

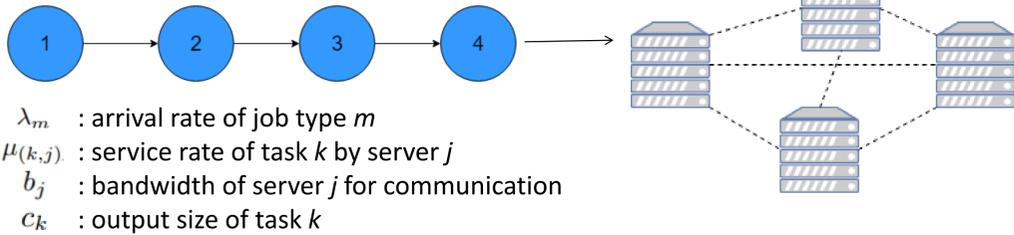
Communication-Aware Scheduling of Serial Tasks for Dispersed Computing

Chien-Sheng Yang¹, Salman Avestimehr¹, Ramtin Pedarsani²

¹University of Southern California ²University of California, Santa Barbara

System Model

M different types of jobs where each type is specified by a chain of tasks. Edge of the chain represents a precedence constraint. In the processing network, there are J servers.



λ_m : arrival rate of job type m
 $\mu_{(k,j)}$: service rate of task k by server j
 b_j : bandwidth of server j for communication
 c_k : output size of task k

Definition: A scheduling policy is throughput-optimal if it stabilizes the system when there exists a stabilizing scheduling policy.

Goal: Find a throughput optimal scheduling policy.

Capacity Region

Capacity region is determined by an LP:

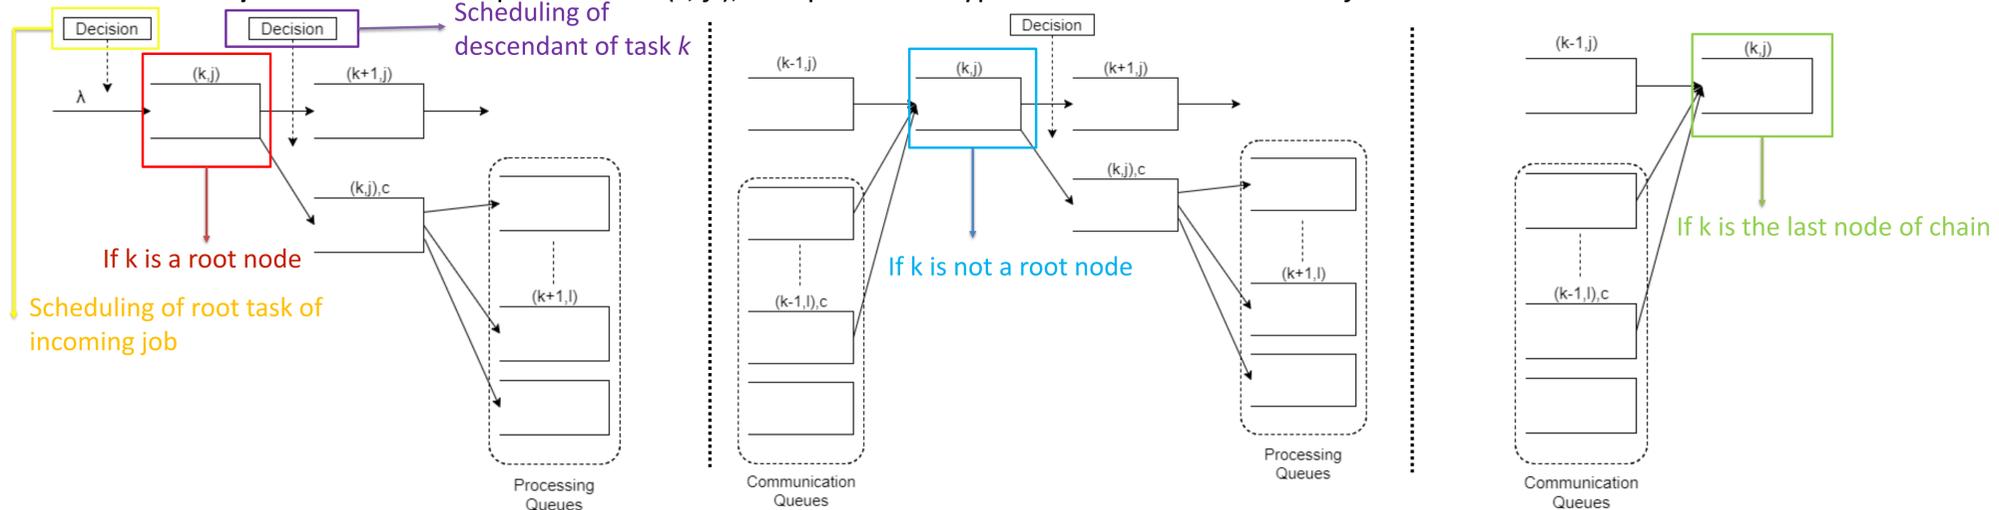
$$\begin{aligned} & \text{Maximize } \delta \\ & \text{subject to } \nu_k \leq \sum_{j=1}^J \mu_{(k,j)} p_{(k,j)} - \delta, \quad \forall 1 \leq k \leq K \\ & \quad \frac{b_j q_{(k,j)}}{c_k} \geq \mu_{(k,j)} p_{(k,j)} - \mu_{(k',j)} p_{(k',j)}, \quad \forall 1 \leq j \leq J, \quad \forall (k, k') \in \mathcal{E} \\ & \quad 1 \geq \sum_{k=1}^K p_{(k,j)}, \quad \forall 1 \leq j \leq J \\ & \quad 1 \geq \sum_{k \in \{1,2,\dots,K\} \setminus \mathcal{H}} q_{(k,j)}, \quad \forall 1 \leq j \leq J \\ & \quad p_{(k,j)} \geq 0, \\ & \quad q_{(k,j)} \geq 0. \end{aligned}$$

Arrival rates are in the capacity region if $\delta^* \geq 0$

Queuing Network

Processing queue: virtual queue called (k, j) for type- k tasks which are processed at server j

Communication queue: one virtual queue called $(k, j), c$ for processed type- k tasks sent from server j to other servers



Equivalent Capacity Region

Corresponding to the queuing network, we define an **optimization problem**:

$$\begin{aligned} & \text{Maximize } \gamma \\ & \text{subject to } r_{(k,j)} \leq \mu_{(k,j)} p_{(k,j)} - \gamma, \quad \forall 1 \leq j \leq J \text{ and } \forall 1 \leq k \leq K. \\ & \quad r_{(k,j),c} \leq \frac{b_j q_{(k,j)}}{c_k} - \gamma, \quad \forall 1 \leq j \leq J \text{ and } \forall k \in \{1, 2, \dots, K\} \setminus \mathcal{H}. \\ & \quad 1 \geq \sum_{k=1}^K p_{(k,j)}, \quad \forall 1 \leq j \leq J \\ & \quad 1 \geq \sum_{k=1}^K q_{(k,j)}, \quad \forall 1 \leq j \leq J \\ & \quad 1 = \sum_{j=1}^J u_{m \rightarrow j}, \quad \forall 1 < m < M \\ & \quad 1 = \sum_{l=1}^J s_{k,j \rightarrow l}, \quad \forall 1 \leq j \leq J \text{ and } \forall k \in \{1, 2, \dots, K\} \setminus \mathcal{H}. \\ & \quad p_{kj} \geq 0, \\ & \quad q_{kj} \geq 0, \\ & \quad u_{m \rightarrow j} \geq 0, \\ & \quad s_{k,j \rightarrow l} \geq 0. \end{aligned}$$

Lemma: The capacity region characterized by the optimization problem is equivalent to capacity region by the LP.

Throughput Optimal Policy

Max-weight policy :

Given the lengths of virtual queues at time n , Max-weight policy allocates the allocation vectors p, q, u and s , i.e.

$$\begin{aligned} & \arg \max_{p,q,u,s \text{ are feasible}} -(Q(n))^T E[\Delta Q(n) | \mathcal{F}^n] - (Q_c(n))^T E[\Delta Q_c(n) | \mathcal{F}^n] \\ & = \arg \min_{p,q,u,s \text{ are feasible}} (Q(n))^T E[\Delta Q(n) | \mathcal{F}^n] + (Q_c(n))^T E[\Delta Q_c(n) | \mathcal{F}^n], \end{aligned}$$

The Max-Weight policy is the choice of p, q, u and s that minimizes the drift of change of a Lyapunov function:

$$V = \sum_{k,j} Q_{(k,j)}^2(n) + \sum_{k,j} Q_{(k,j),c}^2(n).$$

Theorem: Max-Weight policy for the queuing network is throughput-optimal.

Max-Weight policy is rate stable for all the arrival vectors in the capacity region presented in optimization problem.

Max-weight policy makes the underlying Markov process of the queue-lengths positive recurrent for all the arrival rate vectors that are in the interior of the capacity region.

Ongoing Work and Future Directions

- Consider jobs modeled as directed acyclic graphs (DAG) and design queuing network with low complexity.
- Simulations on the real world scenario, ex: Amazon EC2, Digital Ocean, etc.