Analyzing spatially constrained Power of two choice based policies in a two-dimensional distributed service network

Nitish K. Panigrahy[†], Prithwish Basu[¶], Don Towsley[†], Ananthram Swami[‡], Kevin S. Chan[‡] and Kin K. Leung[§]
[†] University of Massachusetts Amherst, MA, USA. Email: {nitish, towsley}@cs.umass.edu
[¶]Raytheon BBN Technologies, Cambridge, MA 02138, USA. Email:prithwish.basu@raytheon.com
[‡]Army Research Laboratory, Adelphi, MD 20783, USA. Email:{ananthram.swami, kevin.s.chan}.civ@mail.mil
[§]Imperial College London, London SW72AZ, UK. Email: kin.leung@imperial.ac.uk

ABSTRACT

We consider a power of two choices based allocation policy where both resources and users are located on a two-dimensional plane. First, we consider a "Spatial Power of two" (*sPOT*) policy which sequentially allocates a user to the least loaded resource among its two nearest resources. We show that sPOT maps to a classical balls and bins allocation policy with bins corresponding to the voronoi regions associated with the second order voronoi diagram of the set of resources. We provide closed form expression for the lower bound on the asymptotic maximum load on the resources and show that the classical Power of two (*POT*) benefits are not achieved by sPOT. Further, when resources are placed in an equidistant manner on a two-dimensional grid, sPOT maps to POT on the Delaunay graph associated with the voronoi tessellation of the set of resources. We show that the associated Delaunay graph is Δ-regular with $\Delta = 4$ and provide expressions for asymptotic maximum load using results from the literature. Finally, we propose a candidate set based sPOT policy (*k-sPOT*) in which a user uniformly at random selects two resources from a candidate set consisting of its k closest resources and gets allocated to the leastly loaded resource. Numerical results suggest that when $k = O(n^\alpha)$ with $\alpha > 0$ and n being the number of users, k-sPOT provide POT benefits.

1. INTRODUCTION

Past few years have seen a noteworthy development in the utilization of appropriated distributed services involving code, data and computational resources. In numerous such systems, for instance, Internet of Things (`IoT`) [2], an enormous number of computational and capacity resources are generally deployed in the physical world. Thus the spatial appropriation of resources assumes a significant role in deciding the overall performance of the service network.

While allocating resources to users on a service network, the maximum number of users assigned to a resource, or *load*, plays an important role in determining the efficacy of the allocation policy. A well known policy for balancing load in parallel and distributed systems is the Power of two (*POT*) choice based policy [7] where users are allocated to the least loaded resource among any two resources chosen uniformly at random from the set of resources. POT provides an exponential improvement in maximum load as compared to the classical randomized policy where each user is randomly allocated to a resource.

While much of the work on resource allocation has focused on balancing the load on the resources, little importance has been given to considering the spatial aspects of user and resources. Byers et al. [4] first studied POT in a geometric setting where resources are picked independently but with nonuniform probabilities proportional to the area of the Voronoi cells surrounding the resources. However, they do not consider the spatial distribution of users to characterize the expected load behavior. In this paper, we apply POT based load balancing paradigm in natural settings for resource allocation where users and resources are distributed in a Euclidean plane.

We consider the placement of users and resources in a two dimensional geographic region. Users and resources are placed uniformly at random across the region. We assume communication between the devices involves non-interfering single hop wireless transmission, otherwise called the *Direct transmission model* in the literature [1]. Our goal is to allocate users to resources under various allocation policies and characterize their asymptotic maximum expected load behavior. We focus on the following policies for allocating users to resources.

- **Spatial Power of two (*sPOT*):** In this policy, each user is sequentially allocated to the least loaded resource among its two nearest resources.
- **Candidate set based sPOT (*k-sPOT*):** Under this policy, each user uniformly at random selects two resources from a candidate set consisting of its k closest resources and is allocated to the leastly loaded resource.

Our contributions are summarized below:

1. Analysis of *sPOT* yielding lower bound expressions for asymptotic maximum expected load.
    - When users and resources are placed uniformly at random on a 2-D region, we model *sPOT* as classical balls and bins allocation policy with bins corresponding to the voronoi regions associated with the second order voronoi diagram of the set of resources..
    - When resources are placed on a two-dimensional grid, we model the situation using POT on the

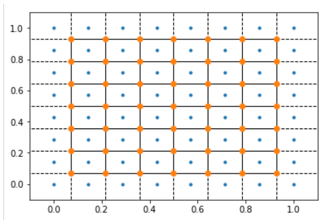

Figure 1: Voronoi Diagram of grid based resource placement

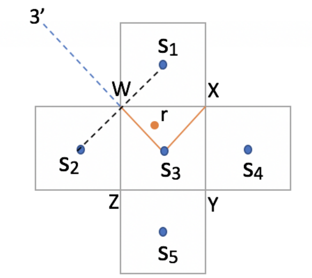

Figure 2: Second nearest region for user r

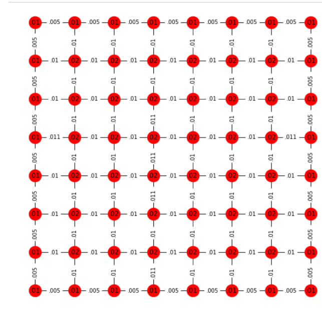

Figure 3: Delaunay Graph associated with grid based resource placement

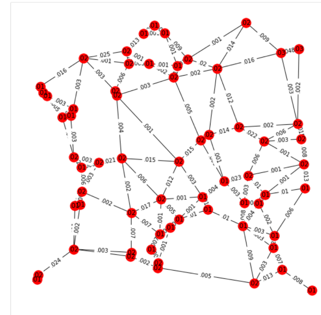

Figure 4: Delaunay Graph associated with uniform resource placement

Delaunay graph associated with the voronoi tessellation of the set of resources.

2. Through simulation we compare the performance of different allocation policies: *sPOT*, *k-sPOT*, *POT* etc.

3. We also study the effect of user and resource mobility on maximum expected load.

2. TECHNICAL PRELIMINARIES

Consider a set of users R requesting service and a set of resources S that can execute these requests. Suppose users and resources are located on a two dimensional plane, $\mathcal{D}$, placed uniformly at random. Denote $m = |R|$ and $n = |S|$. We consider the case when $m = n$. We define the following user allocation strategies.

2.1 Allocation Policies

Policy	Description
Power of One (POO)	Select one of the s resources uniformly at random.
Power of Two (POT)	Select the least loaded resource among two of the s resources chosen uniformly at random.
Spatial POO ($sPOO$)	Select the nearest resource.
Spatial POT ($sPOT$)	Select the least loaded resource among two nearest resources.
Candidate Set Based sPOT (k-$sPOT$)	Select the least loaded resource among two resources chosen randomly from k nearest.

Table 1: List of user allocation policies.

We present various user allocation strategies in Table 1. Note that, the expected load analysis for policies such as POO and POT are very well studied in the literature. Thus we have the following corollaries [7].

Corollary 1. *With high probability, the maximum load on any resource under POO is approximately* $\frac{\log n}{\log \log n}$.

Corollary 2. *With high probability, the maximum load on any resource under POT is* $\log \log n + O(1)$.

Thus moving from POO to POT leads to an exponential improvement in maximum load, also known as *POT benefit.*

We also extend the policies to consider mobility of users and resources under various mobility models. Precisely we consider random way point (*rwp-sPOT*), random direction (*rd-sPOT*) and random walk (*rw-sPOT*) mobility models for users and resources.

2.2 Performance Metric

To evaluate and characterize the performance of various allocation policies, we consider the maximum asymptotic load on all resources as the performance metric. Let $\pi : R \to S$ be an allocation policy. We define load on a resource as follows.

Definition 1. *Load on a resource, s, under allocation policy π, is defined as* $L_s = |\{r|\pi(r) = s\}|$.

2.3 Geometric Structures

We define the following geometric structures, which are useful constructs for analyzing various user allocation policies.

2.3.1 *Voronoi Diagram*

A Voronoi cell around a resource $s \in S$ is defined as the set of points in $\mathcal{D}$ that are closer to s than to any other resource in $S \setminus \{s\}$ [3]. The Voronoi diagram V_S of S is the set of Voronoi cells of resources in S.

2.3.2 *Delaunay Graph*

Define the Delaunay graph, $G_S(X, E)$, associated with S as follows. Assign $X = S$ and add an edge between resources u and v, i.e. $e = (u, v) \in E$ only if Voronoi cell of u and v are adjacent.

2.3.3 *Higher order Voronoi diagram*

A p^{th} order Voronoi diagram, $H_S^{(p)}$, can be defined as partition of $\mathcal{D}$ into regions such that points in each region have the same p closest resources in S.

3. SPATIAL POWER OF TWO POLICY

We now focus on the load behavior of sPOT policy for various resource placement settings. We assume users are placed uniformly at random on $\mathcal{D}$.

3.1 sPOT with Grid based resource placement

Consider the case where resources are placed on a two dimensional grid, $\mathcal{D} : \sqrt{n} \times \sqrt{n}$ with wrap-around. We show

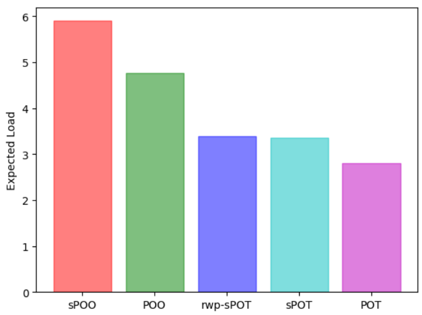

Figure 5: Comparison of various allocation policies.

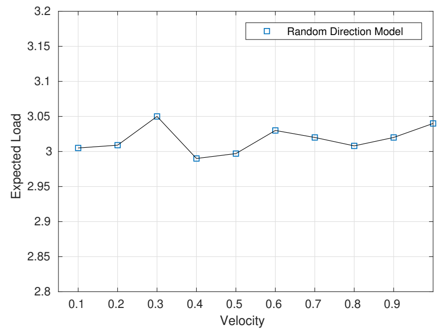

Figure 6: Effect of velocity on expected load.

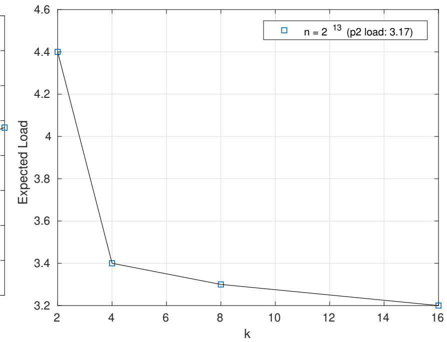

Figure 7: Candidate set size and Expected load behavior.

the Voronoi diagram:V_S for such a setting in Figure 1. The blue points are the resources while the orange points are the Voronoi vertices. Let $A(s, r, l)$ be the event such that a random user $r \in R$ is l^{th} closest to $s \in S$ among all resources in S. Let $B(s_1, s_2, r)$ be the event that the two nearest resources of r are in $\{s_1, s_2\}$. Denote $NN(r)$ as the closest resource of r. Thus we have

$$\begin{aligned}\Pr[B(s_1, s_2, r)] = &\Pr[A(s_1, r, 1)] \Pr[A(s_2, r, 2)|NN(r) = s_1]\\ &+ \Pr[A(s_2, r, 1)] \Pr[A(s_1, r, 2)|NN(r) = s_2]\end{aligned} \tag{1}$$

It is not difficult to show that all Voronoi cells in V_S have equal area. As $\Pr[A(s, r, 1)]$ is proportional to the area of the Voronoi cell surrounding s, we have $\Pr[A(s, r, 1)]$ is equal $\forall\ s \in S$. We consider the following lemma proposed in [6].

Lemma 1. *Given a Δ-regular graph with n nodes representing n bins, if n balls are thrown into the bins by choosing a random edge and placing into the smaller of the two bins connected by the edge, then the maximum load is at least $\Omega(\log\log n + \frac{\log n}{\log(\Delta \log n)})$ with probability of $1 - 1/n^{\Omega(1)}$.*

We now propose the following theorem.

Theorem 1. *Suppose resources are placed on a two dimensional grid, $\mathcal{D} : \sqrt{n} \times \sqrt{n}$ with wrap-around. Let users be placed independently and uniformly at random on $\mathcal{D}$. Under sPOT, the maximum load over all resources is at least $\Omega(\frac{\log n}{\log\log n})$ with probability of $1 - 1/n^{\Omega(1)}$.*

Proof. With out loss of generality (*W.l.o.g.*) consider a user: r placed uniformly at random on $\mathcal{D}$ as shown in Figure 2. Let $NN(r) = s_3$. We now evaluate $\Pr[A(s_1, r, 2)|NN(r) = s_3]$. Clearly, $\Pr[A(s_1, r, 2)|NN(r) = s_3] \propto \text{Area}(\triangle WXS_3)$. We also have

$$\begin{aligned}\text{Area}(\triangle WXS_3) &= \text{Area}(\triangle WZS_3) = \text{Area}(\triangle YXS_3)\\ &= \text{Area}(\triangle ZYS_3)\end{aligned}$$

Thus $\Pr[A(s_i, r, 2)|NN(r) = s_3]$, for $i \in \{1, 2, 4, 5\}$ are all equal. By (1), $\Pr[B(s_i, s_j, r)]$ are all equal for $s_i \in S$ and s_j being any grid neighbor of resource s_i. Now consider the Delaunay Graph, G_S, associated with S as shown in Figure 3. G_S is Δ-regular with $\Delta = 4$. Let $e = (s_i, s_j)$ be an edge in G_S. We denote $\Pr[B(s_i, s_j, r)]$ as the probability for selecting e as mentioned in Lemma 1. Since $\Pr[B(s_i, s_j, r)]$ are all equal, direct application of Lemma 1, proves the theorem. □

We verify the claim: $\Pr[B(s_i, s_j, r)]$ *are all equal*, through simulation as shown in Figure 3. We assign $n = 64$ and empirically compute $\Pr[B(s_i, s_j, r)]$ and denote it as edge probability on edge e on the Delaunay graph. It is clear from Figure 3 that the edge probabilities are almost all equal.

Remark 1. *Note that, Theorem 1 ensures that no POT benefits are observed when resources are placed on a two dimensional grid.*

3.2 sPOT with Uniform resource placement

When both users and resources are placed uniformly at random on $\mathcal{D}$, we can no longer invoke Lemma 1. This is due to the fact that $\Pr[B(s_i, s_j, r)]$ would now not all be equal. This is evident from our simulation results on the corresponding Delaunay graph as shown in Figure 4. Note that the edge probabilities, i.e. $\Pr[B(s_i, s_j, r)]$, are all completely different from each other. Thus we resort to using a higher order structure: the second order Voronoi diagram to analyze the expected maximum load behavior.

Consider the second order Voronoi diagram: $H_S^{(2)}$ associated with the set of resources S. The number of Voronoi cells in $H_S^{(2)}$ is upper bounded by $O(3n)$ [Chapter 3.2, [8]]. We can invoke classical balls and bins argument on $H_S^{(2)}$ as follows inspired by the discussion in [6]. We treat each Voronoi cell in $H_S^{(2)}$ as a bin. Thus there are at most $O(3n)$ bins. If each ball (or user) chooses a bin uniformly at random, then the maximum expected load on a resource pair is lower bounded by $\log n/\log\log(3n)$. Thus one of the resource pair has load at least $\log n/2\log\log(3n)$. However, the area distribution of Voronoi cells in $H_S^{(2)}$ is non-uniform, thus each ball choses a bin from a distribution (say p). But one can think that the introduction of non-uniformity introduces more load imbalance to the system. Thus the lower bound on expected maximum load should remain intact and we have the following conjecture.

Conjecture 1. *Suppose both users and resources are placed independently and uniformly at random on $\mathcal{D}$. Under sPOT, the maximum load over all resources is at least $\Omega(\frac{\log n}{\log\log n})$ with high probability of $1 - 1/n^{\Omega(1)}$, i.e., we do not get POT benefits.*

4. PERFORMANCE EVALUATION

In this section, we evaluate the performance of various allocation policies.

4.1 Simulation set up

We consider $n = 2^6$ resources and equal number of users placed at a unit square uniformly at random. We ran 1000 trials for each of the strategy and take the average maximum load, motivated by the procedure adopted in [4]. For all the mobility models we take the maximum velocity to be 0.1m/s and also consider the boundary effect by reflection.

4.2 Comparison of allocation policies

We compare the performance of various allocation policies in Figure 5. As expected, sPOO performs the worse while POT provides the least maximum expected load. Note that the introduction of spatial aspects into a policy increases the expected load. While sPOT is the second best policy, adding mobility to users does not effect the overall performance of sPOT. We study the effect of mobility in details, in the next Section.

4.3 Effect of mobility on sPOT

In this section we discuss the effect of user and resource mobility on performance of sPOT policy.

4.3.1 Effect of user mobility

We consider various mobility models [5] to study its effect on maximum expected load behavior. We consider the effect of velocity on maximum expected load and present the results in Figure 6. It is evident that even though the velocity was increased ten folds, the change in expected maximum load of the system was negligible. This suggests that mobility of users does not provide POT benefits.

4.3.2 Effect of resource mobility

We also consider mobile resources with static users. However, the expected load more or less remains the same as that of static resource case.

4.4 Improving sPOT load behavior

In Section 3 we showed that for both grid and uniform based resource placement, sPOT does not provide POT benefits. Below we discuss k-sPOT, a modified version of sPOT policy to improve the overall load behavior and achieve POT benefits.

W.l.o.g, we define C_k to be the candidate set (of size k) consisting of k nearest resources for a particular user. In k-sPOT policy, the user selects two resources uniformly at random from C_k and assigns itself to the leastly loaded one. Clearly sPOT and POT are two extremes of the policy k-sPOT with $k = 2$ and $k = n$ respectively. We discuss the effect of candidate set size on expected load behavior as follows.

4.4.1 Candidate set size and Expected load behavior

We report on an experiment where we evaluate expected value of the maximum load across various values of k and n. The results are shown in Figure 7. It is clear from the figure that when $k = O(\log n)$, the expected load value converges to that of POT. Based on the results, we make the following conjectures.

CONJECTURE 2. *If the candidate set does not grow with n, no POT benefit is expected.*

CONJECTURE 3. *If the candidate set size $k = O(\log n)$, POT and k-sPOT have similar expected load behavior.*

5. CONCLUSION

In this work we considered a power of two choices based allocation policy where both resources and users are located on a two-dimensional plane. We analyzed the sPOT policy and provided expressions for the lower bound on the asymptotic maximum load on the resources. We claim that for both grid and uniform based resource placement, sPOT does not provide POT benefits. We propose a candidate set based modified sPOT policy and analyzed the case when it achieves POT benefits. We studied the effect of mobility through extensive simulations.

6. REFERENCES

[1] A. Asadi, Q. Wang, and V. Mancuso. A Survey on Device-to-device Communication in Cellular Networks. In *IEEE Commun. Surv. Tut*, 2013.

[2] L. Atzori, A. Iera, and G. Morabito. The Internet of Things: A Survey. *Computer Networks*, 54(15):2787–2805, 2010.

[3] B. A. Bash and P. J. Desnoyers. Exact Distributed Voronoi Cell Computation in Sensor Networks. In *IPSN*, 2007.

[4] J. Byers, J. Considine, and M. Mitzenmacher. Geometric Generalizations of the Power of Two Choices. In *SPAA*, 2004.

[5] T. Camp, J. Boleng, and V. Davies. A Survey of Mobility Models for Ad hoc Network Research. In *In: Wireless Communications and Mobile Computing*, 2002.

[6] K. Kenthapadi and R. Panigrahy. Balanced Allocation on Graphs. In *SODA*, 2006.

[7] M. D. Mitzenmacher. The Power of Two Choices in Randomized Load Balancing. In *Ph.D. Dissertation, Harvard University*, 1996.

[8] A. Okabe, B. Boots, and K. Sugihara. Spatial Tessellations Concepts and Applications of Voronoi Diagrams. In *New York: Wiley*, 1992.