Encoding for Reinforcement Learning Driven Scheduling

By: Alex Allen & Adam Brohl

Reinforcement Learning Driven Scheduling

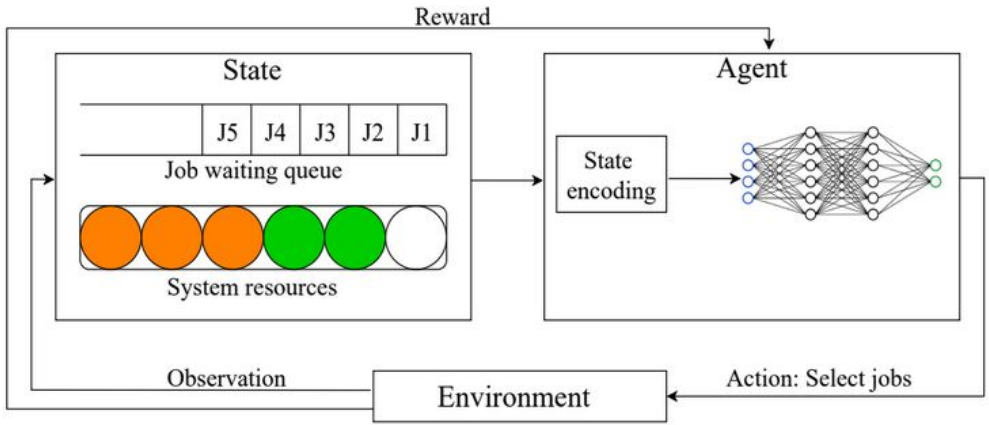

Fig. 1: Overview of RL driven scheduling. In the state, each system resource is represented by a circle. The circles sharing the same color indicate these system resources are allocated to the same running job; blank circles mean free/available system resources.

Existing Encoding Approaches

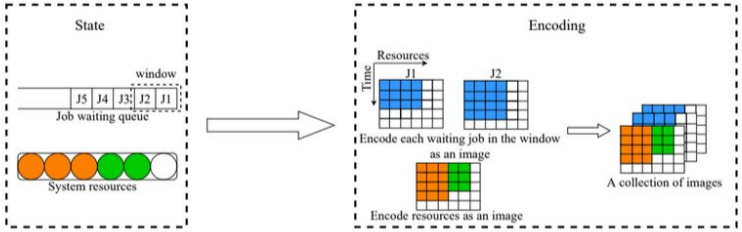

(a) Image-based state representation

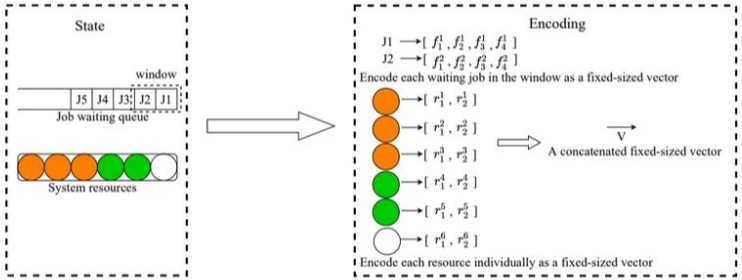

(b) Vector-based state representation

Scalable and Efficient Model (SEM) Representation

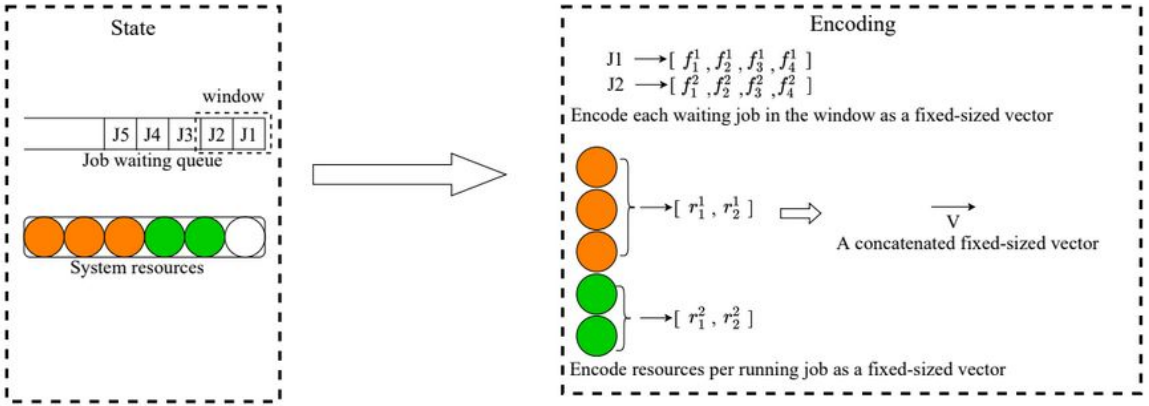

Fig. 4: SEM state representation. It captures system resources and user jobs as a fixed-size input to the RL agent depicted in the general RL driven scheduling shown in Figure 1.

Resource Encoding

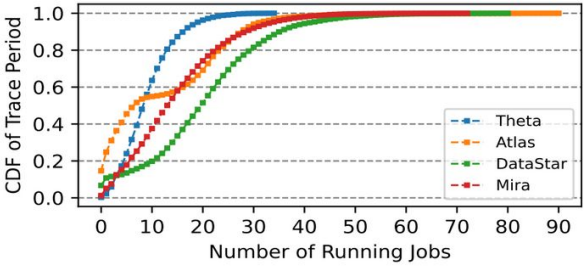

Fig. 5: The relationship between the number of running jobs and cumulative distribution function of trace period.

Table 2: Maximum number of running jobs

Workload	System Size	Maximum Number of Jobs
Theta	4,360	32
Mira	49,152	71
Atlas	9,216	90
DataStar	1,664	81

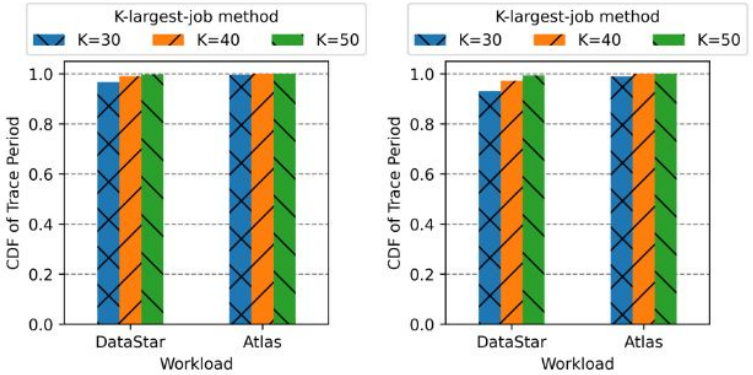

(a) 90% of the allocated nodes (b) 95% of the allocated nodes

Fig. 7: Trace coverage by using different K values

Table 3: Information loss when using the K-largest-job method

Workload	Information loss (node-hour loss %)		
	K=30	K=40	K=50
DataStar	1.22%	0.36%	0.01%
Atlas	0.23%	0.03%	9e-05

Evaluation

Table 4: RL agent configurations under SEM and DRAS

Configuration	Theta		DataStar	
	SEM	DRAS	SEM	DRAS
State Vector Size	268	8,920	280	3,528
Convolutional Layer	134	4,460	140	1,764
Fully Connected Layer 1	200	4,000	200	1,000
Fully Connected Layer 2	100	1,000	100	250
Output	50	50	50	50

Results

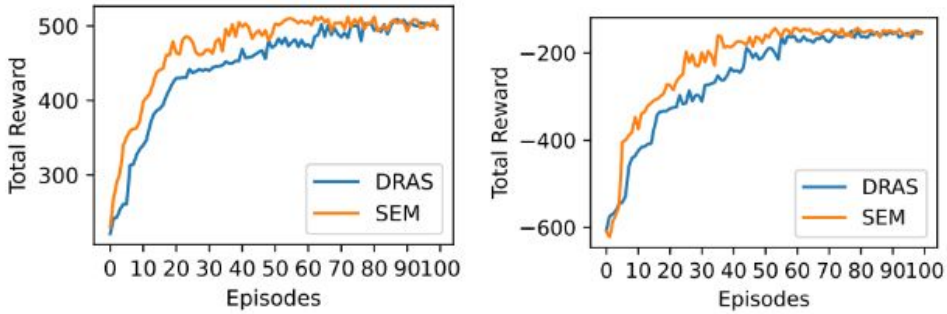

(a) Theta (b) DataStar

Fig. 8: Convergence rate of RL agent using different state encodings.

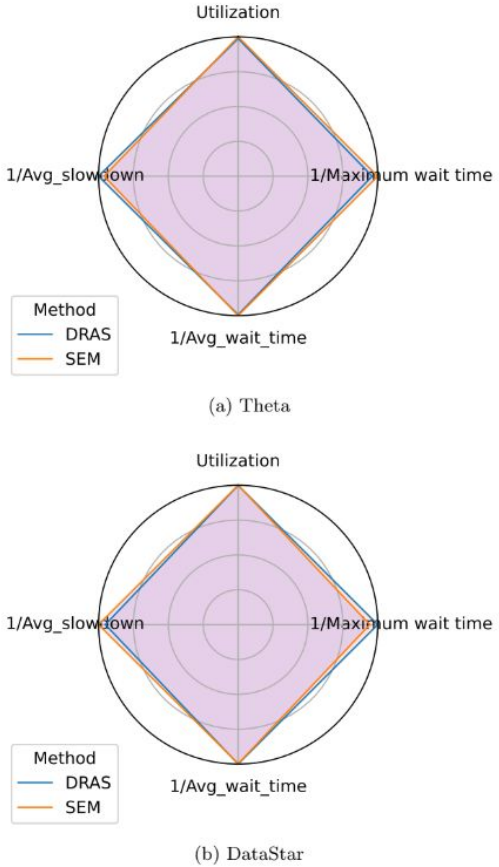

(a) Theta

(b) DataStar

Fig. 9: Scheduling performance by using different state encodings. The plots use Kiviat graphs to provide a comprehensive view of scheduling performance. The larger the area is, the better the overall performance is. It indicates that SEM can achieve comparable scheduling performance as those obtained by the existing state encoding.

Results (Cont.)

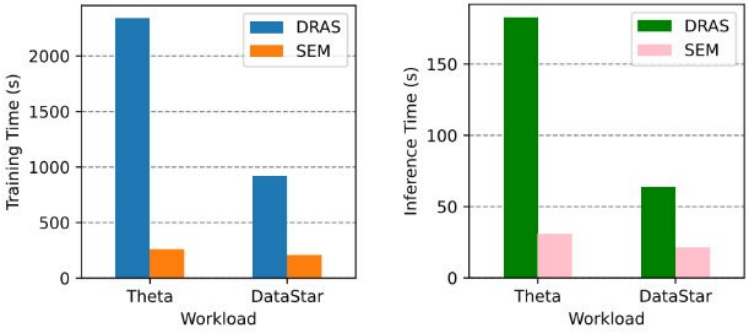

(a) Training time (b) Inference time

Fig. 10: Comparison of training and inference times.

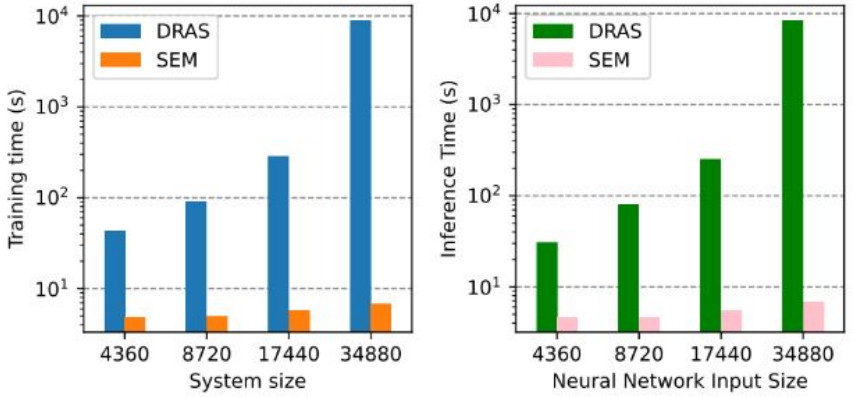

(a) Training time (b) Inference time

Fig. 11: Analysis of training and inference times when system size increases. One thousand jobs are used for training and inference testing separately.

Contributions

- Creation of a new generic state representation called SEM for RL driven scheduling
  - Captures state of scheduling environment in fixed-size vector
  - Uses new method for capturing system state
  - Two methods for creating fixed-size vectors (Zero-padding and K-largest-job)
- Evidence showing that SEM outperforms traditional vector-based models
- Future Work
  - Expanding SEM to work for multi-resource scheduling