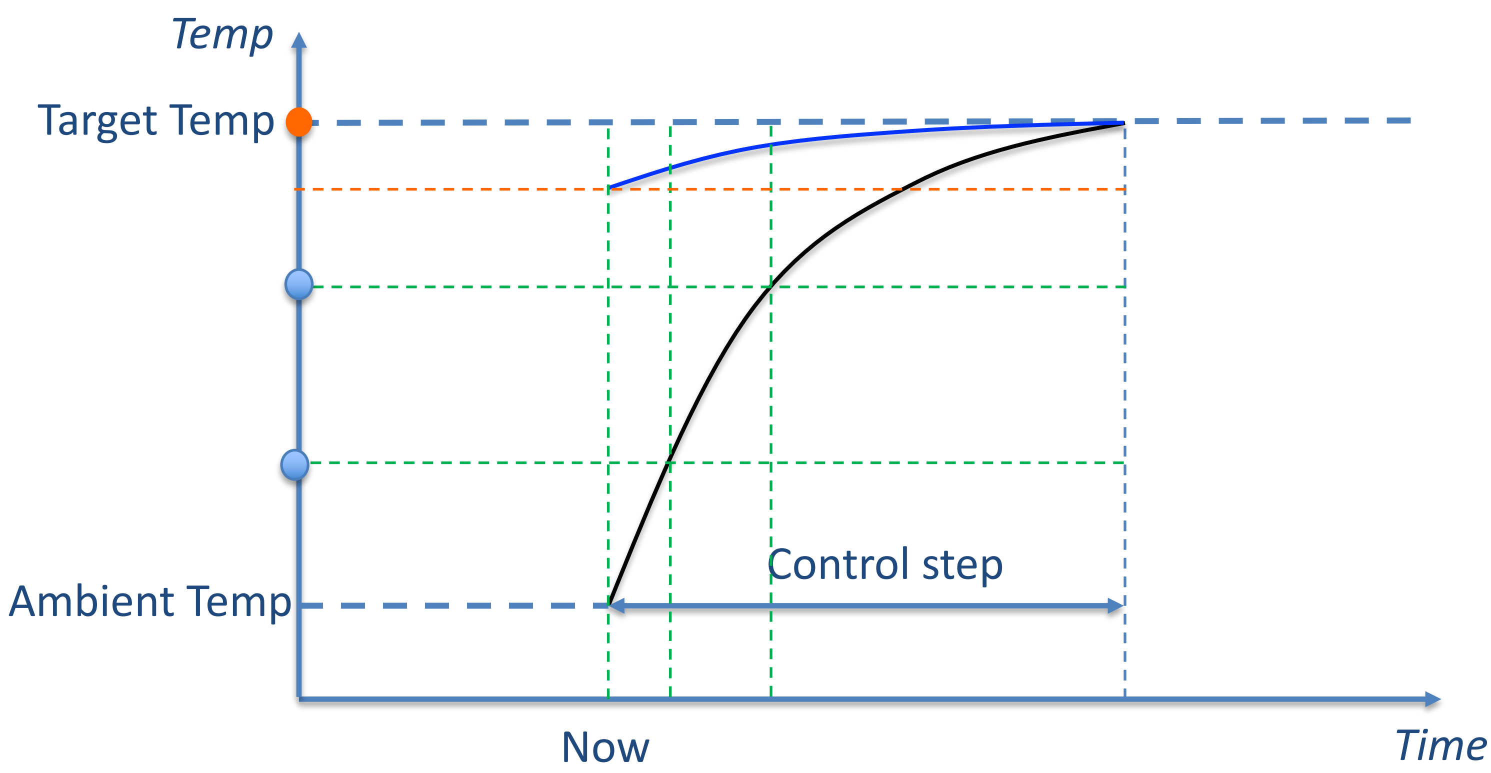

PWL thermal model formulation

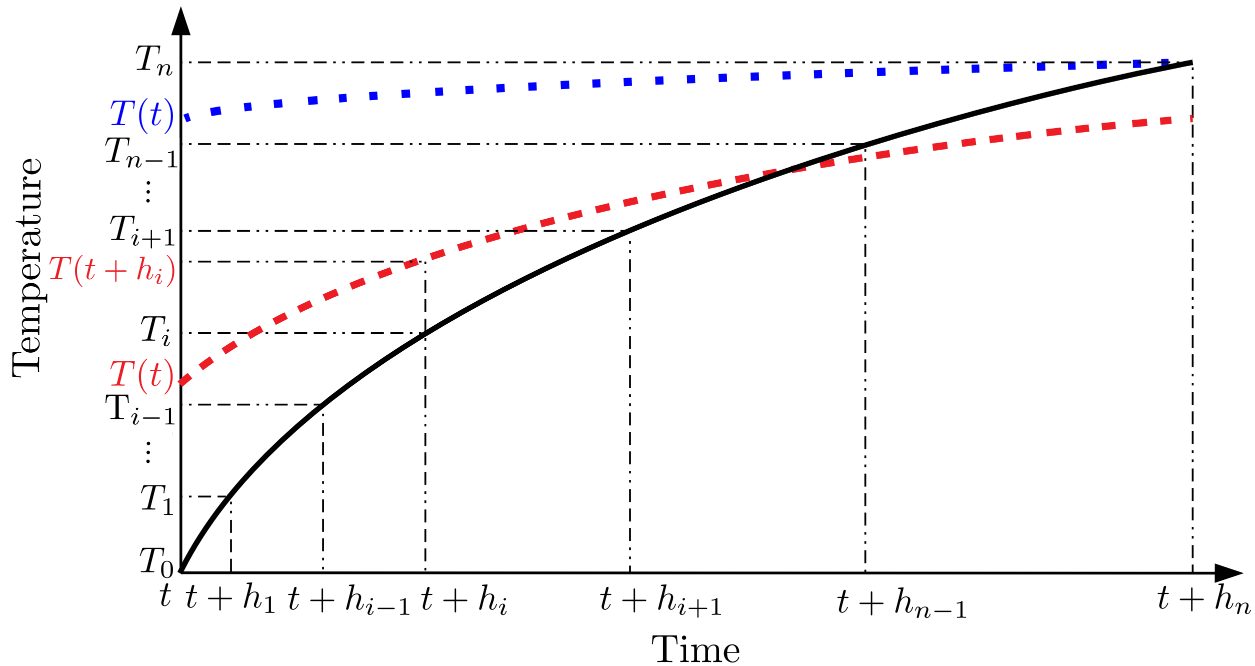

- Generally, the initial temperature $T(t)$ lies between T_{i-1} and T_i, (red dashed line).
  The first expansion point is T_i, corresponding time from t to $t+h_i$,
  and next point is T_{i+1}, corresponding time from $t+h_i$ to $t+h_{i+1}$.

- Specially, the initial temperature $T(t)$ is close to the target temperature T_n (blue dot line).
  Only one segment with target temperature T_n as the expansion point.

PWL thermal model formulation for single control step

- PWL thermal estimation trace using thermal models expanded at T_i T_{i+1}, T_{i+2}, ..., T_n

$$\begin{aligned}
T(t+h_i) &= \widehat{A}_i T(t) + \widehat{D}_i P_d + \widehat{D}_i \widehat{P}_i \\
T(t+h_{i+1}) &= \widehat{A}_{i+1} T(t+h_i) + \widehat{D}_{i+1} P_d + \widehat{D}_{i+1} \widehat{P}_{i+1}, \\
T(t+h_{i+2}) &= \widehat{A}_{i+2} T(t+h_{i+1}) + \widehat{D}_{i+2} P_d + \widehat{D}_{i+2} \widehat{P}_{i+2}, \\
&\vdots \\
T(t+h_n) &= \widehat{A}_n T(t+h_{n-1}) + \widehat{D}_n P_d + \widehat{D}_n \widehat{P}_n,
\end{aligned}$$

Combining these equations

- PWL thermal model for single control step:

$$\begin{aligned}
T(k+1) &= \widehat{\mathcal{A}} T(k) + \widehat{\mathcal{D}} P_d(k) + \widehat{\mathcal{D}}_i \widehat{P}_i + \cdots + \widehat{\mathcal{D}}_n \widehat{P}_n, \\
Y(k+1) &= L T(k+1).
\end{aligned}$$

where $\widehat{\mathcal{A}} = \widehat{A}_n \widehat{A}_{n-1} \cdots \widehat{A}_i$, $\widehat{\mathcal{D}} = \widehat{A}_n \widehat{A}_{n-1} \cdots \widehat{A}_{i+1} \widehat{D}_i + \widehat{A}_n \widehat{A}_{n-1} \cdots \widehat{A}_{i+2} \widehat{D}_{i+1} + \cdots + \widehat{D}_n$,

$\widehat{\mathcal{D}}_i = \widehat{A}_n \widehat{A}_{n-1} \cdots \widehat{A}_{i+1} \widehat{D}_i$

PWL model predictive control framework

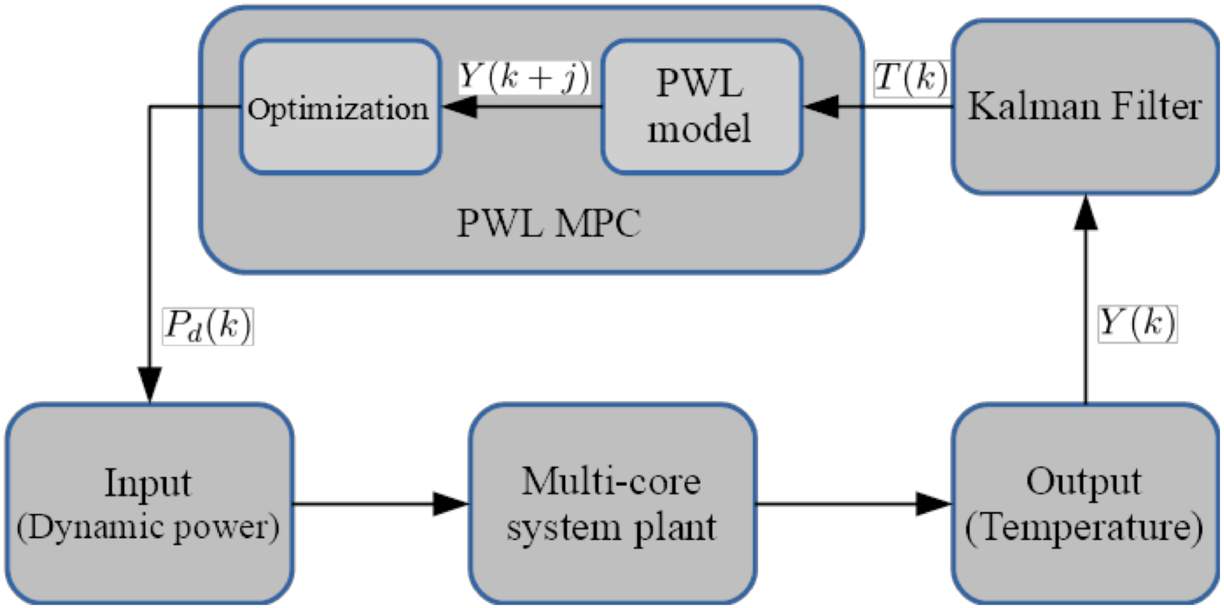

Note:
Kalman Filter is used for temperature estimation of chip and package

Predict temperature trajectory for multiple control steps

- The first control time step $k+1$ (red dashed line):

$$T(k+1) = \widehat{\mathcal{A}}T(k) + \widehat{\mathcal{D}}P_d(k) + \widehat{\mathcal{D}}_i\widehat{P}_i + \cdots + \widehat{\mathcal{D}}_n\widehat{P}_n,$$

$$Y(k+1) = LT(k+1).$$

- Assume temperature prediction is close to the target temperature after the first control step (blue dot line):

$$T(k+j) = \widehat{A}_nT(k+j-1) + \widehat{D}_nP_d(k+j-1) + \widehat{D}_n\widehat{P}_n,$$

$$Y(k+j) = LT(k+j),$$

Predicted temperature trajectory for multiple control steps

$$\mathcal{Y} = FT(k) + V\mathcal{P}_d + \phi_1\widehat{\mathcal{P}} + \phi_2\widehat{\mathcal{P}}_n$$

$$\mathcal{Y} = [Y(k+1)^T, Y(k+2)^T, \ldots, Y(k+N_p)^T]^T$$

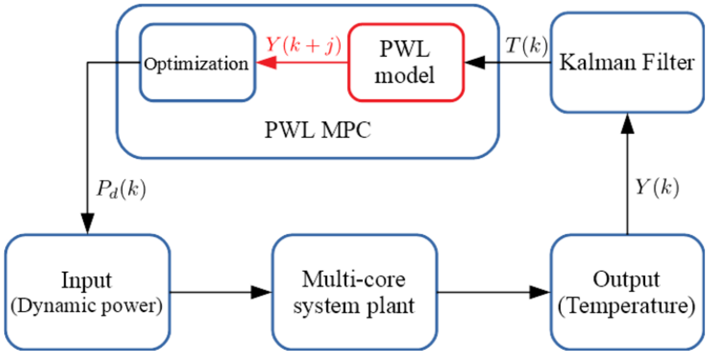

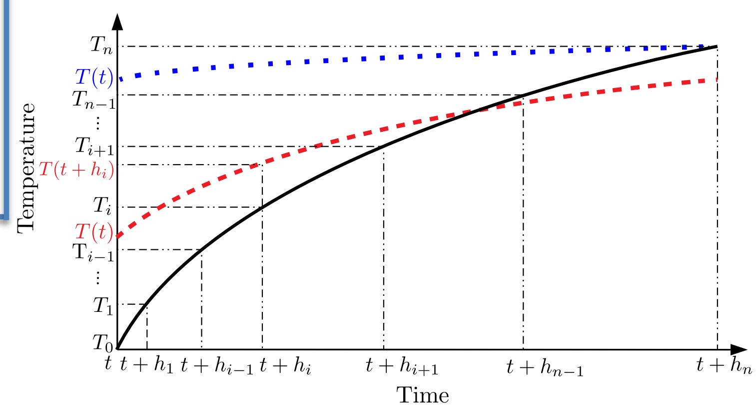

Compute power recommendation

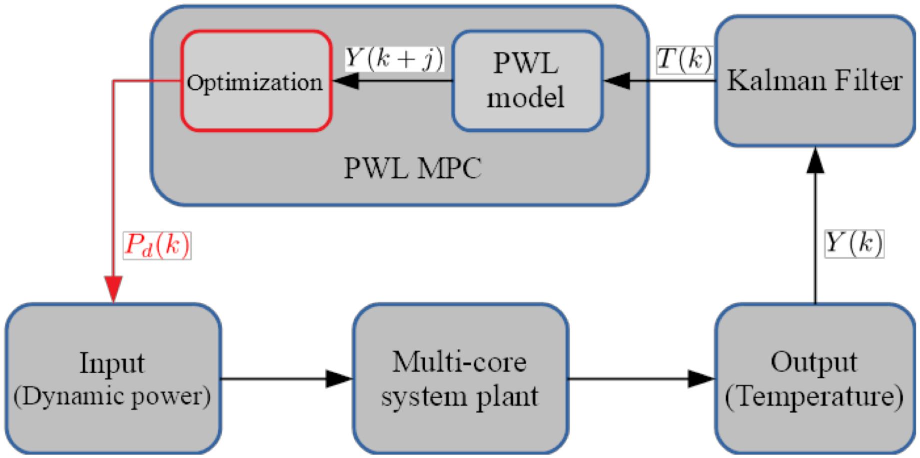

- Minimize regulated cost function: $\mathcal{J} = (\mathcal{Y}_g - \mathcal{Y})^T(\mathcal{Y}_g - \mathcal{Y}) + \mathcal{P}_d^T R \mathcal{P}_d$

➡ Future dynamic Power recommendation $\mathcal{P}_d = (V^T V + R)^{-1} V^T (\mathcal{Y}_g - FT(k) - \phi_1 \widehat{\mathcal{P}} - \phi_2 \widehat{\mathcal{P}}_n)$

The frequencies and task loads will be adjusted according to $P_d(k)$

Outline

- Background
- Problems of leakage-aware DTM
- PWL predictive DTM
- Experimental results
- Conclusion

Experimental setup

- Experiment on a 16-core system
- One thermal sensor for each core
- Ambient temperature $40\,°\mathrm{C}$.
- Target temperature $80\,°\mathrm{C}$.
- Compared with linear model predictive control (called traditional DTM).

C11	C12	C13	C14
C21	C22	C23	C24
C31	C32	C33	C34
C41	C42	C43	C44

Performance comparison

Methods	$N_c = 1, N_p = 1$				$N_c = 1, N_p = 2$				$N_c = 1, N_p = 3$				$N_c = 2, N_p = 3$			
	time (ms)	mem (KB)	difference		time (ms)	mem (KB)	difference		time (ms)	mem (KB)	difference		time (ms)	mem (KB)	difference	
			max	avg			max	avg			max	avg			max	avg
Traditional	1.01	7	6.02	1.35	1.13	13	5.97	1.32	1.22	21	5.86	1.34	1.30	23	5.94	1.37
New (3 points)	1.12	14	1.25	0.84	1.28	23	1.20	0.81	1.39	42	1.21	0.82	1.46	44	1.23	0.85
New (5 points)	1.12	22	1.21	0.79	1.28	36	1.16	0.78	1.39	65	1.18	0.80	1.47	67	1.20	0.82
New (7 points)	1.13	32	1.15	0.77	1.29	51	1.12	0.75	1.40	91	1.09	0.76	1.47	94	1.13	0.79
New (11 points)	1.14	58	1.08	0.75	1.30	90	1.05	0.72	1.42	150	1.11	0.74	1.49	154	1.09	0.75

- Tracking difference: much smaller
- Computing time: similar to traditional
- Memory cost: increases linearly with expansion point number

Hint: balance accuracy and memory cost with proper expansion point number

Control quality comparison

- 5 expansion points
- Both DTM methods activated at 1s, target temperature $80\,^{\circ}\mathrm{C}$.

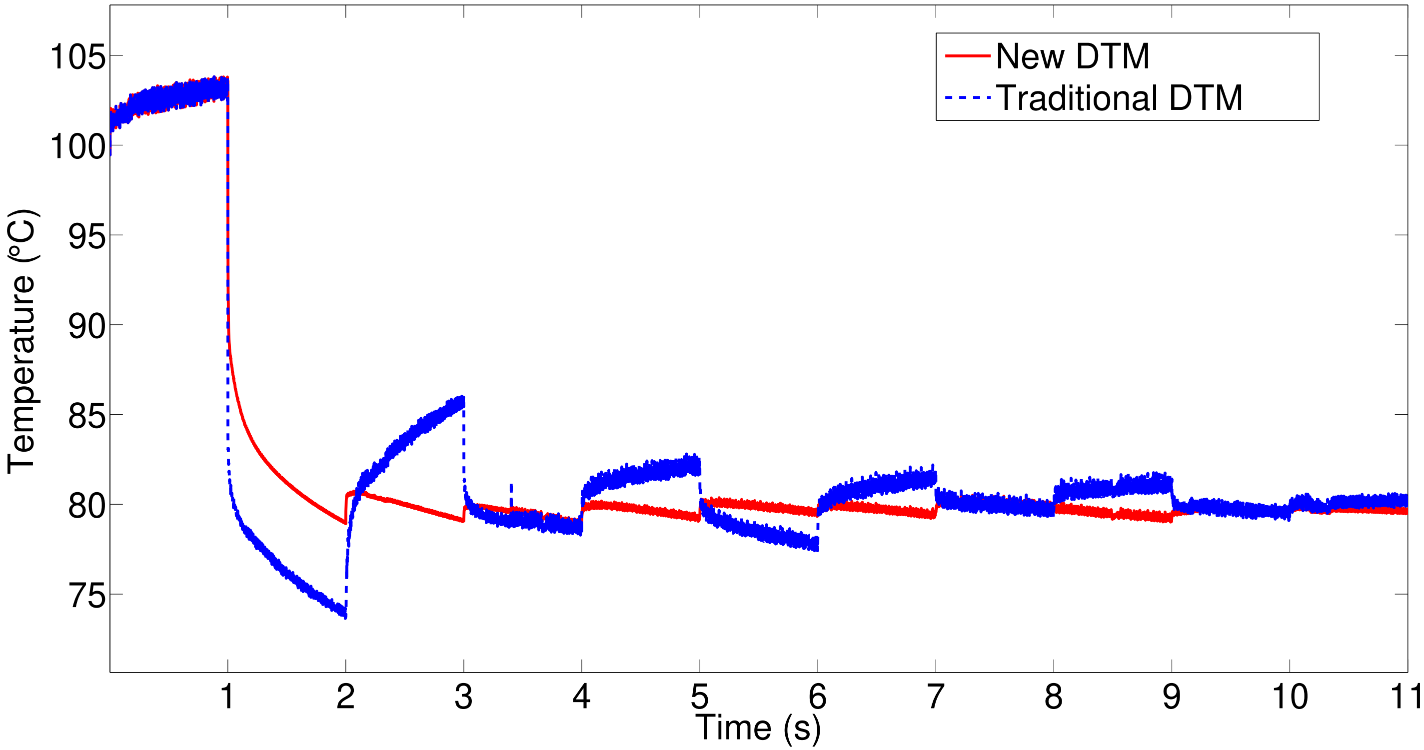

- Traditional DTM shows large tracking overshot.
- New method has extremely smooth control with little overshot.

Conclusion

- A PWL model predictive control based DTM method is proposed.
- The PWL thermal model is concise and can be integrated into the predictive control elegantly.
- A systematic expansion point selection scheme is developed for PWL models.
- We show how to integrate the new PWL thermal model into the predictive control framework.
- The new method outperforms traditional method in control quality.

Thank you!