Holistic Reality Examination on Practical Challenges in A Mobile CrowdSensing Application

Xintong Song*, Fan Ye†, Xiaoming Li* and Yuanyuan Yang†
*School of Electronics Engineering and Computer Science, Peking University, Beijing
†Department of Electrical and Computer Engineering, Stony Brook University, New York
songxintong@pku.edu.cn, fan.ye@stonybrook.edu, lxm@pku.edu.cn, yuanyuan.yang@stonybrook.edu

***Abstract*—Despite significant research efforts and great advances on Mobile CrowdSensing (MCS), building MCS applications remains difficult. In this paper, we develop and run Dining Halls on Live (DHOL), a campus dining population density monitoring system over several months. We make a holistic reality examination, discover key technical and practical difficulties, develop effective solutions and share our experiences and insights. We find two main obstacles on data fusion and incentive design: insufficient data quantity/quality and "irrational" user behavior. We develop effective methods by combining historical and real time data, and allocating a given budget among users to address them. We also conduct a detailed user survey to identify reasons behind interesting discoveries, important practical difficulties in acquiring sufficient users and location data, and share our experiences dealing with them. Our main insight is that insufficient data quantity/quality and "irrational" user behavior demand practical yet effective data fusion and incentive mechanisms, and one must provide values to users to acquire and retain a large user base.**

I. Introduction

As smartphones embedded with various sensors penetrate our lives, the emerging *Mobile CrowdSensing (MCS)* [1] paradigm has become popular for building applications. They leverage large numbers of mobile devices from users to collect data useful for different purposes. A typical example is real time road traffic such as Waze [2], where drivers proactively report and share data about road conditions including traffic jams, police traps and accidents. There has been significant research on different aspects on the MCS paradigm. Among others, two important issues are: data fusion algorithms that generate desired information from raw sensing samples [3]–[5]; incentives based on economic tools and models (e.g., game theory, reverse auction) [6]–[8].

Despite these significant efforts and great advances, building effective, large-scale MCS applications remains difficult. There have not been many successful academic cases. Most existing research work tends to each focus on a specific, standalone aspect, and many elegant results require certain assumptions (e.g., the "economic man") for mathematically provable properties. There has been relatively much less work on a *holistic reality-examination*: to build and operate a large-scale mobile crowdsensing application, what are the most difficult technical problems in reality, how they might be addressed, and how they impact the fundamental assumptions underlying MCS research?

In this paper, we conduct a study to answer the above questions. In particular, we build and operate *Dining Halls on Live (DHOL)*, an MCS-based population density monitoring system for the 10 dining halls in Peking University. During several months' of investigation, we run into problems such as insufficient amounts of users and data, practical financial constraints by our shallow pockets, and unexpected observations such as "irrational" user behavior. We devise practical and effective solutions to solve these problems. After weeks of system operation, we conduct a comprehensive user survey to gain better understanding of users' perspective. Among others, we find that instantaneous, instinctive decision-making dominates rational calculation. Thus users are not necessarily "rational" on data contribution, which is fundamentally contrary to popular assumptions made in many incentive study [6]–[9].

We make the following contributions:

- We find that lack of sufficient data quality and quantity is common. We propose two methods that leverage data characteristics such as historical similarities. They can greatly reduce population density estimation errors even when user reports are very sparse and contain significant noises.
- We find that users are impatient and can be "irrational" when they contribute data. They make instantaneous and instinctive decisions in a fraction of a second, without any careful thinking. Thus the "economic man" model that assumes a rational, calculating person is not realistic. We compare ranking and lottery, two practical and effective reward allocation methods, and show they outperform mechanisms based on the "economic man" model.
- We conduct a comprehensive user study to discover reasons behind observed phenomena, and find practical problems significantly impacting the success of MCS applications: acquiring a large user base and sparse location data. We explain how we handle such issues, and discuss the lessons and implications for MCS application designers.

The rest of this paper is organized as follows: We present an overview of the DHOL application and challenges in Section II. In Section III and Section IV, we discuss the technical problems, our solutions and discoveries in data fusion and incentive mechanisms respectively. We present other practical problems we met in Section V. We discuss limitations of this work and plans for future works in Section VI. Finally we discuss related work in Section VII and conclude in Section VIII.

978-1-5090-1328-9/16/$31.00

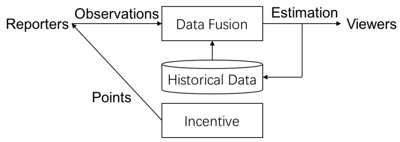

Fig. 1. Architecture of Dining Halls on Live (DHOL).

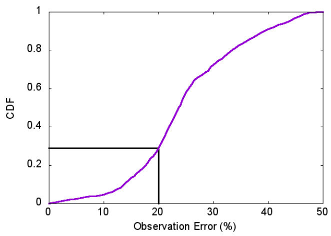

Fig. 2. CDF of observation errors show only 28.9% data have errors less than 20%. Most contain significant errors.

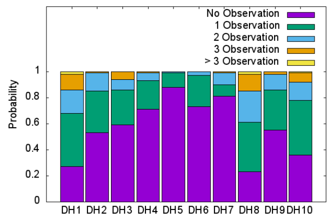

Fig. 3. Probability of amount of reports collected in every 15 minutes. In 95.1% 15-minute time windows there are less than 3 reported observations.

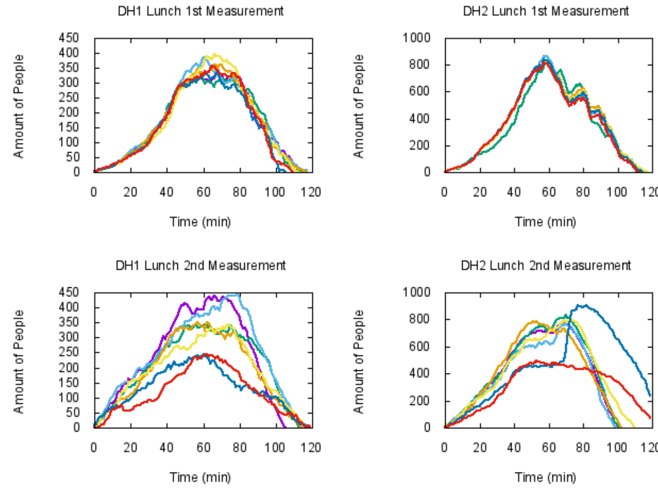

Fig. 4. Two population measurements for dining halls, with several months in between. Each of the 7 curves represents population on one day of the week. The trends are similar within one week, but differ a lot over longer time scale.

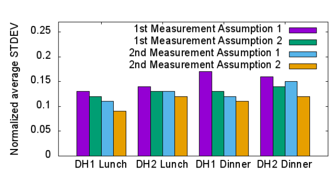

Fig. 5. Validation of the two assumptions in the two data fusion methods using the measured population data.

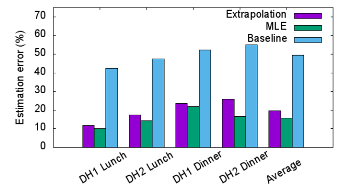

Fig. 6. Estimation errors of different data fusion methods. Both extrapolation and MLE have less errors than static history.

II. DHOL Overview and Challenges

Dining Halls on Live (DHOL) is a mobile crowdsensing based application that monitors the real time population of 10 dining halls in Peking University. As shown in Figure 1, DHOL collects and fuses crowdedness *observations* (a score between 0 and 5) provided by *reporters* to generate real time population *estimation*, which is displayed to *viewers*. We also leverage *historical data* on past crowdedness. The *data fusion* module estimates the real time crowdedness based on both real time observations and historical data. Reporters get *points*, which are accumulated and translated into real cash rewards weekly. The amount of points to be rewarded for each hall at any particular time is decided by the *incentive* module and shown on the UI, so users can decide whether that is enough to motivate them to report. To study how strongly the point amount impacts user decision, we set the amount randomly and observe how much contribution it ensues. [1]

We implement DHOL on Android. The experiment data were collected for 6 weeks after the application is released, during which we acquired 222 users and collected 2,355 observations. We also measured the ground truth population by manually counting people entering and exiting dining halls during lunch and dinner hours. The measurement was conducted twice with several months in between, each lasted for one week. The first was before the application design to give us some inspirations, while the second was in the last week of the 6-week experiment to provide ground truth for performance evaluation. After the 6-week experiment, we conducted a user survey on 103 participants, asking people of their experiences and habits using DHOL, and opinions about incentives.

The two key DHOL components are data fusion and incentive, both are common, important issues for general MCS applications. They both face important technical challenges: 1) The number of observations collected is usually small; such limited sample size cannot guarantee accurate population estimation. The most recent reported observation is usually a few minutes in the past, and random noises are common because different reporters may score the same situation differently. The data fusion module must still generate accurate estimations with such data. 2) Given the limited user base, proper incentive mechanisms that encourage user contribution are critical. Otherwise, the system cannot continuously obtain user data to sustain the operation. However, our financial budgets are limited; we cannot afford to give out large size rewards. The incentive module has to be effective under limited budget. 3) We also have to deal with several important practical problems, such as aquiring and retaining enough users and dealing with undesirable conditions (e.g., lack of Wi-Fi/cellular coverage or GPS signals). These practical problems are not unique to DHOL, but general to other MCS applications as well.

[1]The amount can be based on how urgently the system needs real time observations, e.g., more points for a hall the system knows little.

III. Data Fusion

A. Observations on Data

In DHOL, the data fusion module takes user observations to estimate real time population of dining halls, represented as a score of crowdedness between 0 and 5. During experiments, we find that the observations are noisy and insufficient. There exist quite some errors in user observations (Figure 2). 71.1% of reported observations contain errors larger than 20%. The amount of reported observations is also insufficient (Figure 3). Even in the two dining halls with the most reports (DH1 and DH2), less than 3 reports are collected in 95.1% 15-minute time windows while some time windows have no report.

The errors arise from several sources. Because different individuals have different subjective senses of crowdedness, reporters may score the same situation differently. Besides, some halls are very large that different areas might have different crowdedness. To reduce such noises and make accurate estimations, one common approach is to conduct statistical processing (e.g., average) over large amounts of data. However, insufficient data make direct application of such statistical processing difficult.

Thus we have to design a data fusion method that still makes reasonable estimations under noisy, insufficient data. We got our inspiration by measuring real data and observing their characteristics. Figure 4 shows the raw population for the two measurement periods. We find that for the same dining hall, the trends of population variation over time are similar within one week.[2] However, the trends change significantly over longer time scale. This is caused by larger changes that impact people's dining (e.g., different class schedules between the two semesters, winter vs. summer weather). Besides, sometimes there are irregular outliers due to occasional factors. For example, the sharp increase for a 2nd measurement DH2 curve was caused by hundreds of attendants to an academic conference on Saturday.

B. Similarity based Estimation

The above results show that we can use historical population to make up for insufficient real time data in short term. The significant differences over long terms mean that real time data are indispensable to capture long-term variations. So we combine real time and historical data for the estimation. We propose and compare two methods, one based on extrapolation and the other based on maximum likelihood estimation.

1) Extrapolation method: We use a weighted average of the historical data $h(t)$ and all observation data in the past 15-minutes time window $[t-15, t]$ as the estimation.

We make a very simple assumption: the ratio between real time population ($r(t)$) and historical population remains the same within a time window:

$$\frac{r(t_1)}{h(t_1)} \approx \frac{r(t_2)}{h(t_2)}, \qquad \forall t_1, t_2, |t_1 - t_2| \leq 15min \tag{1}$$

This assumption represents the current as a linear scaling of the past. We verify it (hereafter referred to as Assumption 1) using our ground truth measurement data. We use the population data of each day as real time data, and approximate historical data with the average of the other days in the same week. We use the normalized average standard deviation as the metric. For each 15-minute time window, we scale the real time data with different ratios for the minimum deviation. Outliers are excluded. The results (Figure 5) show that the deviations for both dining halls are less than 0.017, for both lunch and dinner. Considering the significant errors in observation data, the small deviations show the assumption holds very well.

[2]Two curves in the second measurements show significant differences. Those are measured on the weekend right before the final exams, which might have caused different behaviors of students.

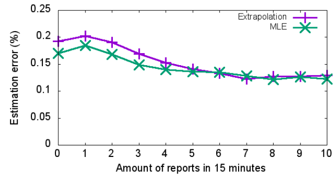

Fig. 7. Impact on performance of data fusion algorithms of the amount of reports collected in the 15 minutes before the estimation taking place.

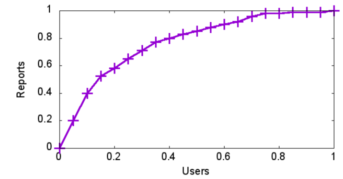

Fig. 8. CDF of users contributing observation reports.

Based on Assumption 1, for each recent reported observation o_i that is reported at time t_i, we have a respective estimate $r_i(t)$ for the current population on time t:

$$r_i(t) \approx h(t)\frac{o_i}{h(t_i)} \tag{2}$$

We then compute a weighted average of $h(t)$ and all the $r_i(t)$s as the estimation $r(t)$ at time t. The weight of each observation is 1, while the weight of historical data is w_h. In our evaluation we find that the estimation error is minimized when w_h is set to around 3.

2) MLE-based method: We use *maximum likelihood estimation* to estimate dining halls' real time population. Letting $r(t)$ denote the current population at time t, we assume:

$$r(t) = \alpha \cdot h(t) + \beta \tag{3}$$

where α is the same scaling factor while β is the constant difference, both are invariant within the same 15 minute time window.

Similarly to how we validate Assumption 1, we do the same to validate this assumption (hereafter referred to as Assumption 2) using measured data. We find the α and β in Equation (3) that minimizes the deviation between the observation and historical data. Figure 5 shows that the normalized average standard deviation is even smaller than those of Assumption 1.

We consider each reported observation o_i as a sample of $r(t)$. We use MLE to find the most likely value of α and β. The likelihood function is defined as:

$$L = -\sum_{o_i \in O} (o_i - r(t_i))^2 - \lambda_1(\alpha - 1)^2 - \lambda_2\beta^2 \tag{4}$$

Finally we use EM [10] algorithm to find the values of α and β that maximizes L. Then we use Equation (3) to estimate the current population $r(t)$. According to our evaluation, the best performance is achieved when λ_1 and λ_2 are set to 3,000 and 0.05 respectively.

3) Results: We compare the estimation error of the two methods. The error is the difference between the estimated crowdedness and the ground truth, normalized by the maximum score 5. We also present a baseline method without historical data, which directly uses the most recent reported observation as the estimation. We can see (Figure 6) that the baseline has the largest error of 49.4% (averaged both dining halls and both mealtime), while the Extrapolation and MLE methods have only 19.6% and 15.8% respectively.

We present how the number of reports in the past 15 minute time window affects the estimation error (Fig. 7). We find that as more reports are gathered, the estimation error decreases but stabilizes after 5 reports. This shows that too many reports have diminishing returns on improving accuracy. The increase in error from zero to one report is because that a single individual's sense of crowdedness may have a large bias, resulting in larger error. As more reports are available, such biases tend to cancel out each other, thus less estimation errors.

To summarize, we find two methods of extrapolation and MLE can make accurate predictions despite noisy and insufficient data. Through this case study, we validate the general approach of exploiting the trend in historical data. Thus one can use combine history and real time data with different weights for much improved estimation.

IV. Incentive Mechanism Design

A. Reverse Auction Mechanism

A significant portion of research on incentive mechanism of MCS has produced excellent results on incentive mechanisms based on *reverse auction* [6]–[8], a model borrowed from economics. It has two stages. First, the platform publishes multiple sensing tasks to users. Users decide which tasks they can conduct, at what desired rewards, and notify the platform of their bids. It usually takes certain time for all users to finish such decisions. Next, after receiving all users' bids, the platform selects a set of users that covers all the tasks with the minimal total rewards. The platform then notifies these winners, pays them the rewards and waits for their data upload.

Despite many theoretically provable, elegant features, reverse auction models requires several conditions: 1) The platform must have significant budget. The platform doesn't have control over total budget, and there is no constraint on how much rewards users can ask. 2) There must exist large numbers of users so the platform has a pool to choose alternative bids providing the same data at different prices. 3) Users must be patient, rational and tolerate. They will spend time carefully thinking through what data to contribute and what price to ask. They tolerate risks that they may not be choosen and paid by the platform after such careful decision-making. Unfortunately, we find these assumptions do not hold in reality, especially for DHOL.

B. Practical and Effective Incentive

We have to devise a practical and effective incentive mechanism for DHOL. It should achieve three goals: 1) We should be able to set a total budget such that rewards would not exceed this limit, regardless of users' data contribution behavior. 2) It should accept all users' data without rejecting anyone who is willing to contribute. 3) The overall process should be simple and fast to suit the instinctive decision-making. The platform cannot keep users waiting. Once a user submits data, he should get a return immediately.

Our incentive design uses "points" that can be converted into real cash. The system shows how many points to reward a report for each dining hall. After seeing the points on the smartphone UI, a user that decides to report will immediately get the points deposited into his account. This ensures all users' contributions are accepted, and the return comes right after the user contribution. To control the total budget, we distribute the cash among users weekly using two rules. In the *ranking* rule, we select the users with the most amounts of accumulated points; in the *lottery* rule, we randomly select users with a probability proportional to their accumulated points. Then we give each of the selected users an equal amount of cash.

Given a total budget, such distribution divides it among winning users. Thus the total rewards are always the same as and never exceeds the budget. The two rules are designed under different hypotheses about user contribution behavior. The ranking rule favors top contributors, assuming that a small number of core contributors can provide sufficient data for the whole system. The lottery rule emphasizes wider participation. It gives any user some chance to win cash, as long as he makes at least one report.

C. Discoveries

Ranking Rule Excels We first compare the results of financial distribution of the two rules. During the 6 weeks' experiment, the ranking rule had 60 winning slots and 22 users were rewarded, each winning 2.73 times on average. This indicates that some users frequently contribute a lot of data and make to the top 10. In the lottery rule, 51 users won the 60 slots, which more than doubles the number of winners in ranking, showing that the lottery rule indeed distributes rewards among much wider user base.

To further examine whether a "20-80" split exists in DOHL, we plot the CDF of user contributions (Fig. 8). Of all the 2, 355 reports, about 60% come from only 20% of users. This indicates that there indeed exists a core contributor group that produces most of the reports. Our survey also confirms the discovery: about half of the users report 1-2 times a week, 1/5 report 3-5 times, 1/10 report 6-10 times and a tiny 3.70% report more than 10 times, while 13% of users never report data. Since the bottom 50% users provide only 20% reports, this suggests that a strategy that favors a core contributor group can keep the system running. Thus ranking is a more sensible option.

Comparison to Reverse Auction We also compare our methods to those based on reverse auction, using both survey and simulation based on parameters from survey. We find that the majority (78.64%) of people prefer our methods because of the speed and certainty. Users also comment that ours are much simpler and avoids the waiting time.

Next we survey the desired weekly (DHOL) and per-report (reverse auction) rewards. 27% participants expect \$0.80 weekly rewards under DHOL's ranking rules, 16% \$1.60, 18% \$3.20, 20% \$8.00 and 18% \$16.00. Our \$0.80 reward thus meets the expectation of 27% people and more budget is needed to obtain data from the rest. For reverse auction, 29% participants ask less than \$0.02 for each report, 16% \$0.03, 12% \$0.08, 11% \$0.16 and 33% equal to or more than \$0.80. Many people (29%, 33%) have either a very low (\$0.02) or very

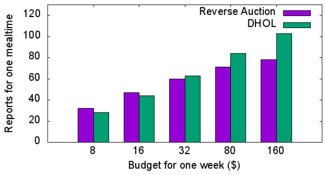

Fig. 9. Number of reports collected under different total budgets using demanded weekly DHOL rewards and per-report reverse auction rewards from user survey.

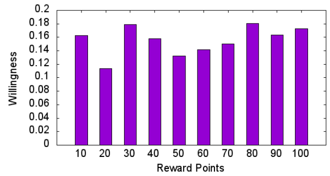

Fig. 10. Relationship between users' willingness to provide reports and the amount of rewards. We define user willingness as the fraction of users who did report.

high ($0.80) price. We conjecture they represent two opposite mentalities: those focusing on winning ask a low price to increase the chance, those focusing on minimizing losses (i.e., "wasted bidding efforts") ask a high price to offset such losses.

Figure 9 compares how many reports we can collect in one mealtime using reverse auction and DHOL's ranking rule. We assume the 103 survey participants are potential reporters, and the minimum rewards needed to make them provide reports are exactly the same as our survey result described above. For reverse auction, we divide the weekly budget among 14 meals (lunch/dinner for 7 days). Then we count how many people we can afford to pay with a greedy method, starting from those with the lowest prices, until the remaining budget cannot pay one more user. For DHOL, we divide the weekly budget by 10 to get the financial reward for those make to the top 10 list (ranking pays only the top 10 contributors). Then we count the number of people whose expectations are equal to or below that reward.

We can see that: 1) when the total budget is small ($16.00 or less), reverse auction collects a little more reports than ranking. Because many (29%, 16%) people ask a small price ($0.02, $0.03), a small budget can still satisfy such users. For ranking, the weekly reward has to meet much higher weekly bars (27% $0.80, 16% $1.60). With a small budget one cannot attract those asking more, thus the slightly less reports. 2) As the total budget grows, ranking shows more and more gain compared to reverse auction. This is because reverse auction needs to pay for all the collected data. As people ask for exponentially higher prices, it requires more and more budget, especially for those asking the highest (33% $0.80). While ranking pays only the top contributors; it gets data from all whose weekly expectations are met but does not need to pay those not in top 10. The results suggest that our ranking rule can attract more reports under DHOL's real user reward expectations.

Irrational User Behavior Our survey data also reveal the lack of "rational" calculation in users' decision to contribute data. From our user study (Tabel I 1-2), we find there is no strong correlation between the reward points and users' contribution willingness. 1) More than half of users seldom or never check the reward points before reporting, 31% sometimes check it, only 17% always check before reporting. 2) About 2/3 people say the points have nothing or a little effect on their decision to provide data; the points have some effects for only 1/3 people, in which only 5% say the points strongly affect their decisions. The above discoveries indicate that there is no

TABLE I
SELECTED USER STUDY RESULTS.

Question	Answer	#	%
1. Do you check the reward points before reporting?	Always	17	16.50%
	Sometimes	32	31.07%
	Seldom	44	42.72%
	Never	10	9.71%
2. How is your decision to provide data affected by the reward points?	Strongly	5	4.85%
	Some	33	32.04%
	A little	31	30.10%
	Not at all	34	33.01%
3. How important crwodedness is to you when deciding which dining hall to go?	Very important	31	30.01%
	Important	49	47.57%
	Normal	15	14.56%
	Doesn't affect	8	7.77%
4. How helpful DHOL is to you?	Very helpful	12	22.22%
	Helpful	25	46.30%
	Not very helpful	15	27.78%
	Not helpful at all	2	3.70%
5. How often is your phone's GPS turned on?	Always	24	23.30%
	Most time	39	37.86%
	Only when needed	27	26.21%
	Never	12	11.65%

"rational" calculation for the majority of users.

We further validate the relationship between the amount of reward points and user willingness to provide reports. We record how many users have viewed the points for reports (shown on the app UI) and how many chose to send reports after viewing. We define user willingness as the fraction of users who did report. Figure 10 shows that the user willingness is not strongly correlated to the amount of rewards. The willingness fluctuates between 15-20% as the points increase from 10-100, without any obvious increase. The points amount does not affect which dining halls people choose to go either. The above shows that the amount of points has little impact on user behavior, further validating the reality of "irrational" users.

In summary, the irrational user behavior that breaks the basic assumptions for many existing incentive mechanisms are wide spread among MCS applications, and our ranking rule is not just simple, practical and shown to outperform representative economic man based mechanism.

V. Practical Problems

1) Acquiring Large User Base: When launching an MCS application, the top priority is to acquire a large user base as quickly as possible. Many MCS applications work only if sufficient numbers of users exist; otherwise the lack of data makes the application useless. DHOL relies on users for both real time and historical data. Only with enough users can all the dining halls be covered most of the time. We believe the key to acquire and retain large user base is that the application itself has to provide valuable, indispensable functions for users. DHOL does provide a valuable service. As shown in Tabel I 3-4, our survey finds that 3/4 participants think the crowdedness level is an important or very important factor in their decisions where to go, while only 8% think it hardly affects their decisions. So the need to know crowdedness level is common among people. We also find that about 70% users feel DHOL is very helpful or helpful. Thus DOHL's service is also effective. As a result, DHOL achieved acquiring 222 users and collecting 2, 355 reports in only 6 weeks.

In summary, we believe that *satisfying users' strong, essential needs* is the key to acquiring and retaining large user base. We have to ask ourselves what value we can bring to them, and what fundamental needs our app can satisfy them.

2) Caveats on Location Data: Locations of where data are collected are critical in MCS applications. In DHOL, location data can be used to check from which dining hall a crowdedness score is reported. However, we find that recording user location at fine spatial and temporal granularity is far from trivial. The phone might fail to obtain a location because of the poor indoor GPS signal and WiFi/cellular coverage, or forbidden by the user for either pravicy concerns or saving battery life. During the 6 weeks' experiments, we collected 4,290 location entries from 86 users, each consisting of a pair of GPS coordinates and a timestamp. This is only slightly over 1/3 of the 222 users. According to our user study (Table I 5), about 26% users turn on the GPS only when they need it, 12% never turn it on, 23.30% and 37.86% users always or at most times turn on their phone's GPS respectively.

In summary, we cannot take fine grained location data for granted. Functions building on fine-grained location traces may not work at all under such conditions. An MCS application has to be prepared to handle sparse location availability.

VI. Discussion

We discuss a few issues related to the complexities in building and operating applications in real environments. We had monetary and manpower resources enough to run only one application. It takes a lot of resources to create the application, promote it, recruit users, maintain the application over long periods of time. In the future we want to create more applications and further validate our observations.

Because we do not have any control over real events and environments, many experiments can not be easily replicated. For example, the weather changes, sudden surge of visitors, can greatly impact the population but there is no way we can repeat the same weather or event. We plan to refine and run the application as a sustainable service to the campus, so that we can re-experience those events and conditions to confirm our discoveries.

VII. Related Work

We briefly discuss related work in data fusion and incentive mechanism due to space limit. More comprehensive survey can be found in [1].

There are several sophisticated techniques dealing with insufficient data quality and quantity, e.g., Compressed Sensing [11] that reconstructs a signal with a small number of samples. However, the technique requires knowledge of the signal's sparsity and centralized control of when to sample. In MCS, although some knowledge of the signal can be learned from historic data, the application can not control when, which users can provide what samples. Adapting such techniques to MCS remains a great challenge.

Much excellent work has been done on incentive mechanisms. Yang et al. [6] propose both platform centric and user centric incentive models, and the latter is based on reverse auction model. The two-stage process is designed to minimize the platform's total cost and it is widely adopted in other work [7], [8]. Zhu et al. [7] consider dynamic smartphones and random arrivals in reverse auction model. Li et al. [8] protect users' account information in reverse auctions.

VIII. Conclusion

In this paper, we investigate two important technical problems of data fusion and incentive common to building and operating mobile crowdsensing applications. Through a case study of *Dining Halls on Live (DHOL)* that monitors real time population density of dining halls in Peking University, we propose effective solutions for data fusion and incentive design by exploiting patterns in historical data and focusing on core contributor group. We also address three important practical problems of acquiring large user base, coarse-grained location data and intermittent Internet connectivity. Based on 6 weeks of actual deployment and subsequent data analysis, user survey, we make several important discoveries, including insufficient data and irrational user behavior, which greatly impact fundamental assumptions in existing MCS work.

Acknowledgement

This material is based upon work supported by the National Science Foundation under Grant No. CSR-1513719 Any opinions, findings, and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the National Science Foundation.

References

[1] R. K. Ganti, F. Ye, and H. Lei, "Mobile crowdsensing: current state and future challenges," *Communications Magazine, IEEE*, vol. 49, no. 11, pp. 32–39, 2011.

[2] "Waze," https://www.waze.com.

[3] R. Gao, M. Zhao, T. Ye, F. Ye, Y. Wang, K. Bian, T. Wang, and X. Li, "Jigsaw: Indoor floor plan reconstruction via mobile crowdsensing," in *Proceedings of the 20th annual international conference on Mobile computing and networking*. ACM, 2014, pp. 249–260.

[4] X. Guo, E. C. Chan, C. Liu, K. Wu, S. Liu, and L. M. Ni, "Shopprofiler: Profiling shops with crowdsourcing data," in *INFOCOM, 2014 Proceedings IEEE*. IEEE, 2014, pp. 1240–1248.

[5] Y. Wang, J. Wang, and X. Zhang, "Qtime: A queuing-time notification system based on participatory sensing data," in *Proceedings of the 2013 IEEE 37th Annual Computer Software and Applications Conference*. IEEE Computer Society, 2013, pp. 770–777.

[6] D. Yang, G. Xue, X. Fang, and J. Tang, "Crowdsourcing to smartphones: incentive mechanism design for mobile phone sensing," in *Proceedings of the 18th annual international conference on Mobile computing and networking*. ACM, 2012, pp. 173–184.

[7] Y. Zhu, Q. Zhang, H. Zhu, J. Yu, J. Cao, and L. M. Ni, "Towards truthful mechanisms for mobile crowdsourcing with dynamic smartphones," in *Distributed Computing Systems (ICDCS), 2014 IEEE 34th International Conference on*. IEEE, 2014, pp. 11–20.

[8] Q. Li and G. Cao, "Providing privacy-aware incentives for mobile sensing," in *Pervasive Computing and Communications (PerCom), 2013 IEEE International Conference on*. IEEE, 2013, pp. 76–84.

[9] L. G. Jaimes, I. Vergara-Laurens, and M. A. Labrador, "A location-based incentive mechanism for participatory sensing systems with budget constraints," in *Pervasive Computing and Communications (PerCom), 2012 IEEE International Conference on*. IEEE, 2012, pp. 103–108.

[10] A. P. Dempster, N. M. Laird, and D. B. Rubin, "Maximum likelihood from incomplete data via the em algorithm," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 1–38, 1977.

[11] "Compressed Sensing," https://en.wikipedia.org/wiki/Compressed_sensing.