Multicast-based Weight Inference in General Network Topologies

Yilei Lin (PSU), **Ting He** (PSU), **Shiqiang Wang** (IBM US), **Kevin Chan** (ARL), **Stephen Pasteris** (UCL)

Objectives

- Validate the performance of distributed analytics deployed through SDN and VNF based on end-to-end measurements
- Formulate and solve a novel problem called the weight inference problem to infer performances at finest granularity

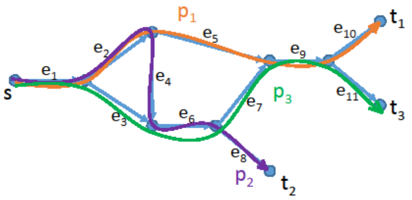

e_1:	$u_{e_1} = 0.0175$	{1,2,3}
e_2:	$u_{e_2} = 0.0071$	{1,2}
e_3:	$u_{e_3} = 0.0094$	{3}
e_4:	$u_{e_4} = 0.0421$	{2}
e_5:	$u_{e_5} = 0.0193$	{1}
e_6:	$u_{e_6} = 0.0363$	{2,3}
e_7:	$u_{e_7} = 0.0478$	{3}
e_8:	$u_{e_8} = 0.0394$	{2}
e_9:	$u_{e_9} = 0.0066$	{1,3}
e_{10}:	$u_{e_{10}} = 0.0247$	{1}
e_{11}:	$u_{e_{11}} = 0.0222$	{3}

Edge weights and edge categories

Technical Challenges

- What is the *finest granularity* of information about link/node performances that can be identified from end-to-end measurements?
- What probing method can reveal this information?
- How to solve the inference problem efficiently after probing?

Approaches

- Cast weight:

$$\phi_C := -\log(\Pr\{X_C=1\}) = -\log\Big(\prod_{e\in\bigcup_{i\in C} p_i} \alpha_e\Big) = \sum_{e\in\bigcup_{i\in C} p_i} u_e.$$

  X_C is the indicator that a k-cast probe sent on paths $\{p_i \in C\}$ successfully reach all the destinations.
- Category A: Set of edges traversed only by paths with indices in A.
- Category weight w_A: sum weight of edges in category A (finest granularity of inferable information)
- The weight inference problem aims at inferring the category weights from the measured cast weights.

$$\sum_{A\in\mathcal{A}:A\cap C\neq\emptyset} w_A = \phi_C, \quad \forall C\in\mathcal{C}.$$

  under the constraint

$$w \geq 0$$

- Theorem 1: If k=n (unconstrained multicast), then the weight inference problem has a unique solution.
- If k<n, we try to minimize $\|w\|_0$ (sparse recovery).
- Theorem 2: If k=O(1), then basis pursuit (BP) based on the simplex method provides an $O(n^k)$-sparse solution.

Algorithm 2 Non-negative Orthogonal Matching Pursuit

Initialization: $L = \emptyset$, $w = 0$ and $r = \phi$
while $\max D^T r > 0$ **do**
  $l \leftarrow \arg\max D^T r$
  $L \leftarrow L \cup \{l\}$
  $w_L \leftarrow \arg\min_{x\geq 0} \|\phi - D_L x\|_2$
  $r \leftarrow \phi - D_L w_L$
end while

Military & Coalition Relevance

- Provide “Looking Glass” service in distributed systems
- Validate communication/computation performances in other domains based on own observations

Results

- nOMP performs the best in reconstructing the measurements.
- nOMP also has the highest accuracy in inferring the category weights.
- nOMP achieves similar accuracy as the simplex method with a much shorter running time.

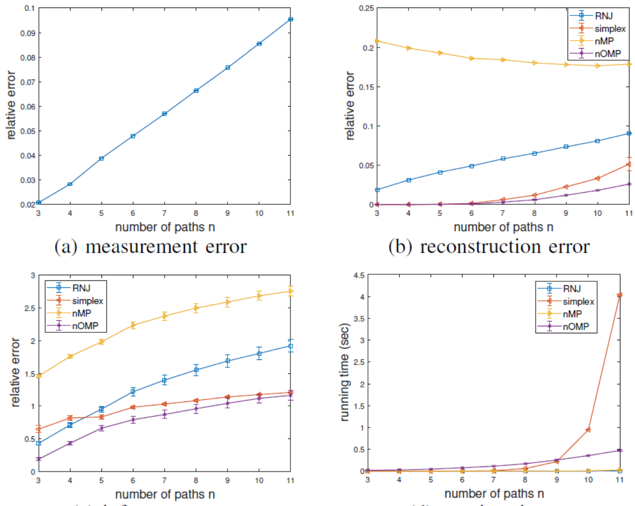

(a) measurement error (b) reconstruction error (c) inference error (d) running time

Relative error as the number of paths n varies (k = 2, 10000 k-cast probes).

- Increasing the “width” of multicast probes significantly increases the accuracy of the inferred category weights.

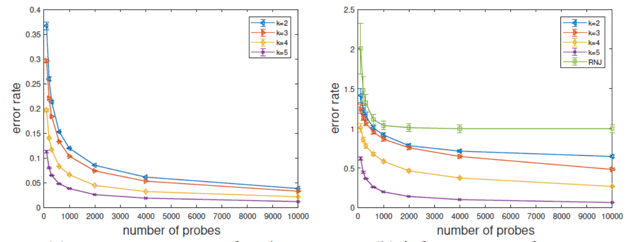

(a) measurement error for ϕ (b) inference error for w

Relative error as k and the number of k-cast probes vary (n = 5).

Summary & Future Work

- Formulated the problem of performance validation for distributed analytics systems as a generalization of the topology and link weight inference problem when routing paths may not follow trees.
- Focused on a novel sub-problem, the weight inference problem, to infer link weights to the finest granularity.
- Proved identifiability under unconstrained multicast and developed efficient heuristics under constrained multicast.
- Future work: apply the result of the weight inference problem to accurate inference of general network topologies.

Publication(s) & Impact

- Papers in INFOCOM 2019 and ICC 2019