Annex A: Analysis and Simulation Results

April 30th, 2015

Siamak Khaledi, Ankit Shah, Matthew Shoaf

1. Understanding the Problem Domain

1.1. Overview

After conducting interviews with our primary stakeholder, a representative of DMVFTA, our team identified the central problems associated with the existing system. These core problems were captured in Stakeholder Requirements 3.1.3 – 3.1.6 in the System Design Document (SDD). These requirements are as follows:

- The system shall de-conflict where vendor preferences are greater than location availability;
- The system shall match vendors to locations in an equitable way, where "equitability" is considered to be a measure of vendor preference values versus assigned locations;
- The system shall not employ auctioneering or bidding to valuate parking spaces monetarily.

It was assessed that the central problem was the result of high demand for a limited number of resources. Given that one of our specified constraints was that the food truck location assignment system could not employ monetary valuation of these resources (e.g. an auction), we determined that it was critical that the system be "fair." This presented us with a challenge as it was difficult to define "fairness".

Initially, we were uncertain of how evenly the existing system distributed high value vendor location assignments. We reasoned that the current system may be distributing high value locations equitably. However, because of design issues and limited resources (i.e. high value locations), vendors may not have perceived the system to be fair. This initial analysis of the problem domain led us to reason that before we made any design decisions, we had to first analyze the existing system to confirm whether there existed a real or perceived "fairness" problem. Differing outcomes of this analysis would lead to significantly different design recommendations for a new and/or improved location assignment system.

1.2. Building the Analytical Structure: Defining "Fairness"

In the previous section, we identified the central role that "fairness," both actual and perceived, played for the success of any food truck location assignment system. This implies that there exists some definition of "fairness" which we can measure objectively, and thus use to analyze the existing system and any alternate system that we might design. Our first real analytical task was therefore *to define the concept of "fairness" as it related to this problem domain.*

After considering various value models, our team settled on a very simple way to define "fairness" relative to food truck vending location distribution. First, we recognized that each vendor had preferences for the existing food truck locations, and that records of these preferences, rank-ordered, are held in the existing location distribution system. Conveniently, the outcomes of vendor location distribution (i.e. where the trucks were actually assigned to vend) are also held in record by the existing system.

From here, we realized that we might define fairness by a simple binary check of whether a vendor received his or her preferred location on a given day. In other words, we simply asked: *did the vendor get what he wanted in this instance, yes or no?* We could apply this methodology at each level (e.g. the vendor's first choice, second choice, and so on), and we need not apply a value curve as it is possible to look at each level independently.

Once these binary comparisons were constructed using historical data, it is now possible to get a sense of how "fair" the system is. Over a given period of time, the more closely the results conformed to a uniform pattern, the more "fair" the system was. That is, a uniform pattern indicated that all vendors were receiving an equal number of their preferences at any given level. Additionally, we could say that the more rapidly a system converged on a uniform pattern, the better it was at assigning locations to vendors in an equitable way.

We now had a quantifiable definition of fairness in this problem domain, and had developed two measures by which to judge any system that operated within. These measures were as follows.

- How closely does the system produce results that converge on a uniform pattern?
- Over what interval does the system produce results that converge on a uniform pattern?

These two measures became the yardstick by which we assessed the existing system and the later proposed Food Truck Location Assignment and Draft System (FTLADS).

1.3. Data Collection and Cleansing

Initially, our team intended to either negotiate with the existing vendor location distribution system's operators to obtain data or to scrape the system's website. Both of these tasks would have proven to be time consuming, but our team's primary stakeholder provided us with a well-developed data set.[1] The complete data set included preferences and outcomes for 22 unique trucks, operated by 14 unique businesses. It covered an interval of nine months from July of 2014 through March of 2015.

Our team conducted some cleansing of the data set. The data was trimmed for the following reasons.

- Some of the vendors had gaps within the interval; that is, they had not made preferences for each month studied.

[1] Our team owes a debt of gratitude to Mrs. Karen Wrege for collecting and building a data set that was well-structured and easy to use. Her efforts truly helped make this analysis possible.

- Many of the vendors had not input preferences for the month of March, 2015.

After trimming the data, we now had a normalized data set with 17 unique trucks, operated by 10 unique businesses, which covered a period of seven months. After a discussion with our primary stakeholder, our team discarded the "truck owner/business" factor, as the system should distribute locations on a "by license/truck," and not "by owner," basis. A sample of our cleansed working data can be found in Figure 1 below.[2]

VSP Number	Vendor Name	Month	Assign/Prefer	Monday	Tuesday	Wednesd	Thursday	Friday
XXXXX	XXXXX	1/15	A	Off	Navy Yard	Franklin	Off	Farragut
XXXXX	XXXXX	2/15	A	Off	Navy Yard	State	L'Enfant	Metro
XXXXX	XXXXX	3/15	A	Franklin	Off	L'Enfant	Off	Farragut
XXXXX	XXXXX	7/14	A	Union	Off	Metro	Off	Farragut
XXXXX	XXXXX	8/14	A	State	Off	Metro	Off	Farragut
XXXXX	XXXXX	9/14	A	Off	Franklin	Off	Navy Yard	Off
XXXXX	XXXXX	10/14	A	State	Off	Metro	Off	Farragut
XXXXX	XXXXX	11/14	A	State	Off	Metro	Off	Farragut
XXXXX	XXXXX	12/14	A	Off	Union	Off	Patriots	Metro
XXXXX	XXXXX	1/15	1	L'Enfant	Metro	Patriots	Metro	Farragut
XXXXX	XXXXX	2/15	1	L'Enfant	Metro	State	Farragut	Farragut
XXXXX	XXXXX	7/14	1	L'Enfant	Metro	State	Navy Yard	Farragut
XXXXX	XXXXX	8/14	1	L'Enfant				

Figure 1: Cleansed Data

1.4. Methodology and Valuation Analysis

Using seven months of normalized data, our team divided the data according to Preference Level (e.g. "First Preference," "Second Preference," and so on). Next we made a simple comparison between a given preference and the assignment, recording the results in a binary fashion. These results were then aggregated according to truck identification number. Figure 2 provides a synopsis of the binary comparison process.

Figure 2: Comparison and Aggregation

VSP Number	Vendor Name	Month	Assign/Prefer	Friday Pref	Friday Assign	FridayGotPref	By VSP Sum
XXXXX	XXXXX	1/15	1	Farragut	Farragut	1	
XXXXX	XXXXX	3/15	1	Farragut	Farragut	1	
XXXXX	XXXXX	7/14	1	Farragut	Farragut	1	
XXXXX	XXXXX	8/14	1	Farragut	Farragut	1	
XXXXX	XXXXX	9/14	1	Farragut	Off	0	
XXXXX	XXXXX	10/14	1	Farragut	Farragut	1	
XXXXX	XXXXX	11/14	1	Farragut	Farragut	1	
XXXXX	XXXXX	12/14	1	Farragut	Metro	0	6

While constructing this analysis, it became clear that some locations were much more valued by vendors than others. We recognized that the uneven valuation of vending locations could lead to a potential flaw in our analysis. Precisely, those vendors who choose high value locations more often are likely to receive their preferred spot less often than those who

2 A full copy of our working data in Microsoft Excel format can be found in Appendix 1 to this Annex.

choose less valued ones. This meant that the system may appear to produce unevenly distributed results, but that this was an artifact of uneven demand, and not a system flaw.

To account for this, we not only focused on overall distribution of results, but also chose to analyze high value locations specifically. Of course, in order to do this, we needed to determine how vendors valued the locations. The data lent itself to such a study; all we needed to do was sum vendor preferences for each location. In fact, we took this a step further, assuming that valuations varied not just by location, but by time. In other words, we wanted to account for the fact that some locations may be more valuable on certain days while different locations might be more valued on others. Figures 3 and 4 provide insight into vendor valuation of the various locations, and locations by time.

Preference Totals *Location*	*1st Preference*	*2nd Preference*	*3rd Preference*	*Total*
Farragut Square 17th St	195	177	191	563
Franklin Square 13th St	66	78	125	269
Union Station	91	84	155	330
L'Enfant Plaza	162	133	120	415
Metro Center	228	292	183	703
Waterfront Metro	22	16	16	54
Navy Yard/Capital River Front	38	39	33	110
Patriots Plaza	67	60	60	187
Virginia Ave (State Dept)	75	59	91	225

Figure 3: Vendor Valuation of Locations (1^{st} – 3^{rd} Preferences)

First Preference *Location*	*Monday*	Tuesday	Wednesday	Thursday	Friday
Farragut Square 17th St	4.76%	6.35%	13.23%	12.17%	66.67%
Franklin Square 13th St	6.35%	16.93%	3.17%	4.76%	3.70%
Union Station	10.58%	13.76%	8.99%	13.76%	1.06%
L'Enfant Plaza	17.46%	26.46%	25.40%	13.23%	3.17%
Metro Center	19.58%	15.87%	37.04%	35.98%	12.17%
Waterfront Metro	7.41%	1.59%	1.06%	1.06%	0.53%
Navy Yard/Capital River Front	9.52%	4.23%	1.06%	5.29%	0.00%
Patriots Plaza	14.29%	10.05%	2.12%	8.47%	0.53%
Virginia Ave (State Dept)	10.05%	4.76%	7.94%	5.29%	11.64%

Figure 4: Two Variable Analysis of Valuation (Day and Location, 1^{st} Preference Only)

Note that Figure 3 is based on *total preference count*, while Figure 4 is based upon *percentage of preference for each day*. Consider Metro Center as an example. Over the interval studied, the Metro Center vending location was chosen as a "First Preference" a total 228 times, and its aggregate "First," "Second," and "Third" ratings came to 703. Figure 2 shows us that Metro Center dominates on Wednesdays and Thursdays, receiving 37% and 36% of the total "First Preference" ratings for those days.

From our analysis of valuation, we now had a better idea of how the nine different locations were valued. In general, we concluded the following about vendor valuation.

- *The Farragut Square, Metro Center and L'Enfant Plaza locations were far more valued than the other six possibilities.* The weakest of these three, L'Enfant Plaza held a 25% advantage over the next most valued location, Union Station. L'Enfant Plaza was selected as a 1st, 2nd, or 3rd Preference 7.7 times more than the weakest overall performer, Waterfront Metro.
- *Vendor preferences varied considerably at different times of the week.* While Metro Center is the most valued location during the middle of the week, the Farragut Square location overwhelmingly dominates on Fridays.

These initial results of vendor valuation of various locations provided us with the ability to assess how the existing system performed at assigning locations generally, but also to see how well it performed at distributing high value locations specifically.

1.5. Results Analysis of Existing System

Our analysis showed that the overall distribution of vending locations as a function of First, Second and Third preferences *did not conform to a uniform distribution* during the interval studied. However, the system did attempt to provide those vendors who got worse results for their First Preference with better results for their Second and Third Preferences. We interpret this result to mean that the existing system was likely attempting to apply an value-based optimization algorithm toward the problem set, although we cannot confirm this as we did not have direct access to the system's code set. See Figure 5 for the overall results.

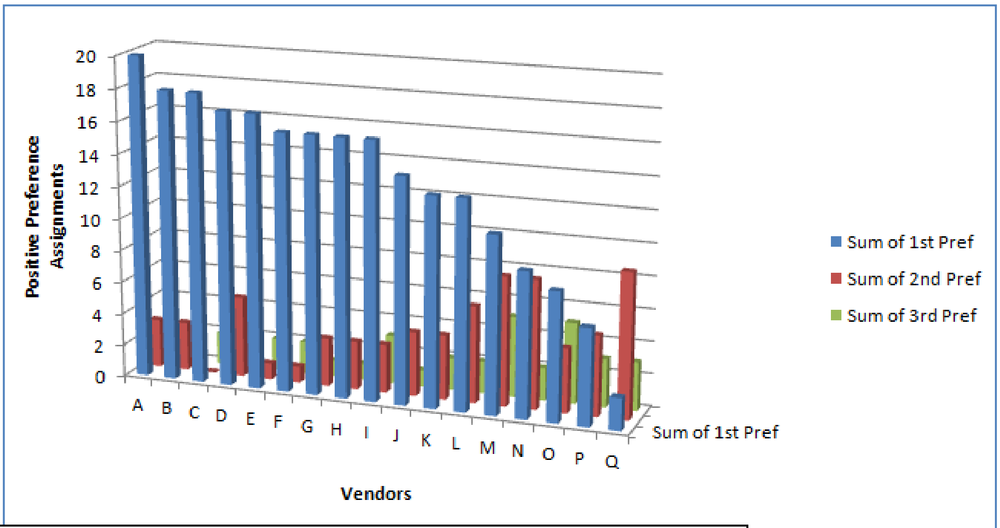

Figure 5: Existing System Assignment Distributions (Valuation Independent)

Figure 5 shows how the system distributed all preferences, regardless of whether the vendor chose high or low value locations. The results appeared uneven, but it should be noted that this could be an artifact of highly divergent valuations. One question remains: how did the system perform against those locations which are highly valued?

Our valuation results provide us with a very useful sub-set of data by which to answer this question. We can focus on those vendors who chose Farragut Square as their First Preference, specifically on Fridays, and see how the system distributes those locations. A look at Figure 4 shows that the Farragut-Friday cell holds by far the highest concentration of vendor First Preferences.[3] Therefore, to determine how the system performed we focused on only those vendors who chose Farragut as their First Preference on Fridays and looked at how the system distributed locations against those preferences. Figure 6 shows the results of this analysis.

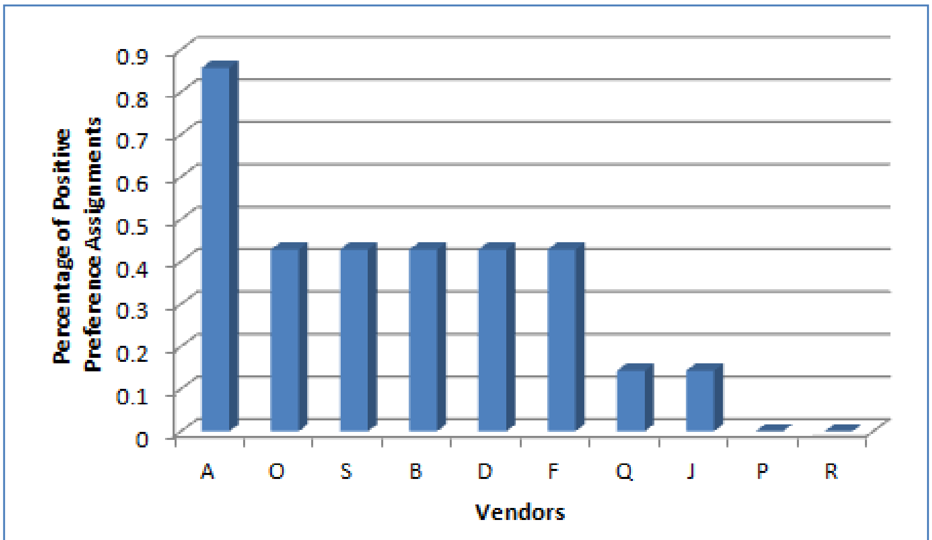

Figure 6: Existing System Assignment Distributions (Farragut-Friday, First Preference)

As we see from Figure 6, the system does not evenly distribute high value vending distributions. In fact, if anything, the most highly valued combination, Farragut-Friday, was distributed *more* unevenly than the broader set of location assignments. Our group did not attempt to explain why this is the case, as we did not have access to algorithms that the existing system used in order to match locations to vendor preferences.

Nevertheless, our analysis confirmed that the system did not evenly distribute location assignments, both in the general case and taking into consideration variance in valuation by vendors. This analysis heavily influenced our design; instead of focusing on usability and

[3] Note that we conducted the same analysis for vendors' First through Fourth Preferences. The First Preference set was the most unevenly distributed; lower valued sets were more evenly spaced among all points in the matrix. Among all combinations, preferences were most concentrated for the combination of First Preference, Farragut-Friday.

interface, we opted to develop a new method for assigning locations to vendors based upon their preferences. To test the effectiveness of our proposed solution we could replicate the study above. To do this, we executed a simulation of our system using an equivalent data set (i.e. equal numbers of vendors over an equivalent time interval) and recorded the results. The next section of this annex discusses the specific methodology and results of that simulation.

2. Validation of the Proposed Algorithm

2.1. Overview

In order to validate the proposed algorithm, it is important to test the outcome of the drafting mechanism against the requirements derived through stakeholder discussions and the evidence of lack of fairness from current system analysis. The validation looks to confirm the new system offers equal chances to vendors to get their preferences. The algorithm also needs to allow for an equal chance of getting those preferences on all different days of the week. A certain cycle length is needed in order to achieve fairness in providing the above.

2.2. Achieving Perfect Fairness

Using the developed algorithm, an initial wheel is generated to provide 252 draft numbers in an order that takes system constraints into account. Suppose the movement of a given vendor on the wheel is being monitored. Using the proposed algorithm, in the first 21 days after generation of the wheel, the vendor receives a number from each group of numbers explained in the design document exactly once, getting these numbers on different days of the week. This process is shown in Figure 7.

	Group	Mon	Tue	Wed	Thu	Fri
month 1	1-12	1				
	13-24		1			
	25-36				1	
	37-48	1				
	49-60				1	
	61-72					1
	73-84		1			
	85-96					1
	97-108			1		
	109-120			1		
	121-132			1		
	133-144		1			
	145-156		1			
	157-168	1				
	169-180				1	
	181-192					1
	193-204			1		
	205-216	1				
	217-228					1
	229-240	1				
	241-252				1	

Figure 7: Distribution of Draft Numbers after Month 1.

As the vendor continues to visit the numbers generated by the wheel through the next month, he will visit different numbers from the same groups of numbers, however this month he receives these numbers on the weekday immediately after which he got a number from the same value the previous month. For example, if he received a number between 13 to 24 on a Tuesday in the first month, in the second month he receives another number between 13 to 24 on a Wednesday. At the end of the second month he will have completed visiting the second set of numbers, but not received any 2 numbers belonging to the same level on the same weekday. Figure 8 shows the distribution of draft numbers at the end of month 2.

	Group	Mon	Tue	Wed	Thu	Fri
month 2	1-12	1	1			
	13-24		1	1		
	25-36				1	1
	37-48	1	1			
	49-60				1	1
	61-72	1				1
	73-84		1	1		
	85-96	1				1
	97-108			1	1	
	109-120			1	1	
	121-132			1	1	
	133-144		1	1		
	145-156		1	1		
	157-168	1	1			
	169-180				1	1
	181-192	1				1
	193-204			1	1	
	205-216	1	1			
	217-228	1				1
	229-240	1	1			
	241-252				1	1

Figure 8: Distribution of Draft Numbers after Month 2.

This process continues through the next three months, every month giving the vendor exactly one number from each level and never providing the number on a day previously given for that draft level. At the end of month 5, equal to 105 working days, the vendor will have received numbers from all levels on all five working days of the week, exactly once. This is shown below in Figure 9.

	Group	Mon	Tue	Wed	Thu	Fri
month 5	1-12	1	1	1	1	1
	13-24	1	1	1	1	1
	25-36	1	1	1	1	1
	37-48	1	1	1	1	1
	49-60	1	1	1	1	1
	61-72	1	1	1	1	1
	73-84	1	1	1	1	1
	85-96	1	1	1	1	1
	97-108	1	1	1	1	1
	109-120	1	1	1	1	1
	121-132	1	1	1	1	1
	133-144	1	1	1	1	1
	145-156	1	1	1	1	1
	157-168	1	1	1	1	1
	169-180	1	1	1	1	1
	181-192	1	1	1	1	1
	193-204	1	1	1	1	1
	205-216	1	1	1	1	1
	217-228	1	1	1	1	1
	229-240	1	1	1	1	1
	241-252	1	1	1	1	1

Figure 9: Distribution of Draft Numbers after Month 5.

The fairness cycle thus finishes at the end of month five, or 105 days, after which a new wheel would need to be generated.

In addition, because the mechanics of the wheel algorithm means all vendors are visiting the same numbers on the wheel at different times, and coupled with the fact that there are only 21 levels of numbers, once the 105 day cycle is completed, all vendors will have been exposed to the same behavior from the wheel regardless of their starting position on the wheel,. The proposed algorithm therefore offers perfect fairness in providing chances to receive equally valued draft numbers to all vendors given a cycle length of 105 days.

2.3. Simulation

In order to further illustrate the practical outcome of the proposed algorithm, our team performed a simulation. Through this simulation, 17 vendors were placed on the wheel, and were assumed to input the popular preference choices identified through data analysis. After receiving the draft numbers and depending on the availability of the locations given the number of vendors ahead of each vendor every day, a similar analysis to what was described in section 1 was performed to display the frequency of receiving first, second and third preferences by the 17 vendors.

The results indicated that, at all times, there would be a difference of no more than one in the number of times different vendors received their first , second and third preferences. Figure 10 displays the frequency of receiving the first three preferences over five months for a sample of 17 vendors.

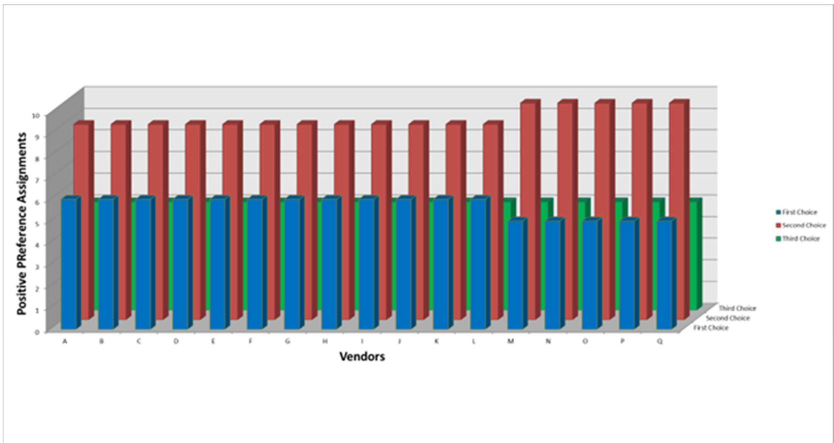

Figure 10: FTLADS Assignment Distributions

2.4. Comparison with the current system

In summary, through validation and simulation, we demonstrated that the proposed system is mathematically guaranteed to give equitable assignments after 5 months in the worst case scenario; that is if all vendors always asked for the same assignments. The proposed system may, however, provide equitable assignments sooner, should vendors have different rankings for locations. Perhaps even more impressively, while recent data on current system behavior indicated a lack of fairness after 8 months, the simulated results suggest perfect fairness is achievable after 5 months using the proposed FTLADS algorithm.