Automated Lookahead Data Migration in SSD-enabled Multi-tiered Storage Systems

Gong Zhang, Ling Liu

Georgia Institute of Technology

Lawrence Chiu, Clem Dickey, Paul Muench

IBM Almaden Research Center

Multi-tiered Storage Architecture

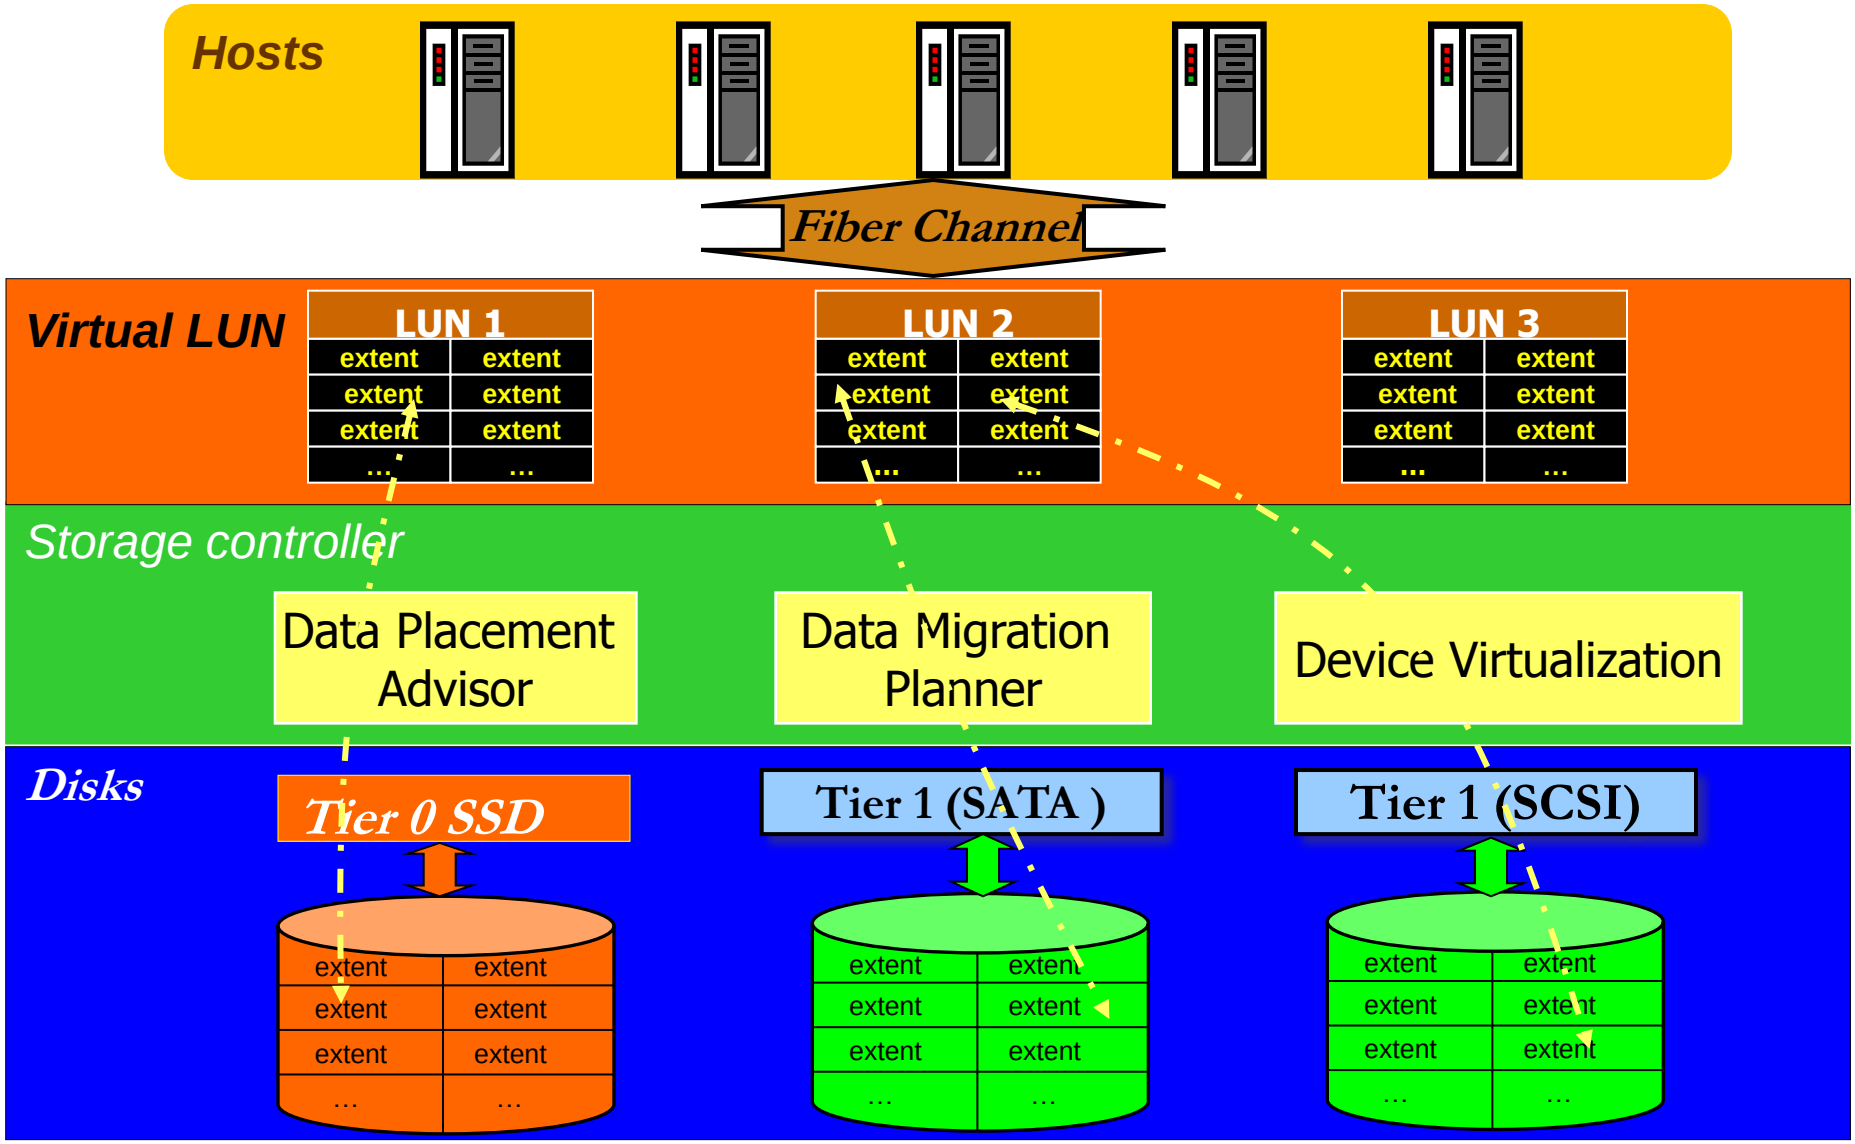

SPC-1 Heat Map

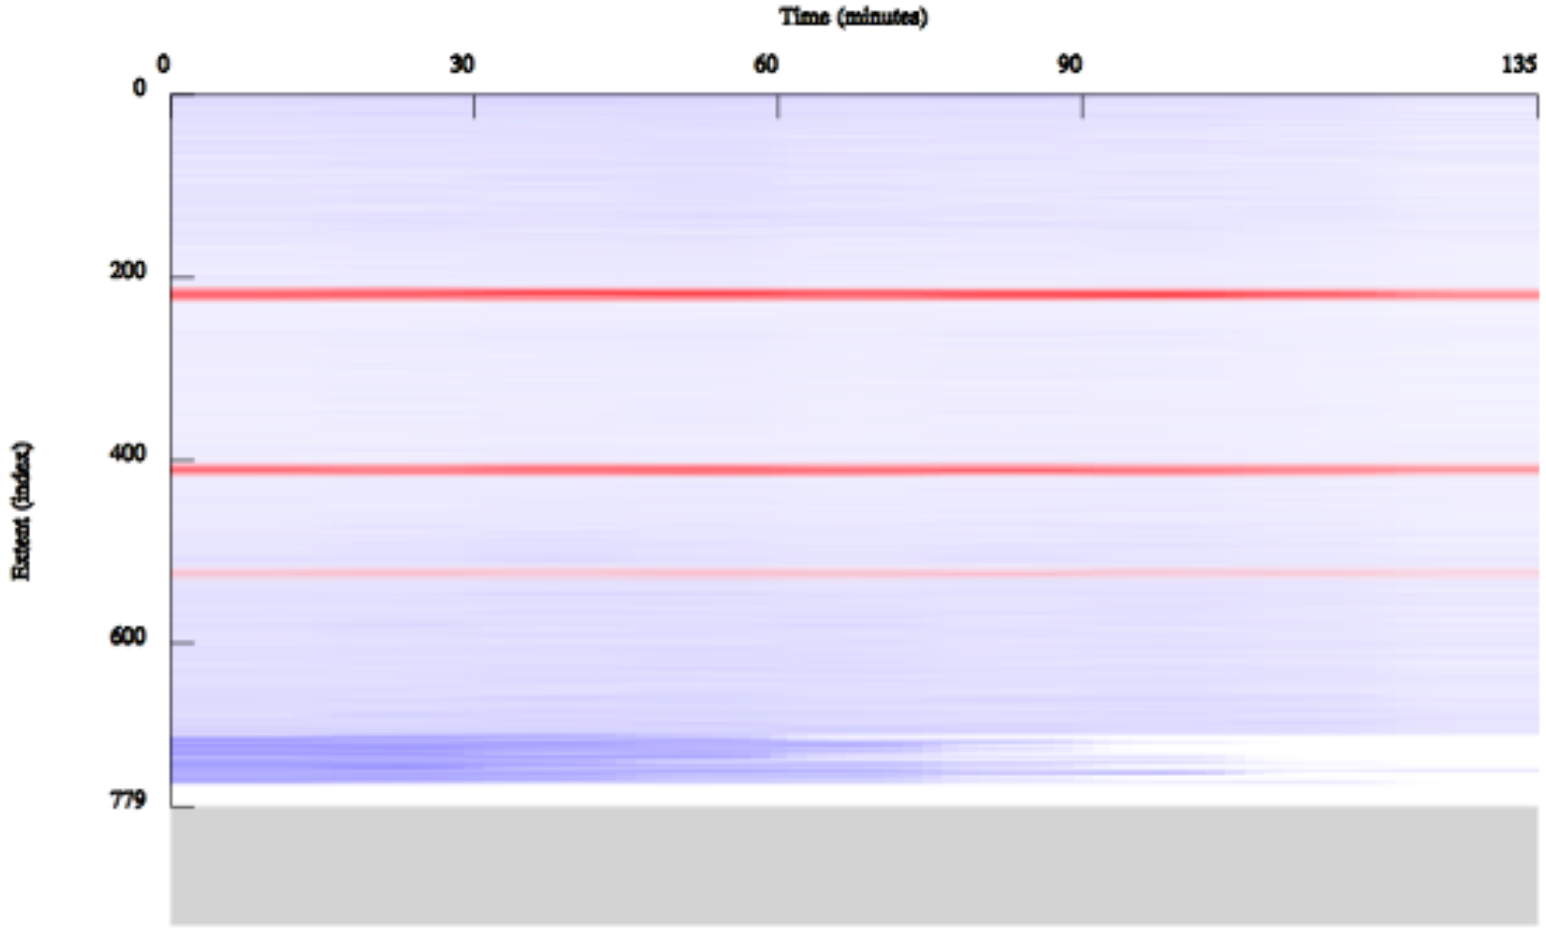

TPC-E Heat Map

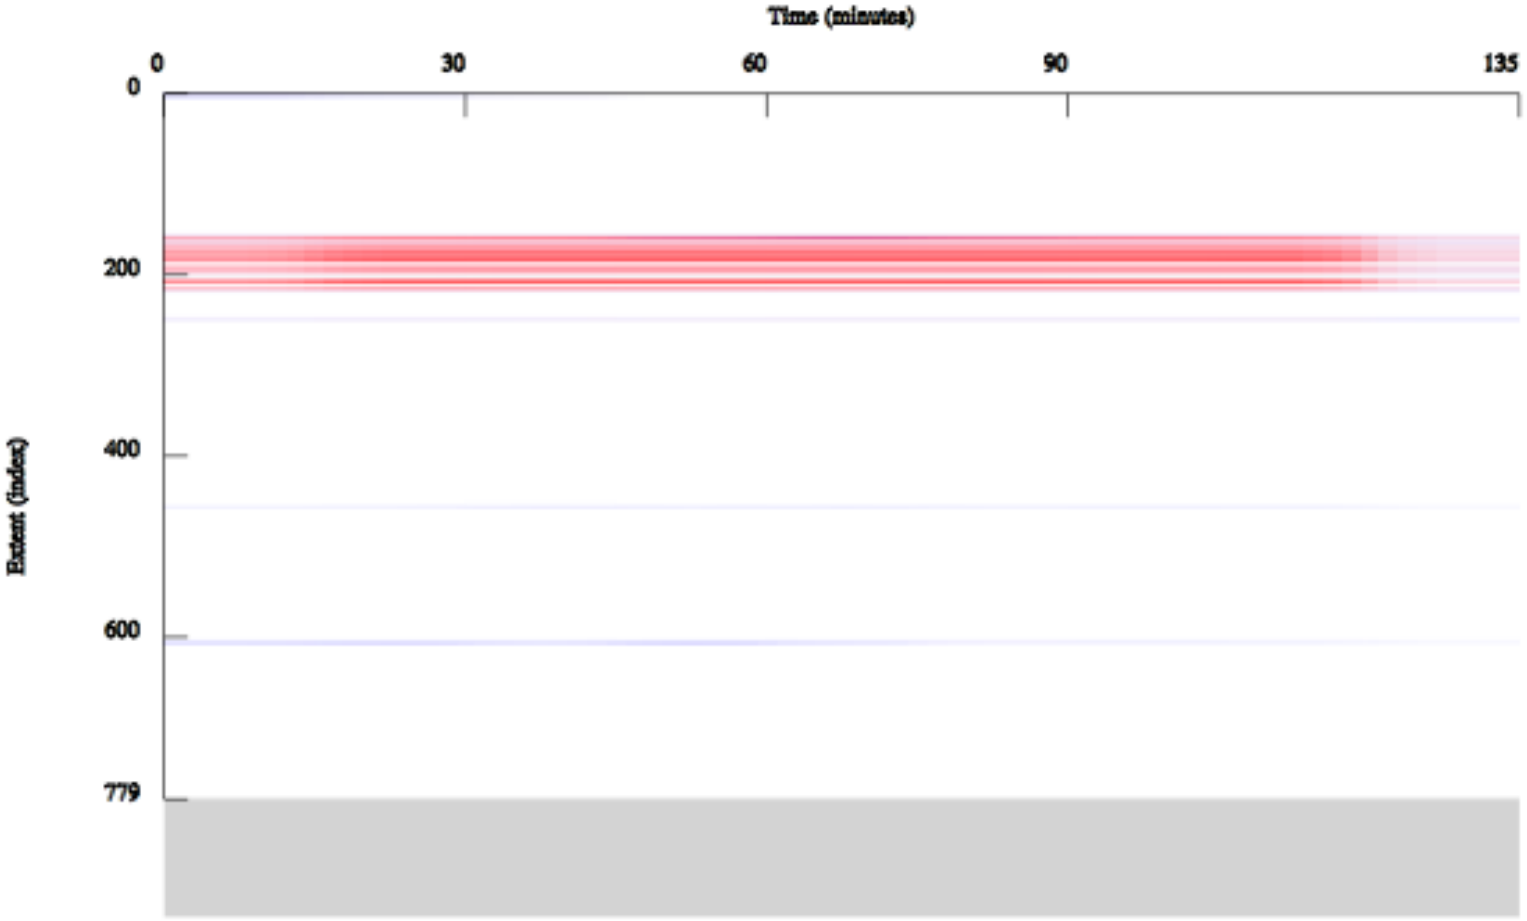

Data Migration

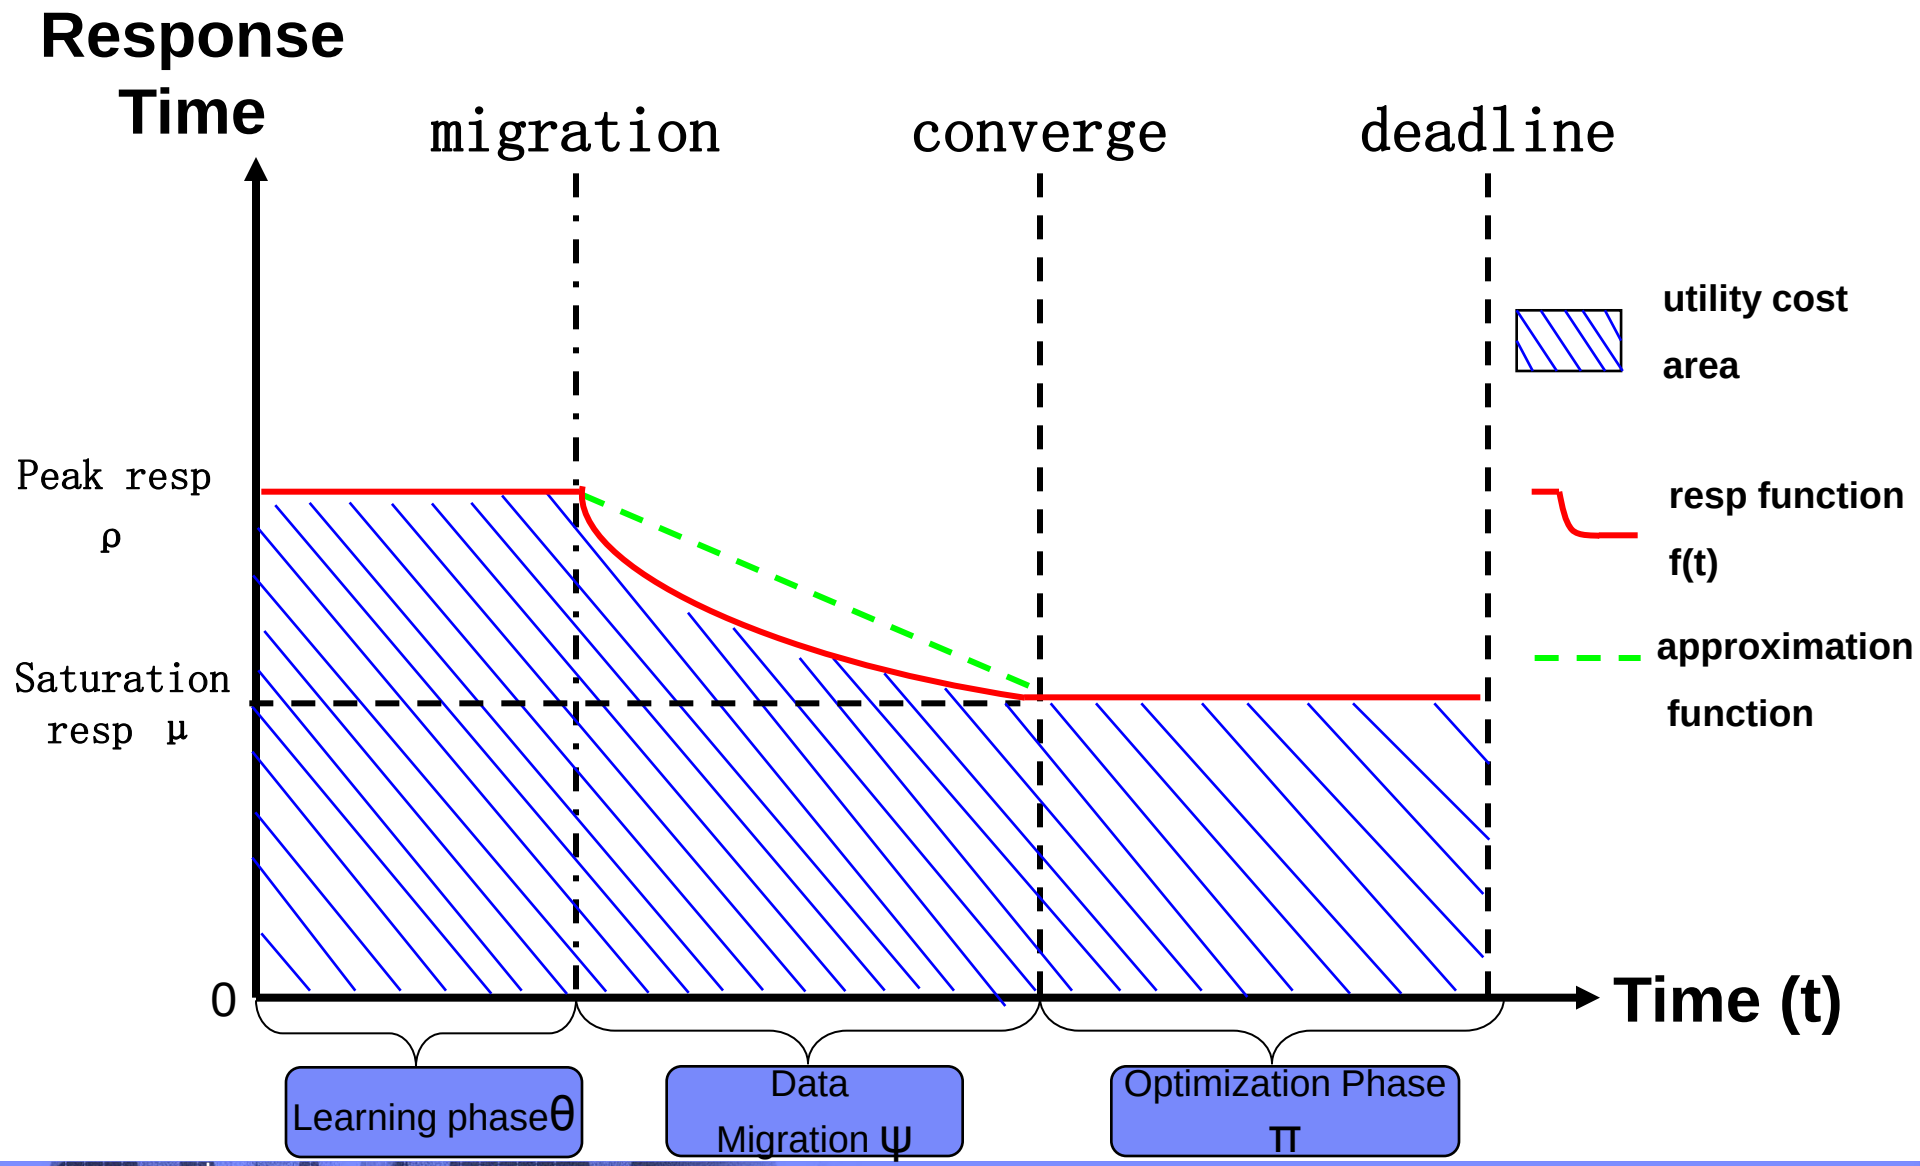

Data Migration

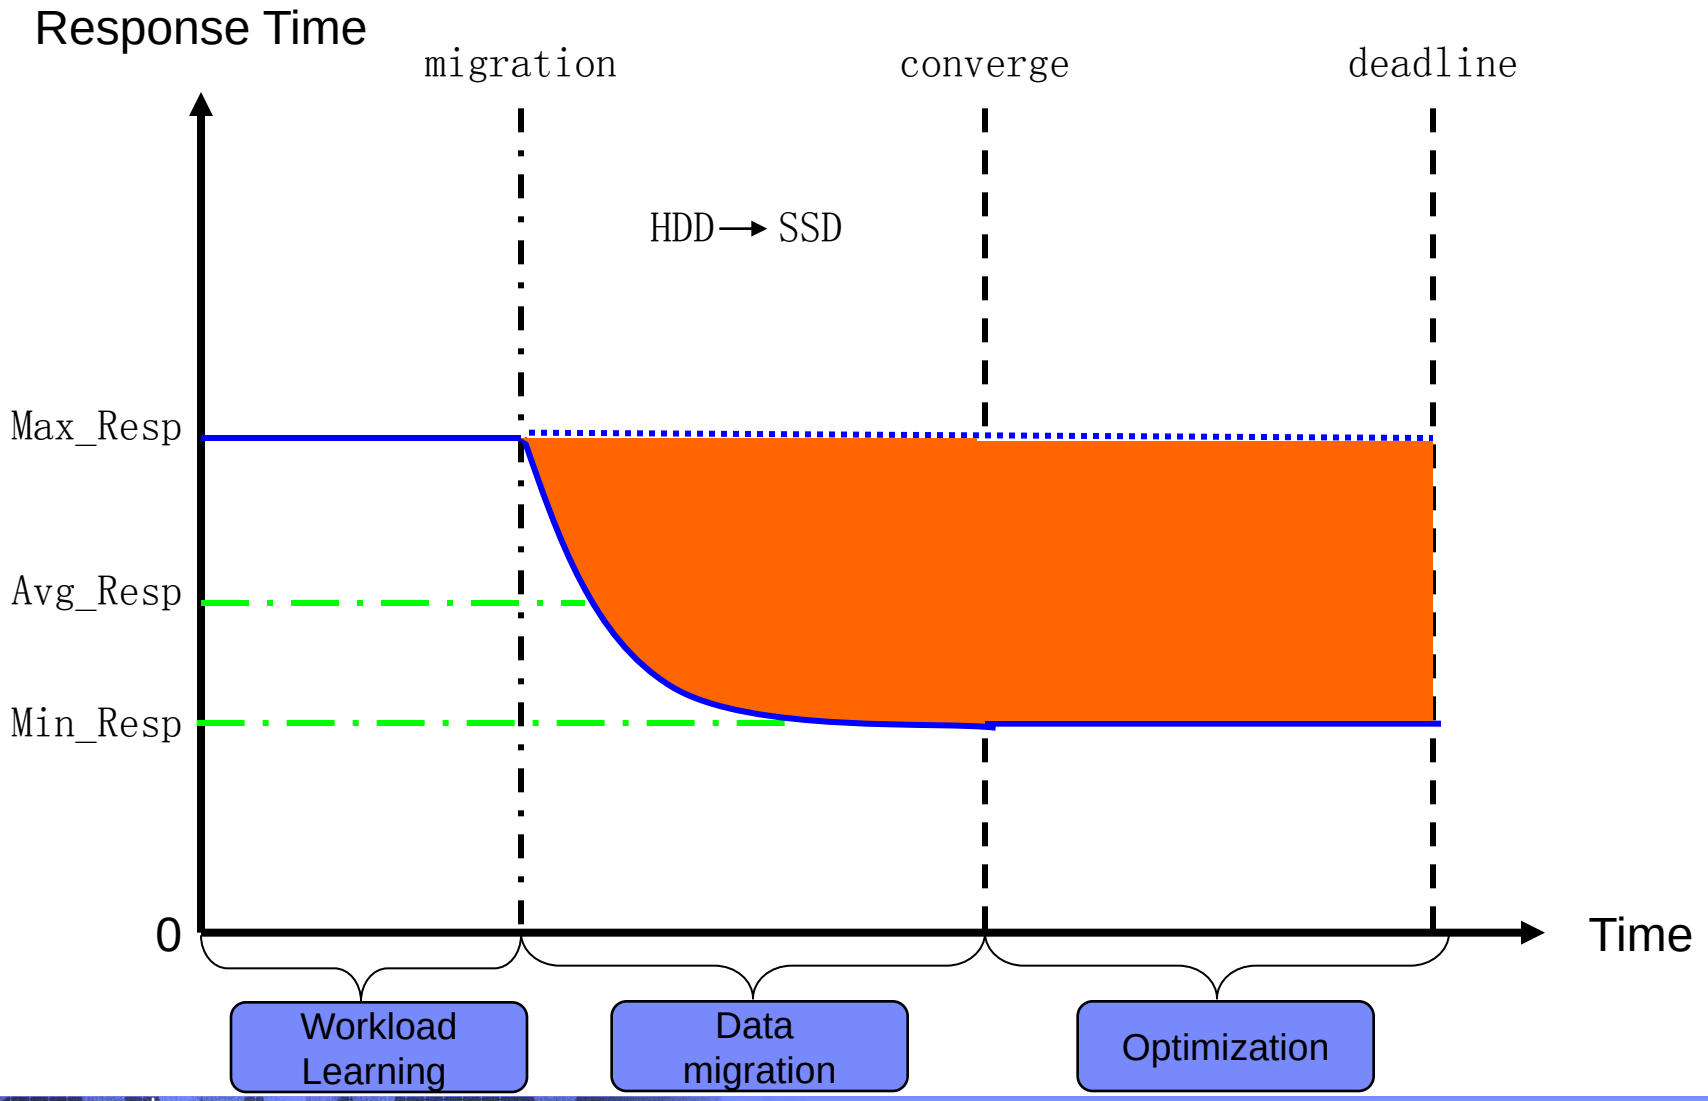

Optimization 1: Reducing Learning Phase

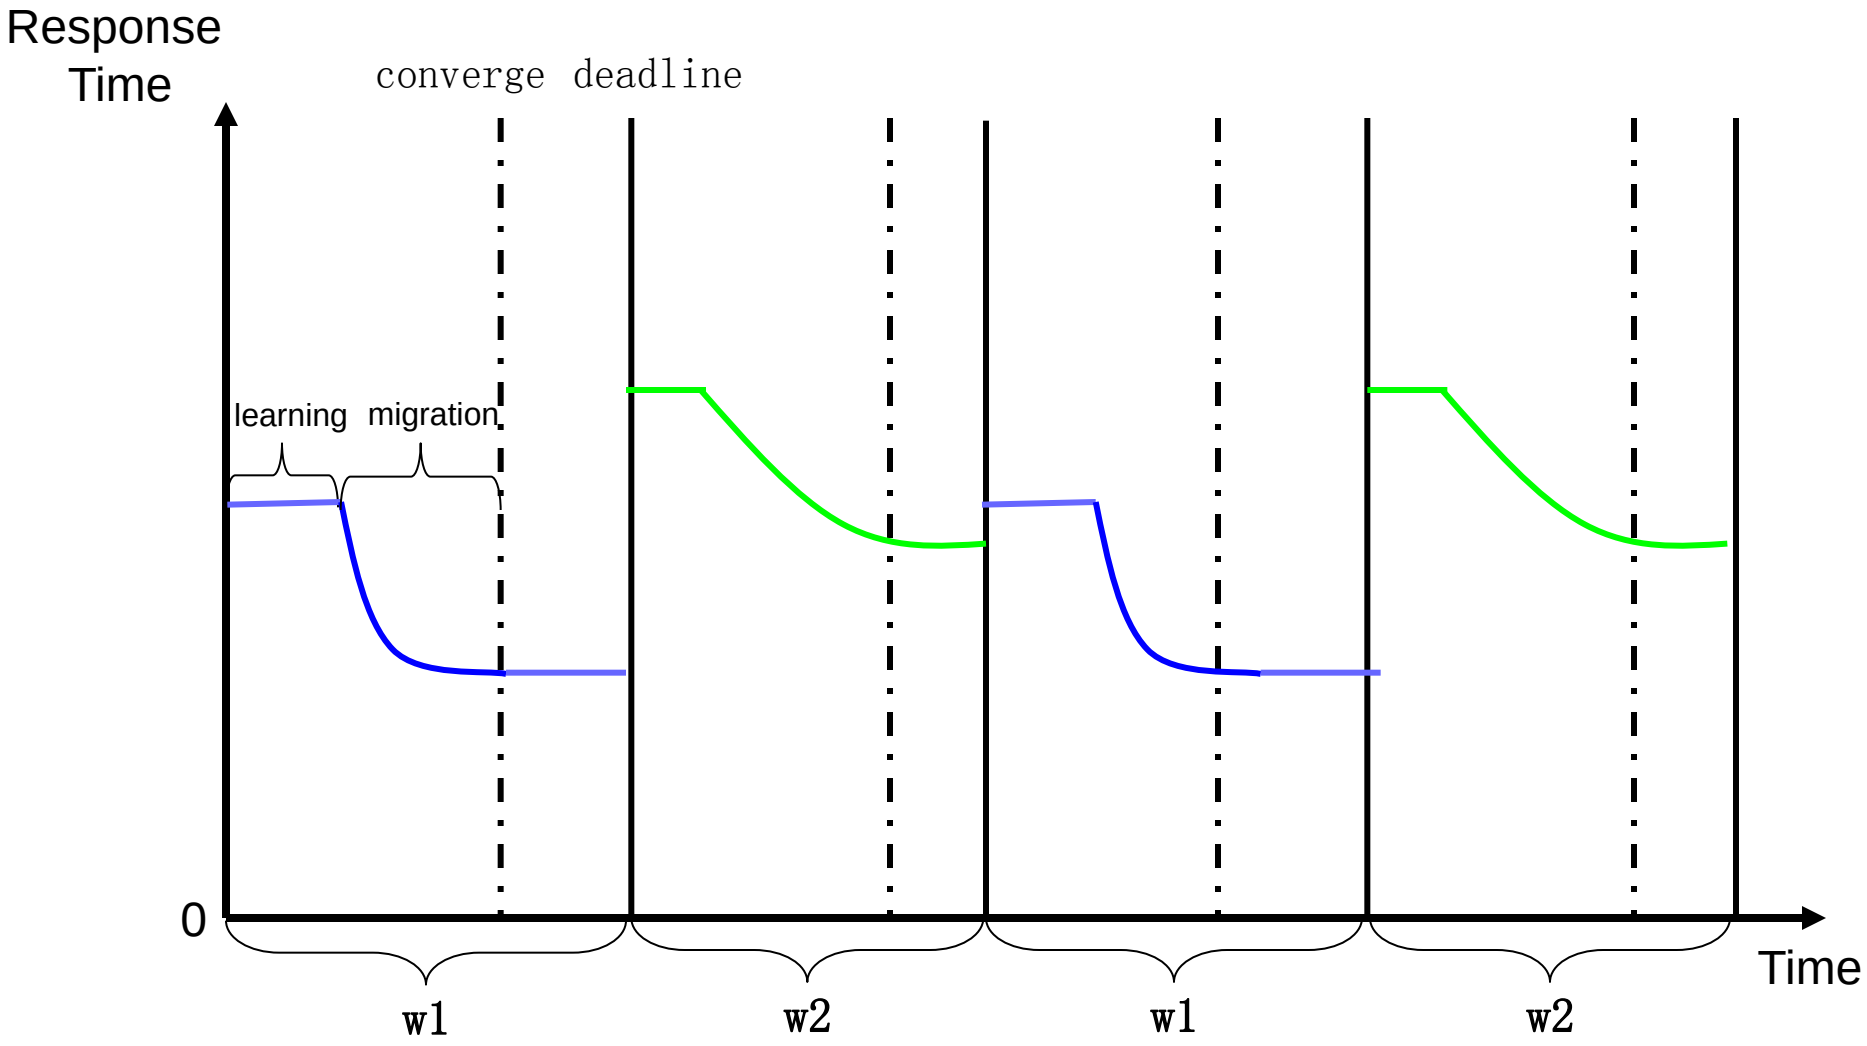

Optimization 1: Reducing Learning Phase

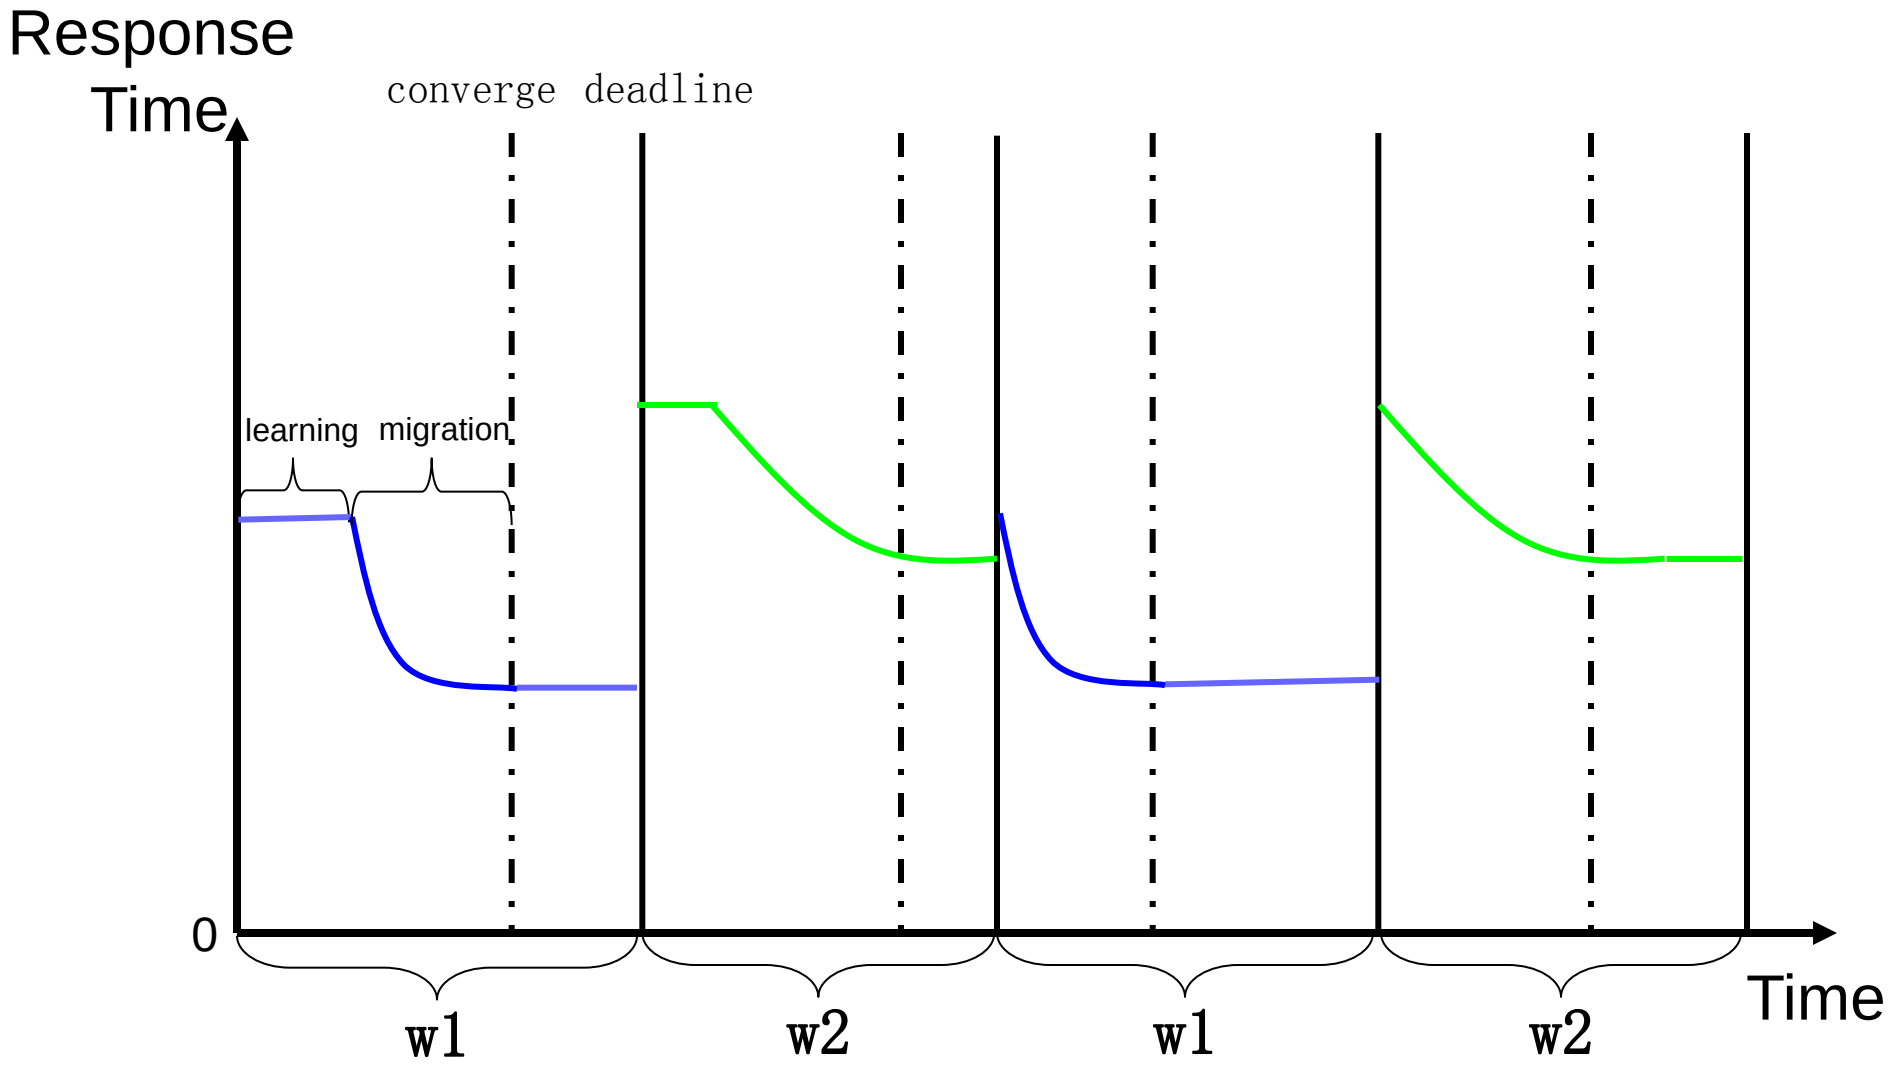

Optimization 2 :Lookahead Migration

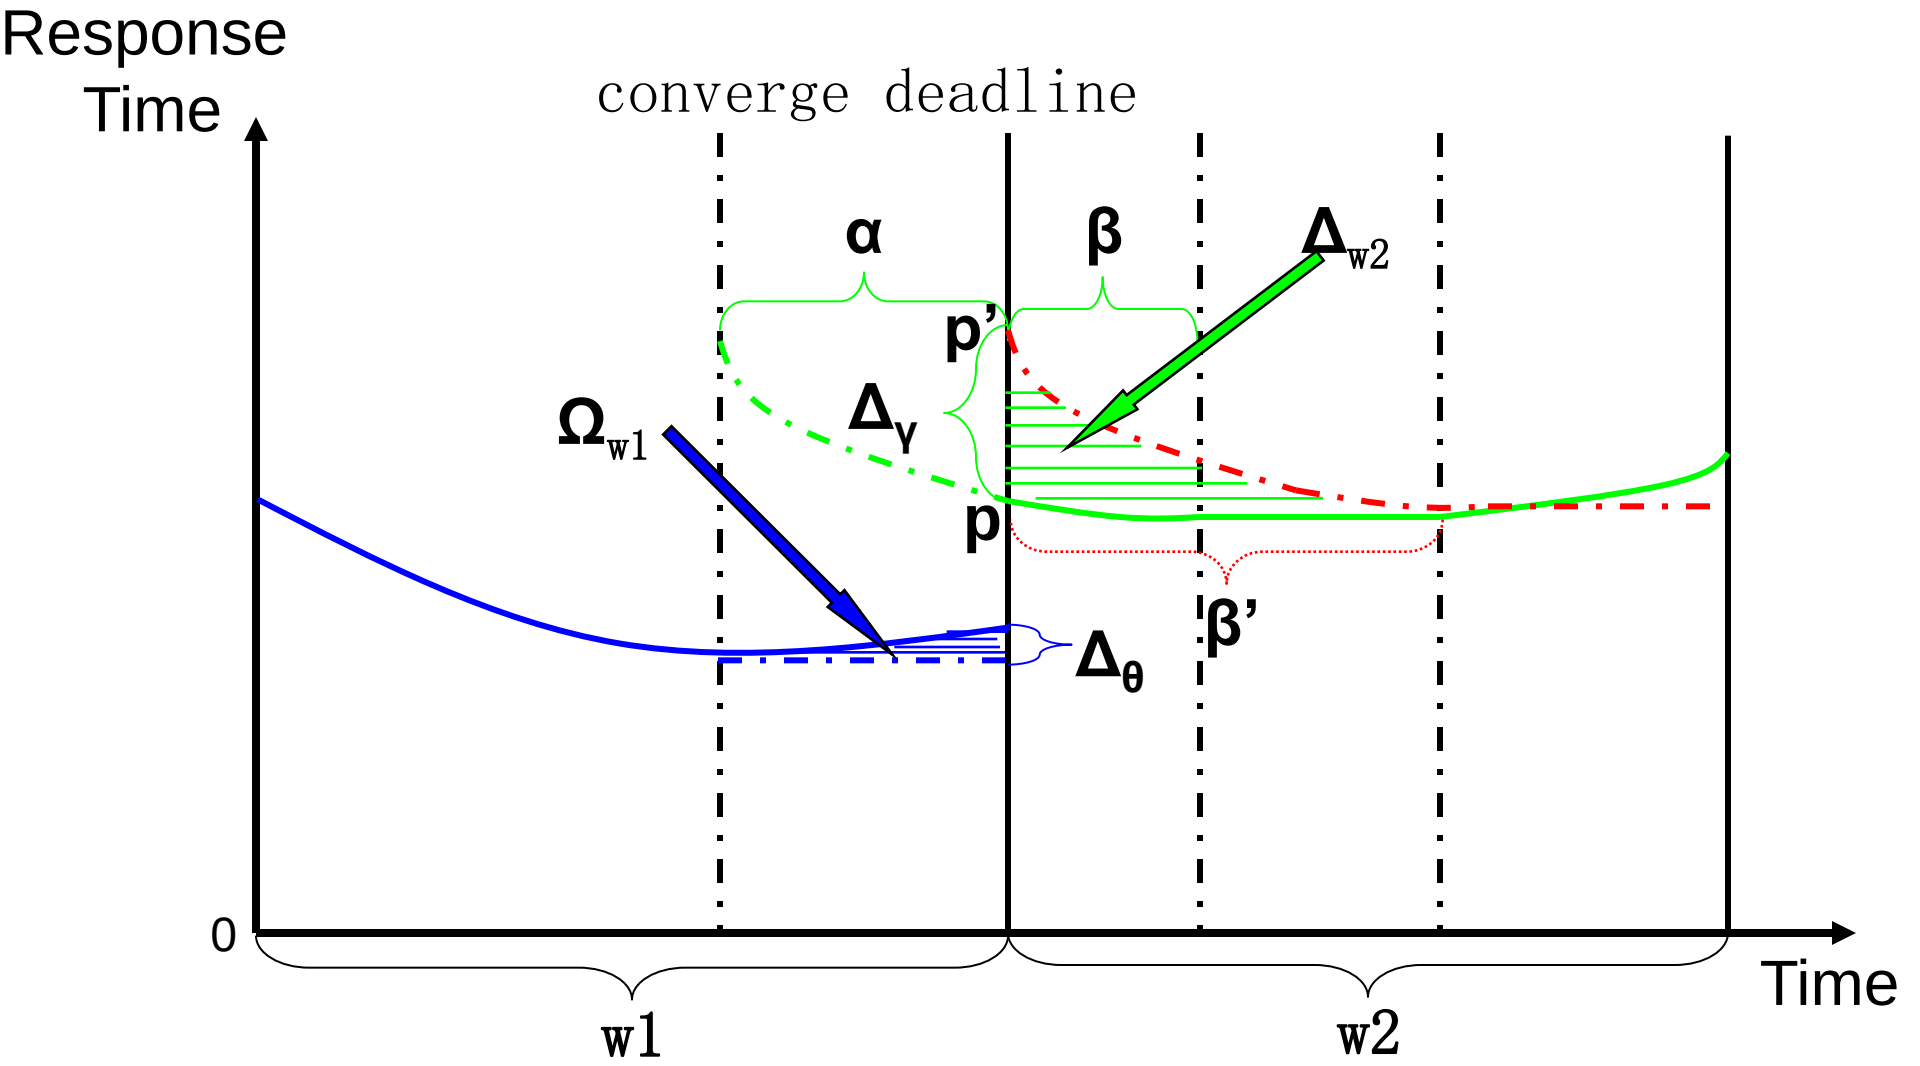

Optimization 3: Adaptive Lookahead Data Migration:

- **Constant lookahead length is not optimal**
- **IO density distribution change**
  - IO density change drives the changes of hot extents in terms of quantity and heat distribution
- **New lookahead length is demanded**
  - Hot extent changes drives the adaptation of lookahead length
- **Greedy algorithm to compute near optimal lookahead length in an approximation manner**
- **New lookahead length is computed based on the learned IO performance profile and constraints**

Greedy Algorithm: max(diff(Δw2-Ωw1)) near optimal lookahead length

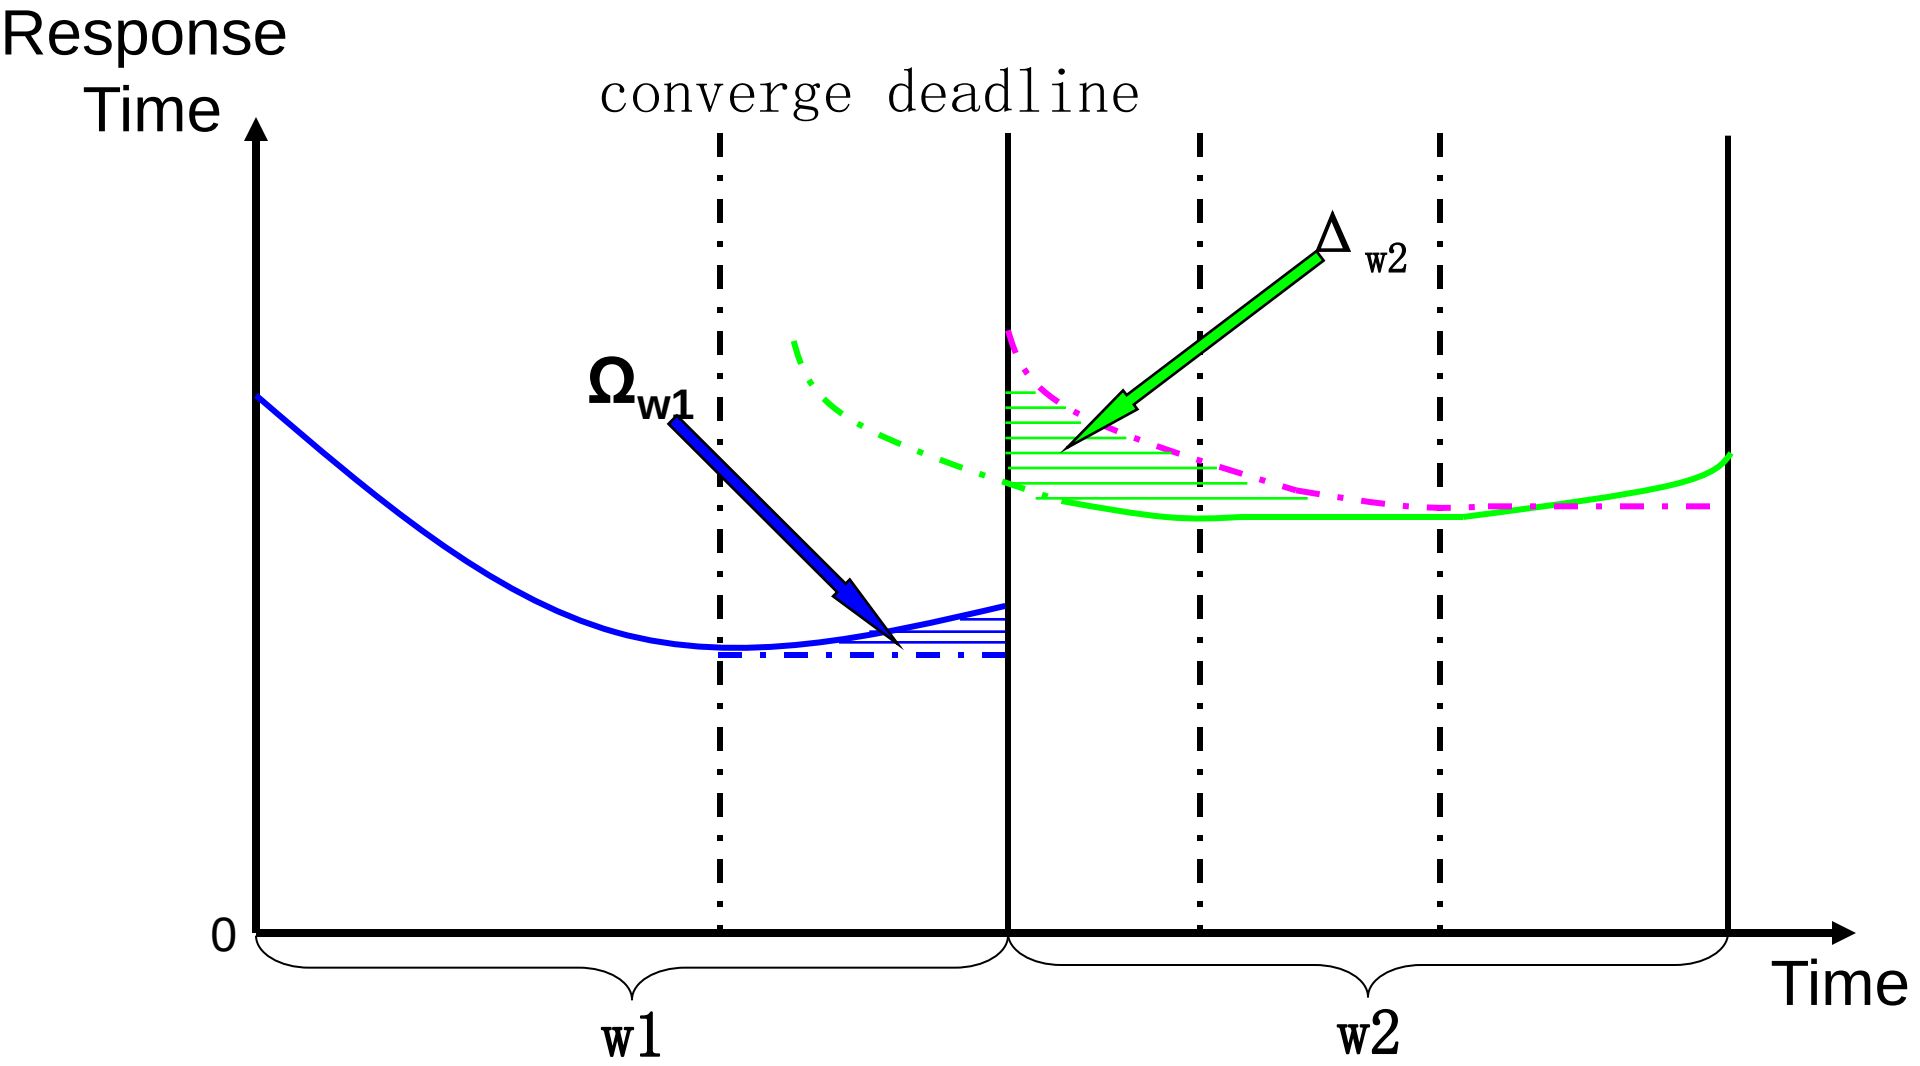

Adaptive Lookahead Migration Computation

- **Greedy algorithm is limited by the increment granularity**
- **Adaptive lookahead data migration: computing optimal lookahead length through bandwidth, SSD Size, etc.**

$$\begin{aligned}\Gamma(\alpha) &= \Delta_{w2} - \Omega_{w1} \\ &= (\rho_2 - \mu_2)\alpha - \frac{\rho_2 - \mu_2}{2 \times \psi_2}\alpha^2 - \frac{\rho_1 - \mu_1}{2 \times \psi_1}\alpha^2 \\ &= (\rho_2 - \mu_2)\alpha - (\frac{\rho_2 - \mu_2}{2 \times \psi_2} + \frac{\rho_1 - \mu_1}{2 \times \psi_1})\alpha^2\end{aligned}$$

$$\alpha = \frac{\psi_1\psi_2(\rho_2-\mu_2)}{\phi_1(\rho_2-\mu_2)+\phi_2(\rho_1-\mu_1)}$$

$$\Gamma_{max} = \frac{3\phi_1\phi_2(\rho_2-\mu_2)^2}{2\psi_1(\rho_2-\mu_2)+2\psi_2(\rho_1-\mu_1)}.$$

Response time reduction with learning phase learning phase reduced migration

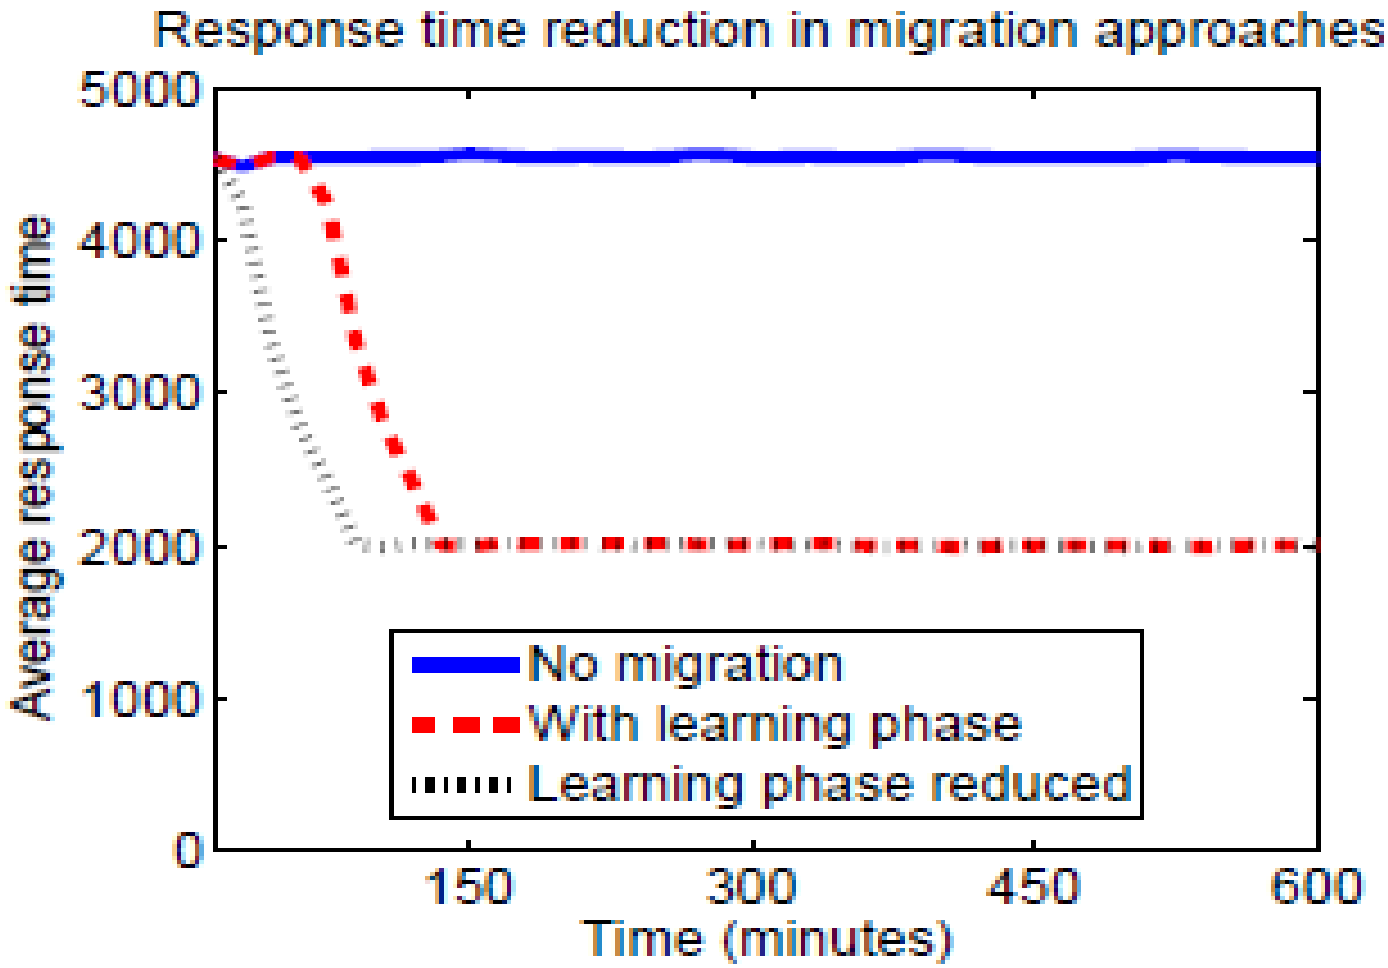

Lookahead length on response time reduction

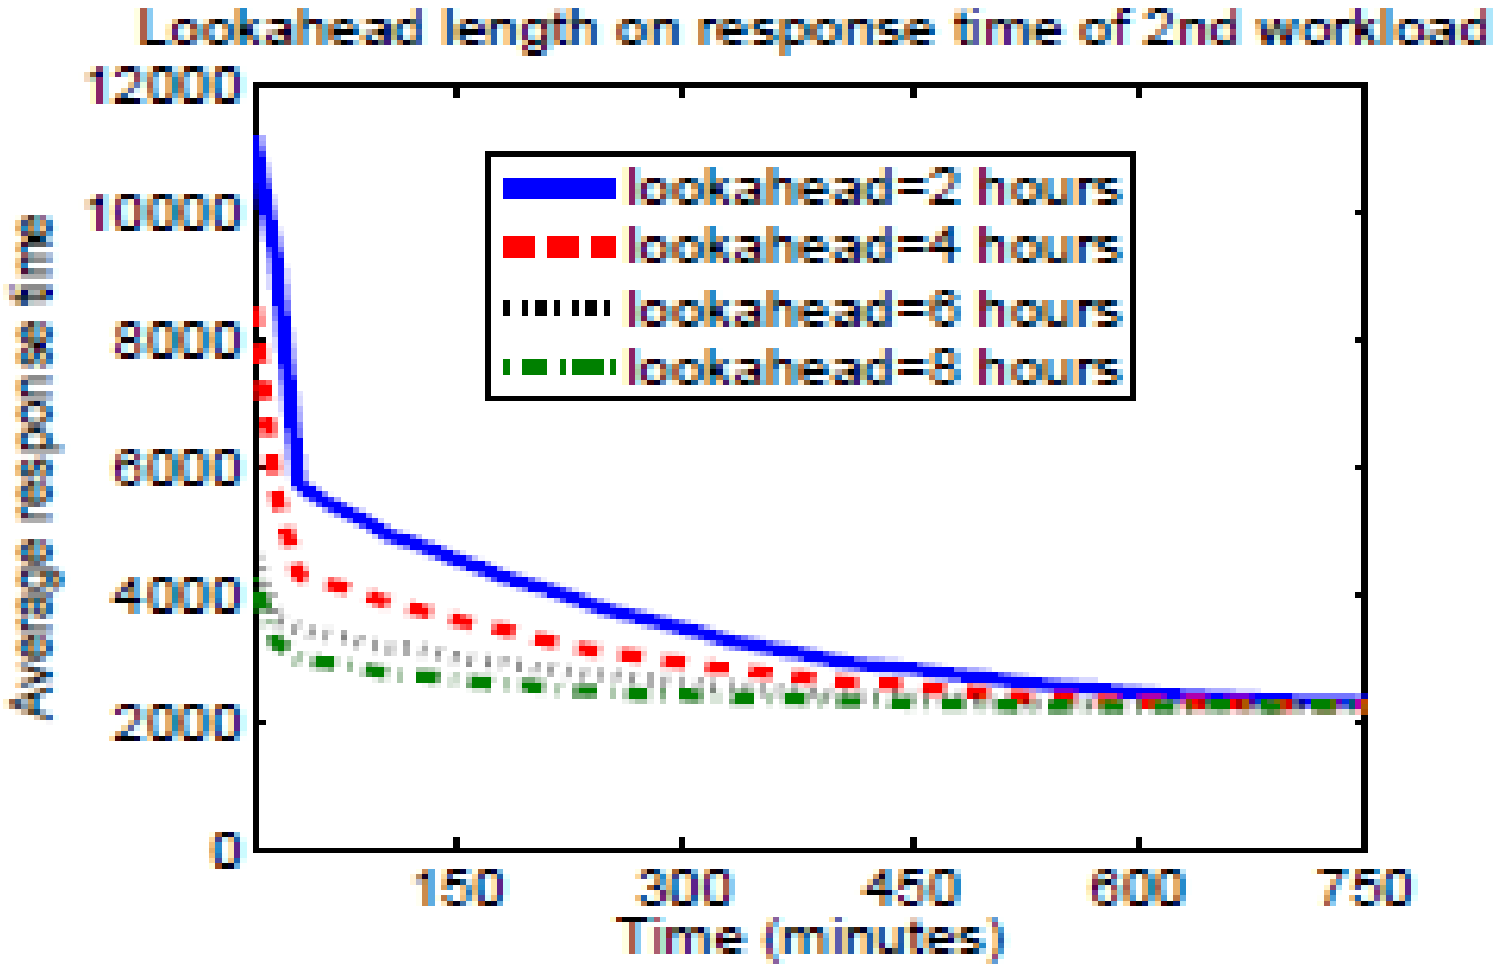

Compute optimal lookahead length

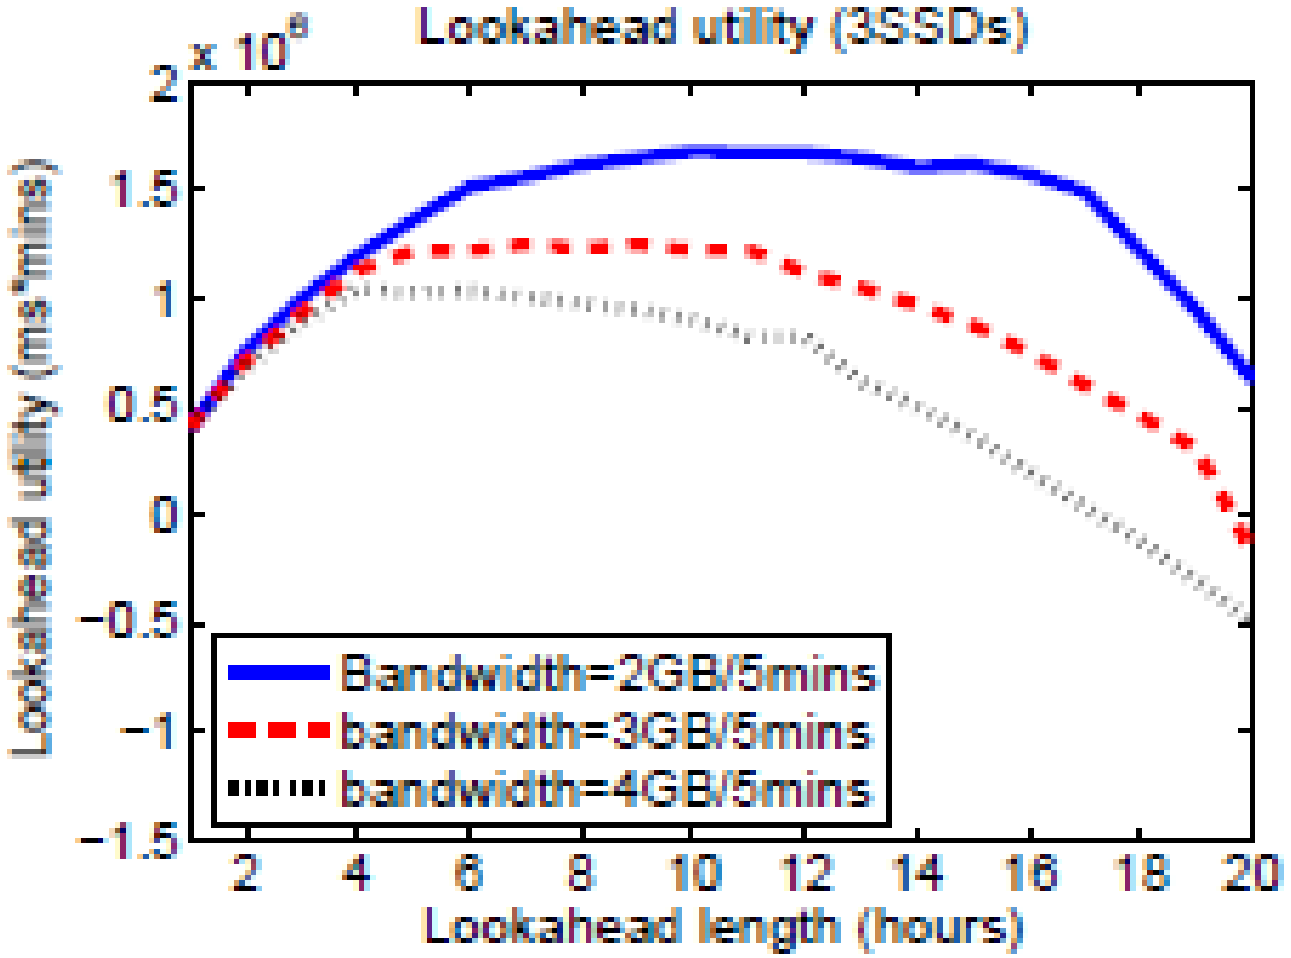

Questions and Thanks

?

SPC-1 Average Response Time is only Reduced with Concurrent Migration

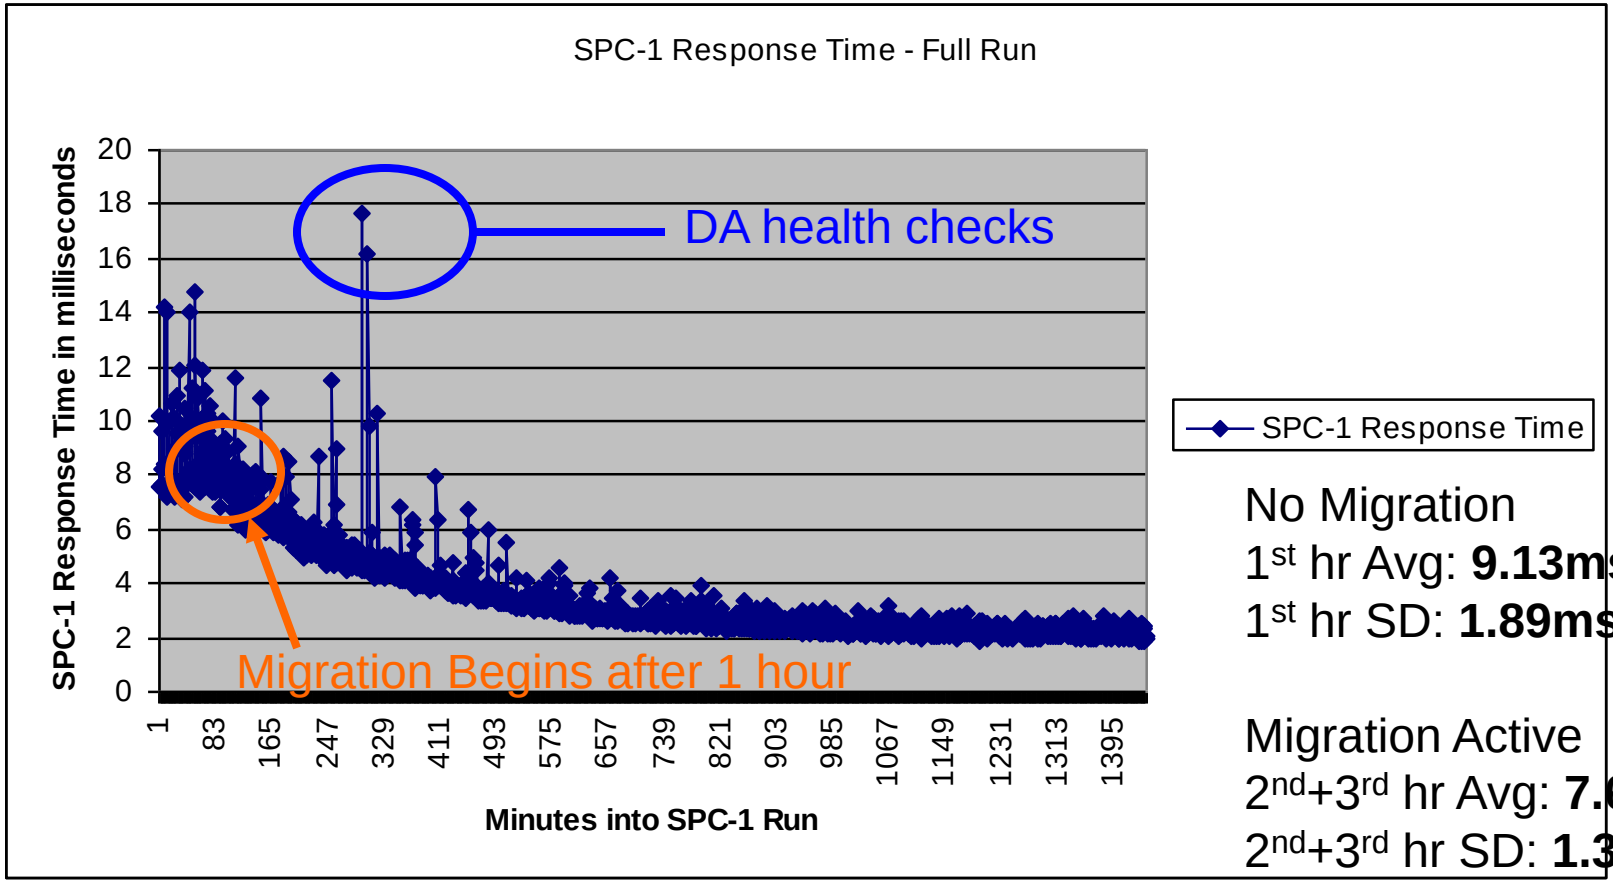

No Migration
1st hr Avg: **9.13ms**
1st hr SD: **1.89ms**

Migration Active
2nd+3rd hr Avg: **7.61ms**
2nd+3rd hr SD: **1.31ms**

What type of workloads are good matching workloads?

- **Sensitivity to migration**
  - Stable heat pattern or dynamic heat patttern?
  - Heat distribution among all extents
  - Hot extents ratio
- **Workload cycle duration**
  - Exception: migration not finished before deadline if too short workload time and limited bandwidth

The impacts of SSD size

- **The ratio between SSD and HDD**
- **The ratio is too small**
  - SSD only holds the hottest extent and creates limited migration impacts
- **The ratio is too large**
  - Extra unnecessary cold extents migration
  - Bandwidth waste
- **Convergence point**
  - Time point differentiate necessary migration and unnecessary migration

How to guarantee the migration deadline

- **Migration deadline must be guaranteed**
  - SSD size
    - Saturated
    - unsaturated
  - The data volume to be migrated
    - Impacting extents: extents migrated before convergence
    - Different on different workload
  - The allocated bandwidth for migration
- **Migration deadline must be guaranteed**
  - Unfinished migration
    - Deadline is not ensured
  - Migration finished earlier
    - SSD is full
    - SSD is not full and a mixture of different workloads

Impacts of migration on response time

- **Migration reduces response time**
- **Approximation function**
  - Linear approximation function
  - Non-linear approximation function
- **Convergence process**

Lookahead migration

- **The factors impacting lookahead migration**
  - Peak response time
  - Saturation response time
  - Migration time length
- Measurement
  - **Utility cost**

Future Work and Conclusion

- **Future Work**
  - Multiple workloads (>2 workloads)
  - Parallel workloads
  - Smarter IO density monitoring
  - Workload/IO density prediction
- **Conclusion**
  - Lookahead Migration further reduces response time and improves system resource utilization

Deadline Aware Data Migration in Multi-tiered Storage Systems

- **User scenario**
  - Daytime workload: OLTP
  - Night time workload: Batch processing
  - SSD is scarce resource
- **Constraints**
  - Repeating periodical workload cycle
  - Finish migration by deadline
  - Relatively Stable IO profile
  - Diminishing migration return
- **Challenge**
  - How to improve data migration performance gains in migration deadline bounded multi-tier storage environment?
    - Response time
    - Resource utilization

Optimization 1: Reducing Learning Phase

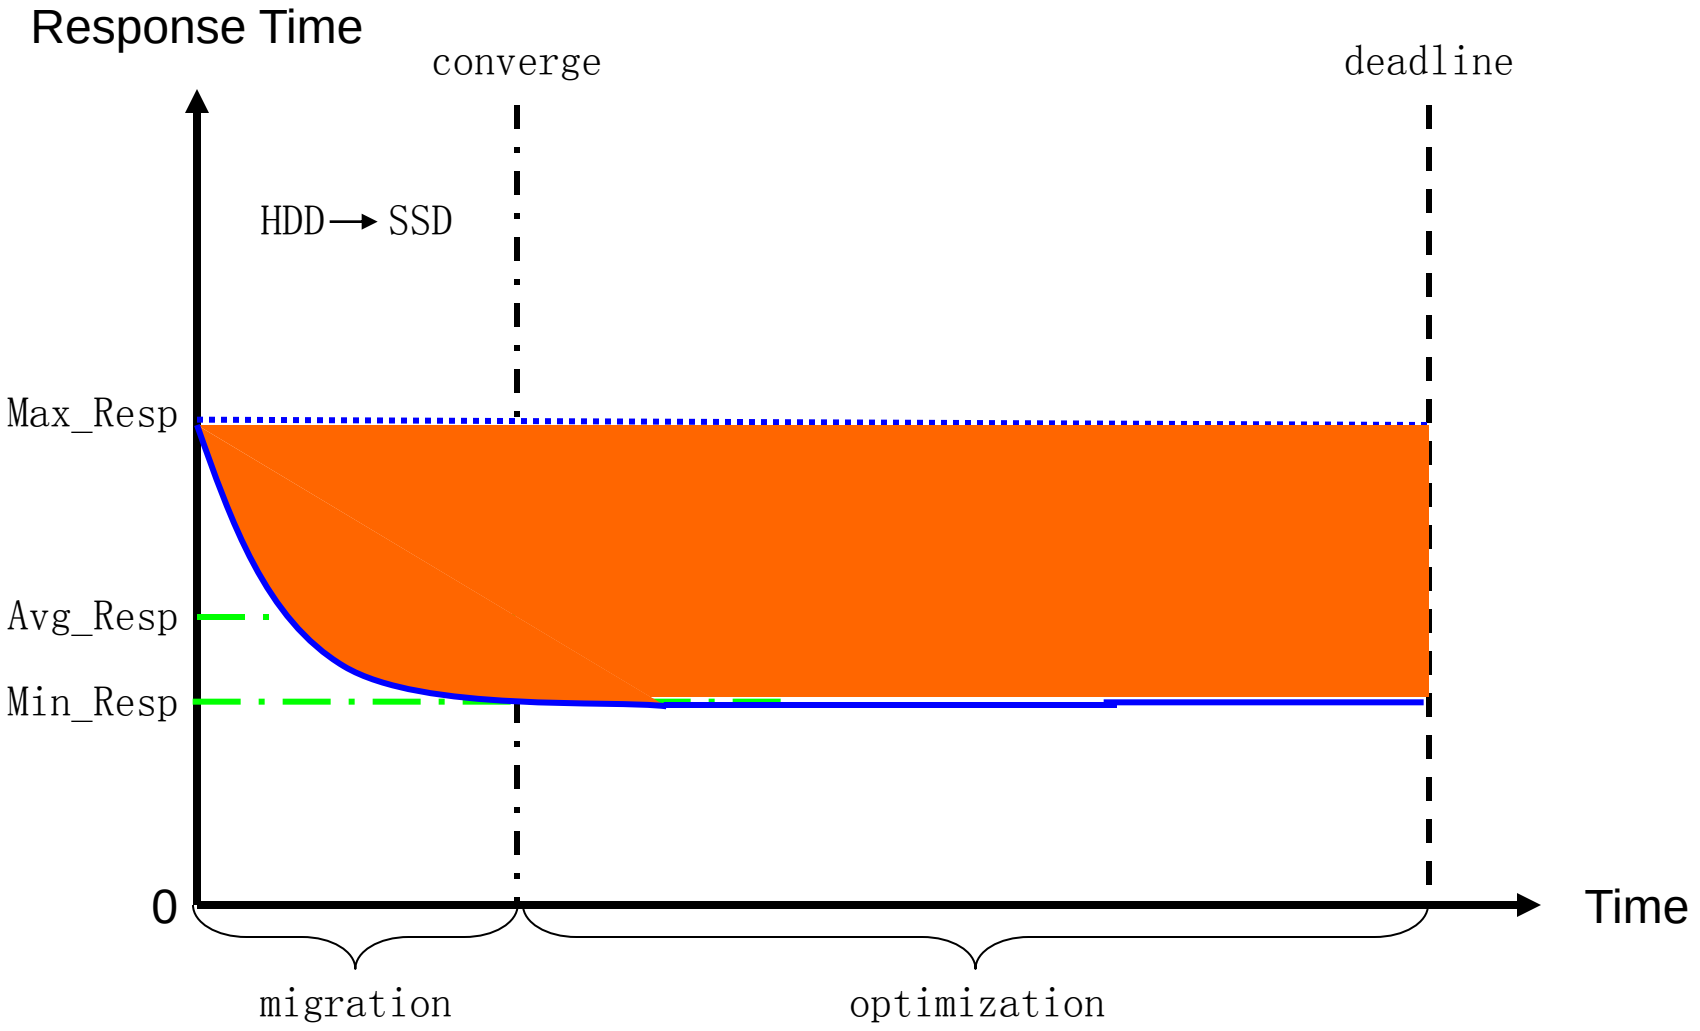

Static Data Migration

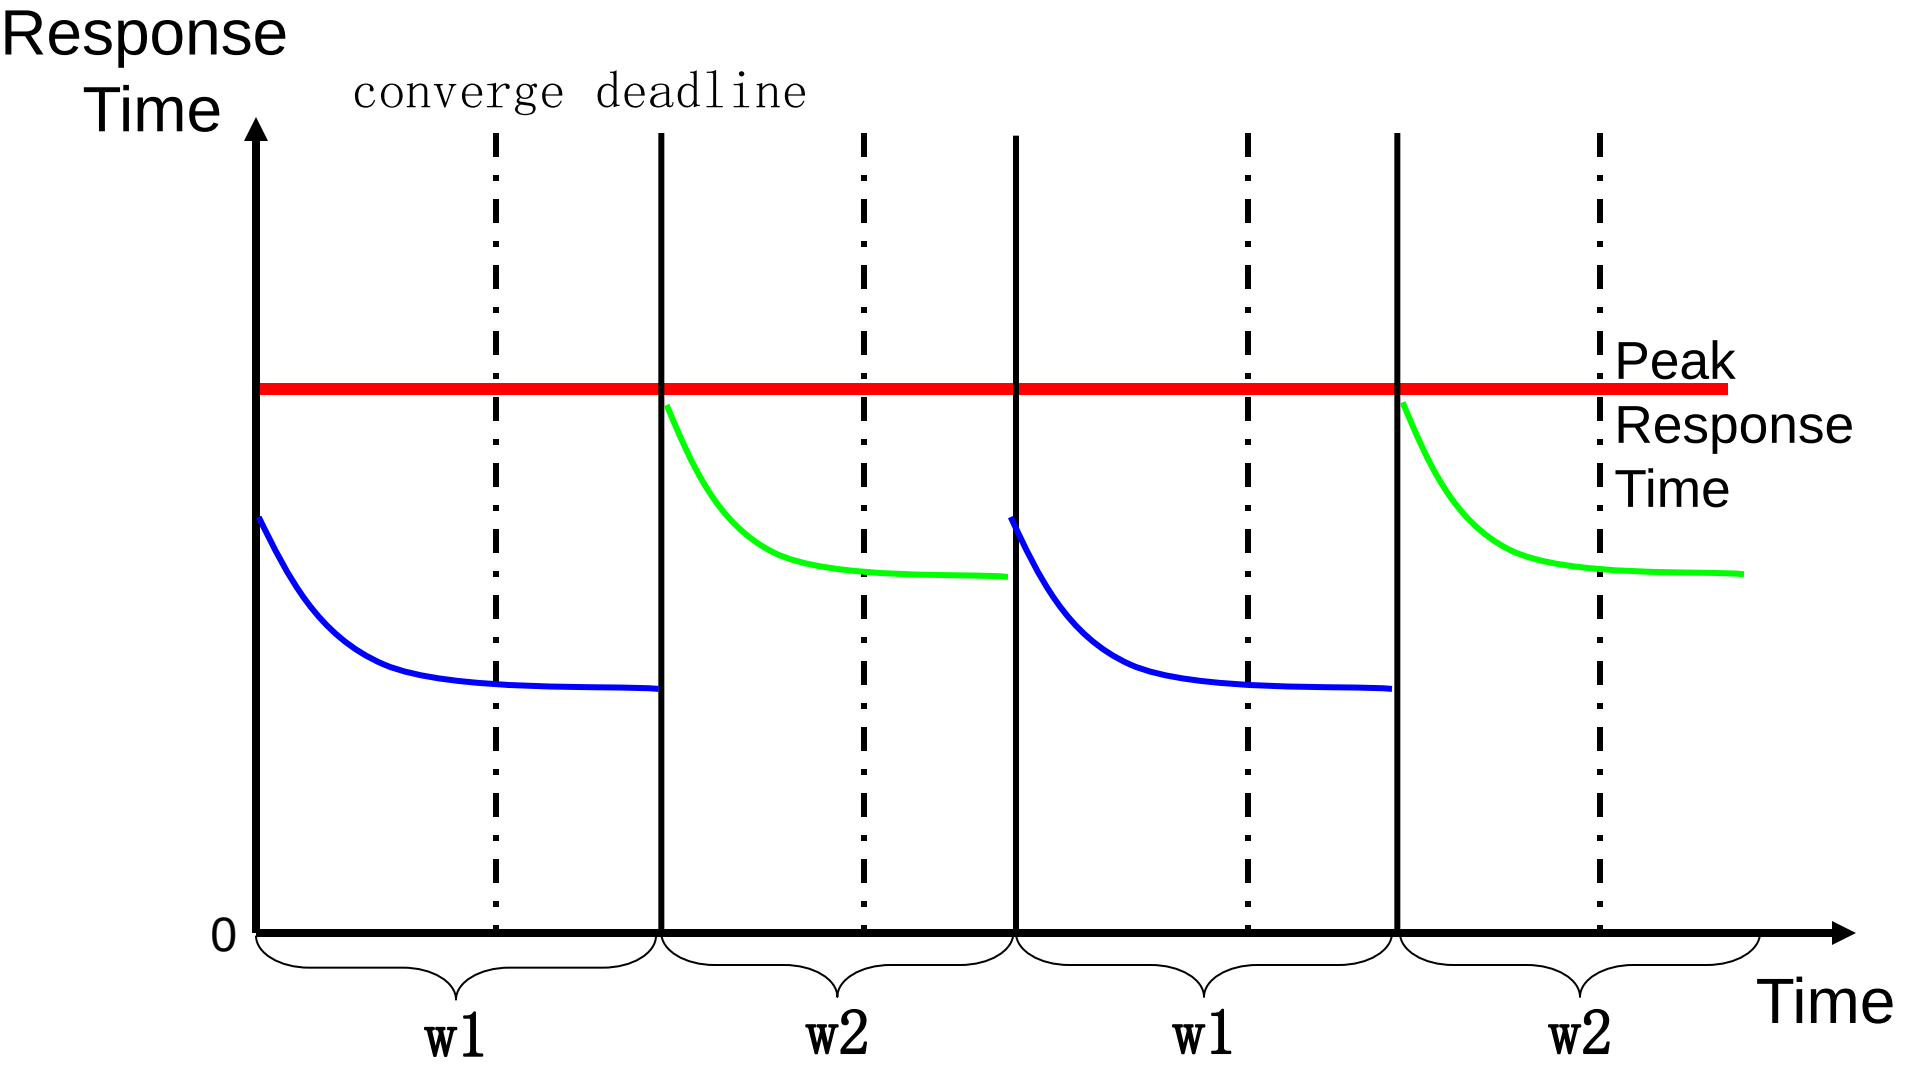

Lookahead Data Migration