Resource Allocation in 1-D Distributed Analytics Network

Nitish K. Panigrahy (UMass), **Prithwish Basu** (Raytheon BBN), **Don Towsley** (UMass), **Ananthram Swami** (ARL), **Kevin S. Chan** (ARL) **and Kin K. Leung** (Imperial)

Issue: How to allocate distributed analytics tasks to servers?

Distributed Analytics Network: Requests & servers on 1-D line

Goal: To characterize the expected distance between a request and its allocated server

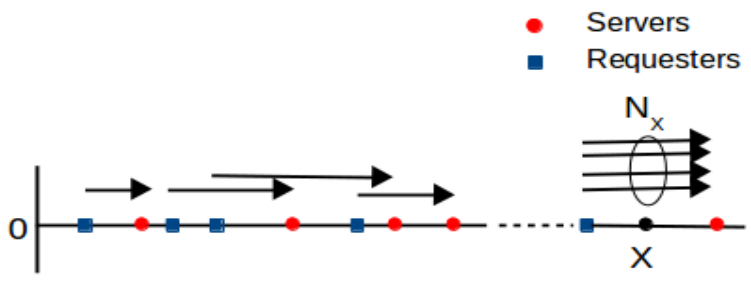

Figure 1: Allocate requesters to servers on the 1-D network under Move to Right (MTR) policy.

Consider a distributed-analytics network:

- R: Set of requesters of analytics tasks
  - Poisson-distributed on a 1-D line
  - λ: Requester density
- S: Set of servers providing resources
  - G_b: Inter-server distance distribution
  - $C_j = c$: Capacity of server j
- Performance Metric: *Request Distance* travelled by a request to allocated server

Unidirectional Matching

Allocate (match) each request to nearest free server to the right of the request (MTR policy)

Poisson Matching ($G_b \sim Exp(\mu)$)

- $c = 1$: Maps to M/M/1 queue
  - *Request Distance ~ Sojourn time*
  - *Avg. Request Distance* = $1/(\mu - \lambda)$
- $c > 1$: Maps to bulk service M/M/1 queue
  - *Avg. Request Distance* = $r_0 / \lambda (1 - r_0)$
  - r_0: root of $\mu r^{c+1} - (\lambda - \mu)r + \lambda = 0$

General Matching ($c = 1$)

- Map to M/G/1 queue with exceptional service G_e for 1st job in each busy period

$$G_e(x) = \frac{\lambda e^{\lambda x}[B(\infty) - B(x)] - \lambda B(\infty)}{1 - \lambda B(\infty)},$$

$$B(x) = \int_0^x G_b(z) e^{-\lambda z} dz$$

Bidirectional Optimal Matching

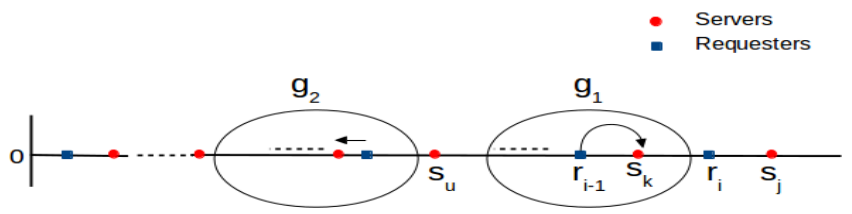

Figure 2: Formation of Groups in a Bidirectional Optimal Allocation.

Properties of Optimal Allocation

- Group: Consecutive allocated requests, servers separated by unallocated servers
- Optimal solution is set of groups
- Allocations do not cross each other

Dynamic Programming Solution (DPS)

- i: Request index, j: Server index
- $d(i,j) = \min\{|R[i] - S[j]| + d(i-1, j-1), d(i, j-1)\}$
- Time complexity: $O(|R|^2)$

Linear Expected Running Time Soln. (LERS)

- Shift left until either more expensive to perform shift or no unallocated server to left of current group
- Time complexity: $O(|R|g^2)$, g: Maximum group size
- Conjecture: g_{avg} is independent of |R|

Simulation Results: Group Size vs. |R| in LERS (Left); Comparison of Allocation Policies (Right)

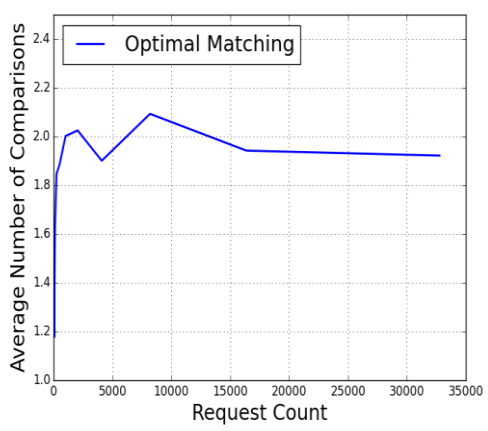

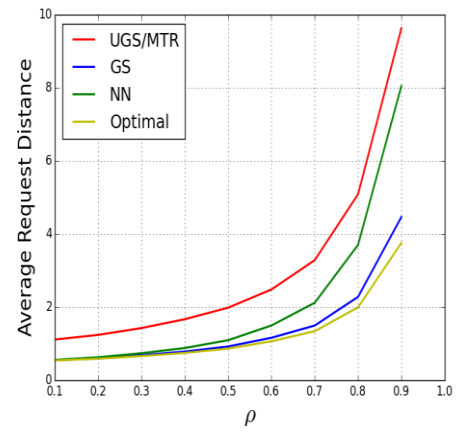

Distributed Analytics and Information Science International Technology Alliance
Annual Fall Meeting, September 2018