Design and Optimization of Spatial Organizations For Context Exchange and Surveillance

Agisilaos-Georgios Ziotopoulos, *Student Member, IEEE*, Gustavo de Veciana *Fellow, IEEE*,

E-mail: `ziotopou,gustavo@ece.utexas.edu`

Abstract

Context-awareness is one of the pillars of ubiquitous computing. Modern environments may have huge quantities of contextual information that is mainly relevant/useful when a user/mobile is situated at particular locations. This paper investigates organizations for exchanging localized context and surveilling the corresponding areas that minimize bandwidth/energy requirements – also implicitly maximizing the 'relevance' of data that is transferred to the users. We propose a simple model where context is associated with spatial regions, cells, and is transferred to/from mobiles as their trajectories traverse the cells. Key elements of the model are the granularity (e.g., area) of the cells and the amount of context associated with each cell. The model uses random tessellations to generate spatial partitions of the space into cells of varying granularity, and an associated context content function to capture possible ways in which contextual data associated with a cell might scale. We discuss a simple taxonomy for such scalings giving a crude idea of how representative applications might behave. Based on this model we analyze optimal flat aggregative and hierarchical organizations for mediating context exchange revealing some fundamental principles relating the optimal granularity of cells to different types of applications' spatial content scaling.

Index Terms

J.9.a Location-dependent and sensitive, C.3.h Ubiquitous Computing, J.9.d Pervasive Computing and C.2.8.b Architectures

I. Introduction

Context-awareness refers to the ability of applications to recognize the environment in which they are executing and is considered a key pre-requisite for ubiquitous computing. Its wide applicability has led to many radically different, but equally valid, manifestations: e.g., knowledge of which services are available in a space, the sensor values describing the local environmental conditions, personal data about people present in a room for a social-networking application or a person's shopping preferences in a mall to be used by a targeted advertisement application, etc., can all be considered as context. Spatial context content e.g., knowledge of which gas stations are in the neighborhood or a person's shopping preferences in a mall to be used by a targeted advertisement application, is a special type of contextual information that exhibits strong locality. Exchanging this type of data has become increasingly prevalent in ubiquitous computing. Thus, we believe that optimizing such contextual exchanges between mobile users and applications is a problem of primary importance. Our focus in the sequel is on contextual information that is encoded, stored and available for use, but *tied* to an area, i.e., is relevant only to users/applications at certain spatial locations.

How often should a mobile user interact with its environment to exchange contextual information? Is it preferable to frequently exchange small amounts of context or bulk amounts of context less often? Should location information come from a shared infrastructure mechanism or is it better for mobile users to individually calculate their location? These are the types of questions that we aim to address in this paper. The answers will depend primarily on the scaling characteristics of the information exchanged. Intuitively, one might expect that e.g. temperature measurements, taken from adjacent sensors will be highly correlated. Therefore, temperature contextual information need not be acquired from all sensors but only from a selected subset of them appropriately distanced from each other. Additionally, the mechanism used to perform the exchange can complicate the picture, e.g. the use of a wireless protocol adding high packetization overheads to each message exchanged can impact the most efficient ways to realize context exchange.

Clearly, the organization of contextual exchanges will have a major impact on the performance and the lifetime of a ubiquitous computing system. For example, if users carry laptops, PDA's, phones etc., to access contextual information, energy consumption will be a primary concern. Shorter and less frequent exchanges of context result in reduced energy consumption which in turn will affect the required size of batteries, their cost and the time horizon until pausing for recharging power. One cannot overemphasize the importance of these issues for ubiquitous computing especially as new exciting avenues such as wearable computers and smart-objects emerge, see e.g., [1], [2]. Understanding some of the trade-offs in designing such infrastructure is the goal of this paper.

In order to quantitatively assess the relevant merits of particular architectures for context exchange we propose a simple, but novel model, capturing the salient features of such systems, see Fig. 1. We suppose space is partitioned as a tessellation with each cell corresponding to a region around natural points of interest whose context will be treated as a coherent entity. This is a first-order approximation that can be used as a guide when designing/optimizing spatial organizations for context exchange. Similar approaches have been successfully applied to the problem of network design, see e.g., [3], [4]. We will in the sequel present a simple taxonomy capturing different ways in which the amount of contextual data associated with a cell may be modeled depending on the character of the associated applications.

Our model assumes a mobile traversing cells of the tessellation, a surveillance mechanism that is part of the space's infrastructure notifies the mobile about these events and the mobile in response acquires/transmits the context relevant within each cell. However, each exchange of contextual data incurs a cost in terms of bandwidth or energy plus some overhead. Our goal is to find organizations that minimize such costs. In particular, we wish to determine the granularity that cells should have. For example, inside a mall one could exchange context at a floor, shop or even finer granularity. Thus, on one hand, if we exchange context more frequently from small cells we exchange only the context that we need, but might incur a higher overhead. On the other hand, if we exchange context less frequently from larger cells, a mobile user will download irrelevant context from fine-grain cells it will *not* actually visit in addition to the useful context that comes from cells to be visited, but the overhead is amortized. As such, depending on the manner in which context content scales and the nature of exchange overheads, one might expect to find optimal cell granularities.

Related work. Context acquisition is a problem of recognized importance in the ubiquitous computing community as well as in the networking community, see e.g., [5],[6] and [7]. However, to the best of our knowledge this is the first paper to focus on quantitative, albeit simplified, arguments for spatial context exchange using formal tools. Geometrical modeling of spaces and their use in ubiquitous computing is not a new idea, see e.g., [8]. The key difference between this work and others focusing on geometrical modeling is that we use stochastic geometry to capture typical architectures with a few parameters to infer results applicable to a large variety of scenarios.

Context aggregation is widely recognized as an efficient policy for coping with the scalability issues challenging ubiquitous systems, see e.g., [9] and [10]. In [9] context aggregation is modeled with the help of a graph connecting context sources, context operators and context sinks. Context information flows uni-directionally from sources to sinks along the edges of the graph in the form of context update events. Key insights drawn from the lessons learned from deploying Solar, a concrete realization of their proposed architecture, are provided. In contrast, our work takes a more detailed view on the scalability characteristics of context to assess the limitations of aggregative architectures and allows for modeling bi-directional context exchanges between mobile users and applications. Our work shares with [11] a concern with context aggregation for resource limited mobile devices. They favor a distributed approach to context aggregation as opposed to our approach, which uses different aggregation scales. The work in [12] and ours share a concern with scaling issues of context but we choose to focus on studying context exchange while [12] focuses on the quality of the acquired contextual information. [13] describes the use of context for improving network performance and makes a similar assumption as this paper, that context data is exchanged by the same mechanisms used for communication.

Other research in the area of context acquisition is [14],[15], [16],[17]

Contributions and organization. The key contributions and organization of this paper might be summarized as follows. In Section II we formally propose a simple first-order stochastic geometric model, based on cells from random Voronoi tessellation, for a spatial organization of context exchanges to/from mobile users/terminals. In order to argue quantitatively about the relevant merits of different architectures, we also propose a taxonomy for how context content scales with the area of a cell for various applications.

In Sections III and IV we consider a flat organization, which aggregates contextual data via cells, and exchanges the associated data as a batch when a mobile traverses a cell. For this scenario we show that when the context content of a cell scales slower than the square root of its area then aggregative organizations increasingly reduce overall costs, i.e.,bandwidth/energy. However, when the context content of a cell grows faster than the latter then it is either not beneficial or there is an optimal granularity for aggregation which depends critically on the exchange overheads and the characteristics of context content scaling. In other words, the case where context content scales roughly as the square root of the area, seems to be a critical case in considering optimization of context exchanges to mobile users.

In Section V we consider hierarchical organizations for context exchange. In particular, if the context content of a cell scales sub-linearly with the area, one expects to have contextual redundancy in the space. Thus, it is natural to exchange shared context using a coarse organization of data, while data specific to a given location is exchanged at finer scales. We will show that such hierarchical organizations are indeed beneficial, but once again the benefits relative to aggregative organizations depend critically on the context content scaling characteristics. In Section VI we elaborate on different mechanisms to surveil a space's cells and consider their contribution to the energy costs. The paper concludes with Section VII.

II. Modeling context regions and scaling

Spatial context is usually formed around designated points of interest in the environment e.g. information about art exhibits targeted to visitors of a museum. The corresponding regions formed are hardly ever regular, usually the more the context of a region e.g. the information about an important exhibit, the bigger the size of the region formed around it e.g. art masterpieces are often allocated more space than other exhibits. The resulting partition of the space is very similar to the Voronoi tessellation formed by the points of interest as nuclei. To express the multitude of possible configurations of regions a stochastic approach is needed.

Stochastic geometry, [18], has recently proven to be a useful tool for modeling the architecture and performance of communication networks as well as the role of mobility, see e.g., [3], [4]. The general idea is to develop simple first-order models, i.e., which depend on a few parameters, that capture the salient features of the problem at hand, allowing one to roughly consider optimizing system designs. This is the character of the model we consider below.

A. Model for context regions

We shall start by considering a non-hierarchical, 'flat', partition of the environment into cells. After a user crosses a cell's boundary, an exchange of context takes place. The cell's localized context is transferred to the user and the user's cell-specific context is transferred to the application(s) serving the cell. We model such a partition based on the cells of a Voronoi tessellation induced by a homogeneous Poisson point process on the plane which we very briefly describe next, see additionally [18]. The geometry of spaces found in the real world is far too complex to be described by a single model. We think that homogeneous stochastic Voronoi tessellations form a *reasonable* first-order model controlled by a single parameter that is amenable to optimization.

A Poisson point process Π on the plane with intensity λ is defined as a random set of points $\Pi = \{x_0, x_1, \ldots\}$ such that:

- for every set $S \subset \mathbb{R}^2$ the number of points in it follows a Poisson distribution with rate $\lambda|S|$, where $|S|$ is the area of S;
- and for $S_1, S_2 \in \mathbb{R}^2$ such that $S_1 \cap S_2 = \emptyset$ the number of points in S_1 and S_2 are statistically independent.

The Voronoi tessellation $V(\Pi) = \{V(x_0), V(x_1), \ldots\}$ corresponding to the Poisson point process Π is defined as a collection of closed and convex cells $V(x_i)$ covering the plane where

$$V(x_i) = \{x \in \mathbb{R}^2 \quad | \quad \|x - x_i\|_2 \leq \|x - x_j\|_2, \forall x_i, x_j \in \Pi\}.$$

The intensity λ of the Poisson point process captures with a single parameter the granularity of the cells of a tessellation – the average area of a cell is given by $\frac{1}{\lambda}$. Higher values of λ lead to finer grain cells, while lower values of λ correspond to a tessellation with coarser cells. Additionally, we shall assume that a tessellation induced by a Poisson process Π_f with intensity λ_f models the natural, *finest grain*, spatial organization of context in the environment. We consider a second *independent* Poisson process, Π_a with rate $\lambda_a < \lambda_f$, modeling a coarser *aggregative* view to study the potential benefits of exchanging context from larger contextual cells. For the remainder of the paper we will refer to these tessellations as the 'finest grain' and 'aggregative' tessellations respectively, these are exhibited in Fig. 1.

B. Model for context content of a cell

Each cell of a tessellation is associated with a certain amount of context to be exchanged. In general, this amount may depend on the size and shape of the cell. For example, a cell with bigger area might be expected to have a higher number of services in it. Or, in the case of a library or supermarket, contextual content may depend on the perimeter of the shelves storing books, items, etc.

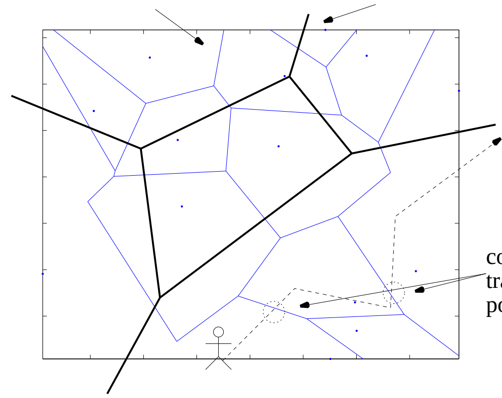

Fig. 1. 'Finest grain' and 'aggregative' Voronoi tessellations modeling contextual spaces.

Definition II.1. ***(Context content function)*** The context content *function* $c : K \rightarrow \mathbb{R}^+$ *where* K *denotes the set of bounded convex sets, models the amount of context associated with a region in the plane. We assume this function is translation invariant.*

Depending on the specifics of the application considered, the amount of context content can refer to a cell's context transferred to a mobile and/or a mobile's cell-specific context that is transferred to the application(s) serving the cell.

Definition II.2. ***(Context scaling)*** *Consider* $A \in K$*, we say a context content function* $c(\cdot)$ *is:*

- additive *iff* $c(A) = \sum c(A_i)$;
- sub-additive *iff* $c(A) < \sum c(A_i)$;
- super-additive *iff* $c(A) > \sum c(A_i)$;

for any partition $A_1, \ldots, A_n \in K$ *of* A*. Note that since* $c(\dot{})$ *is translation invariant, an additive context content function must be proportional to the area of a set.*

Examples of additive, sub-additive and a super-additive context content function are: $|A|$, $\sqrt{|A|}$, and $|A|^2$ where $|\cdot|$ denotes the area of a set.

A sub-additive context content function might arise in situations where there is spatial redundancy in contextual information, e.g., as might be the case when context comes from geographically adjacent sensors or cells that are part of a shared hierarchical structure, e.g., the same building. In this case, shared context can be 'compressed' resulting in a context content function that is sub-additive. An example of an additive context scaling would be a situation where a number of services are spatially deployed, say with intensity 1 service per unit area. Consider the context content of a cell with area $|A|$. If we merely want to collect the status of the services in the cell, e.g., operational or not, the amount of context per region would be proportional to its area $|A|$. However, if contextual information should include both the status and location of each service we will need information proportional to $\log_2(|A|) * |A|$, since to describe the location of a service inside a space of area $|A|$, within a precision of a unit area, we need at least $\log_2(|A|)$ bits. This then is an example of a super-additive context content function.

In practice, for complex environments context content functions may grow arbitrarily with cell size, i.e., they need not fit neatly into the above taxonomy. Consider a company with offices on multiple floors which distributes printers with a density of 1 unit per floor. For a productivity application which needs to know the locations of the available printers, acquiring context at the office level or floor level will produce the same context but acquiring context at the building level will not produce a linear increase in the amount of context for administrative reasons, e.g., printers from offices in other companies will not be available. Therefore, context with respect to area may appear to be constant for small scales and increase sub-linearly for larger scales. Our idealized models for the context content function capture only some basic characteristics of such systems.

For mathematical ease, and to capture a range of possible context content functions we introduce the following assumption.

Assumption II.3. ***(Context content model)*** *The context content of a cell is a function of its shape. The average context content in a typical cell (as seen by a mobile) in the aggregative Poisson Voronoi tessellation* $V(\Pi_a)$ *with associated intensity* λ_a*, is denoted by* $c(V_a)$ *and given by*

$$c(V_a) \triangleq \frac{c(V_f)\lambda_f^\alpha}{\lambda_a^\alpha}, \quad \textit{where } \alpha > 0,$$

and $c(V_f)$ *is another constant interpreted as average context content of a typical cell in the fine grain Poisson Voronoi tessellation* $V(\Pi_f)$ *with intensity* $\lambda_f > \lambda_a$*.*

Note that above we refer to a *typical* cell as a cell seen by a mobile where we will assume in the sequel that mobiles have stationary trajectories independent of the tessellation. Under these circumstances one can show that the area of typical cells is proportional to $\frac{1}{\lambda_a}$, i.e. the area of a randomly sampled cell under the Palm probability. To avoid these considerations we directly model the typical average context content as seen by mobiles.

The polynomial model chosen is continuous at $\lambda = \lambda_f$ and fulfils through a single parameter, α, our stated assumption of expressing various context content scalings that depend on the shape of an average cell i.e. the area scaling as $\frac{1}{\lambda}$, the perimeter or diameter scaling as $\frac{1}{\sqrt{\lambda}}$, etc. In practice, the exact scaling of the context content function can be quite complex but we think it is unlikely to be exponential. Polynomial functions are *reasonable* first order bounds for the context content scalings of most of the use cases we have identified. Inspired by use cases from the environmental sensor domain we observe that there is correlation in spatial context. Use cases from the monitoring, control or social networking domain exhibit combinatorial context scalings. These observations corroborate our choice of polynomial context content scaling.

The parameter α depends on the context content characteristics of a particular space/application. Intuitively, the higher the amount of spatial redundancy in the relevant context, the lower the value of $\alpha < 1$. For $\alpha = 1$ the average amount of context content in a typical cell is proportional to the average area of a cell of the $V(\Pi_a)$ tessellation. More generally, for $\alpha > 0$ the context content growth is proportional to a power of the average area – with cells of higher average area, containing more content, but scaling in different ways. Our results apply to all context content scalings, but the focus is mainly on $\alpha \leq 1$.

Our Poisson-Voronoi model for cells allows for a stochastic amount of context content in each cell through our assumption that context is a function of the shape of each cell. Assumption II.3 models the *average* amount of context content in an aggregative cell and will be used for optimizing our chosen cost function.

C. Mobility model

The time instances at which context exchanges happen depend on the specifics of each application. For example, a ubiquitous application that serves a particular area and operates based on proximity will perform context exchanges as soon as the mobile is within a certain range. A mobile that wishes to acquire the context from an entire space, will exchange the context as soon as it enters the space, see Fig. 1. The intensity of such events depends on the specific characteristics of users' mobility. We shall assume a generic homogeneous model for mobility.

Assumption II.4. ***(User mobility)*** *We assume mobiles are initially distributed as a Poisson process with* *intensity* λ_0*, their motions are* stationary *and* independent *with mean velocity* v*. A user's trajectory is assumed to be sufficiently smooth, i.e., continuous and piecewise differentiable. The context associated with a cell is exchanged when a user crosses a cell boundary.*

Recent advances in mobility modeling suggest that human mobility might follow something akin to Levy-random walks, see e.g. [19]. Our assumption on the users' mobility is fairly generic and acceptable. A more critical concern with the model is the assumption that the mobility patterns are independent of the spatial organization of contextual information. Most likely, mobility and context content would be linked to actual physical structures. This is a simplification required to attempt to study some of the fundamental properties of the problem. Under this assumption one can show the following fact, see e.g., [3].

Fact II.5. *The intensity of cell boundary crossings of a homogeneous Poisson Voronoi tessellation with rate* λ *seen by a typical user moving at mean speed* v *is*

$$\frac{4 * v}{\pi}\sqrt{\lambda} \;\; \textit{crossings/unit time}. \tag{1}$$

Note that the assumption that context exchanges happen when a mobile crosses a cell boundary is not restrictive as long as a context exchange occurs at some point when the mobile is within the cell.

D. Cost model

The nature of the ubiquitous computing paradigm is such that communication will take place via a wireless medium. It is plausible to define the cost associated with an architecture for context exchange based on the bandwidth or energy expended to perform such exchanges. As a *first-order* approximation, both bandwidth and energy, might be roughly proportional to the total amount of context exchanged, including for example, protocol and packetization overheads. The following model captures these salient features.

Assumption II.6. ***(Cost model)*** *The cost to exchange* d *units of contextual data from a cell is*

$$h + K * d. \tag{2}$$

We assume the energy cost for exchanging context is, to a first order, proportional to the amount of data and overhead.

The parameter K can model overheads that are proportional to the amount of data e.g. packet overheads. In the sequel we will assume without loss of generality that $K = 1$. The effect of $K \neq 1$ can be evaluated by scaling the context content function in Assumption II.3. The parameter h can model e.g. fixed protocol overheads, the cost to authenticate with a new cell, the cost to wake-up a terminal from sleep mode, etc.

This is a natural first order assumption, but can be argued more systematically for specific communication architectures for exchanging data, as long as context exchanges occur at random locations relative to the communications access points used to mediate such exchanges.

III. Analysis of Aggregative Tessellations under Additive Context Content Scaling

In this section we explore the benefits of using an aggregative organization to perform bulk context exchanges versus doing this at the finest grain. Recall that these two organizations are modeled via a coarse aggregative tessellation $V(\Pi_a)$ with intensity λ_a and fine grain tessellation $V(\Pi_f)$ with intensity λ_f where $\lambda_a < \lambda_f$. Each time a user/mobile crosses a coarse grain cell in the aggregative tessellation the entire context content associated with the cell is exchanged. Under the fine grain organization, users/mobiles see context exchanges as they cross fine grain cells, and thus see them more often.

In this section we focus on a simple case. We assume that the context content function is additive, and so proportional to the *average* area of an aggregate cell $1/\lambda_a$. In particular we will take $\alpha = 1$ in Assumption II.3 and model the typical context content seen by mobiles at different levels of granularity as

$$c(V_a) = \frac{\lambda_f c(V_f)}{\lambda_a}. \tag{3}$$

The key idea for our analysis is simple. Under our assumption for users' mobility, the intensity of cell boundary crossings, and thus of context exchanges, is proportional to the square root of the intensity of the cells. Each context exchange corresponds to an average cost including overheads and data exchanged. Thus, in the case of the aggregative organization the total cost incurred per unit time is proportional to

$$\sqrt{\lambda_a} * (h + c(V_a))$$

with a similar form for the fine grain case. In order to have cost savings under the aggregative organization versus the fine grain the following inequality must hold

$$\sqrt{\lambda_a} * (h + c(V_a)) < \sqrt{\lambda_f} * (h + c(V_f)) \tag{4}$$

where $c(V_a)$ is related to $c(V_f)$ by Eq. 3. The following result, which is derived in the Appendix, summarizes when aggregation is indeed beneficial.

Theorem III.1. *Under Assumptions II.4,II.6 and an additive context content function, i.e., Assumption II.3 with* $\alpha = 1$*, the aggregative organization with intensity* λ_a *can achieve a reduced cost relative to the fine grain organization with intensity* λ_f *if* $c(V_f) < h$ *and* λ_a *satisfies*

$$\lambda_a \in \left(\lambda_f (\frac{c(V_f)}{h})^2, \lambda_f \right). \tag{5}$$

The optimal rate for the aggregative tessellation, i.e., which minimizes the cost is given by

$$\lambda_{a,opt} = \frac{c(V_f)}{h} \lambda_f. \tag{6}$$

Let x *denote the overhead ratio* $x \triangleq \frac{h}{c(V_f)}$ *then the maximum relative cost reduction is given by*

$$|1 - \frac{\sqrt{\lambda_{a,opt}}(h + c(V_a))}{\sqrt{\lambda_f}(h + c(V_f))}| = 1 - \frac{2\sqrt{x}}{1+x}. \tag{7}$$

We can interpret this result as follows. Only when the average context content per cell $c(V_f)$ of the finest grain tessellation is less than the overhead h, will use of the aggregated context cells be beneficial. This might have been expected since one can amortize overheads through aggregated context exchanges. Note, however, that using increasingly aggregated contextual regions eventually hurts. Indeed there is a lower bound λ_a if cost reductions are to be achieved. Intuitively this lower bound can be explained as follows. If aggregative cells are too large, one will be exchanging context to users that is not actually relevant to them, i.e., associated with spatial regions they will in fact not visit, so finer grain organizations are best. Observe that the optimal rate for the aggregative tessellation is always inside the allowable region since for $\frac{c(V_f)}{h} \in (0, 1)$ we have that $1 > \frac{c(V_f)}{h} > (\frac{c(V_f)}{h})^2$.

Fig. 2 exhibits an example of the results. For $c(V_f) > h$, e.g. $c(V_f) = \frac{4}{3}, h = 1$, the average cost per unit of time for a typical user using aggregative cells always exceeds that of the finest grain organization. For the case where $c(V_f) < h$, e.g. $c(V_f) = \frac{1}{10}, h = 1$, one can see the interval for λ_a in which using an aggregative tessellation is beneficial. The left and right boundaries of the interval as well as the location of the optimal rate of the aggregative tessellation are as predicted by Theorem III.1. The relative cost reduction by using the aggregative tessellation with the specific parameter values used in Fig. 2 is 42.5%.

The results of Theorem IV.1 can be constrained by the limited space resources found in typical mobile devices. Excessively large aggregative organizations may contain too much context content

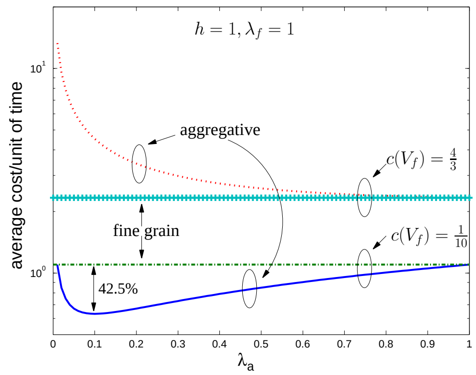

Fig. 2. Context aggregation using aggregative tessellations with additive context scaling.

per cell to be downloaded to a mobile device. A designer trying to define the optimal achievable scale of aggregation should choose the biggest aggregative cell whose context fits in the space provided by the mobiles to be used. This policy is guaranteed by the single mode of the cost function in Theorem IV.1.

A. Selective Context

Up to now we have assumed that each time a user/mobile crosses a cell border of the underlying tessellation a context exchange occurs. However, this need not hold in practice. Indeed a user/application interacting with a fine grain organization could select exactly in which cells/services it has an interest. As a result, the cost associated with exchanges from the finest grain tessellation may be lower. A simple enhancement to our model capturing this phenomenon would be that a context exchange with a fine grain cell occurs only with probability p, where p captures the users' selectivity and only a fraction q of the context of each fine grain cell is acquired. We further assume that p and q are independent from each other.

In that case the following fact holds, for a proof see the Appendix.

Fact III.2. *Under the assumptions of Theorem III.1 and assuming that a mobile interacts with the cells of the fine-grained organization with average probability* $p \leq 1$ *and acquires a fraction* $q \leq 1$ *of a cell's context, a necessary condition for an aggregative organization to achieve savings compared* *to a fine-grained organization is*

$$p * (x + q) > 2 * \sqrt{x} > 2, x > 1 \tag{8}$$

where x is defined as in Theorem III.1.

Observe that for $p = q = 1$ the condition in Fact III.2 is consistent with $x > 1$ as required by Theorem III.1 and ensures that the maximum relative cost in Eq. 7 is less than 1. The optimal scale of aggregation is still given by Theorem III.1. Due to space limitations, we will present the results for the new range over which aggregation is a win, in Section IV for generic $\alpha \neq 1$.

The interpretation of the previous fact is that if the context associated with fine grain cells is of sufficient interest to users one can still benefit from exchanging the aggregated context. Fig. 3 (left) shows the relationship between the probability of interest p and the ratio of the overhead of each exchange to the average amount of context of a typical cell of the finest grain tessellation x, for different values of q. In the $x > 1$ regime where there is more overhead than context associated with a fine-grain cell, the minimum necessary probability of interaction with a fine-grained cell p, diminishes as the overhead increases. Indeed, when the overhead is extremely high, using an aggregative tessellation is the correct strategy and the actual amount of context exchanged does not play a key role. q, the fraction of context of interest in a fine-grained cell, plays a key role. If q is too low, it might be impossible to achieve savings, see the curves above the $p = 1$ line in the invalid region. For large values of x the role of q is diminished. In that case, there is an excessive amount of overhead that dominates the cost and the amount of context transfered does not matter much. Additionally, observe that for $q = 1$ we can always achieve savings through an aggregative tessellation.

Fig. 3 (right) shows the relationship between the minimum fraction of a fine-grain cell's context necessary to be acquired for an aggregative tessellation to be beneficial and the ration of the overhead of each exchange to the average amount of context of a typical cell of the finest grain tessellation x, for different values of p. Observe that as the overhead increases the minimum fraction of context necessary to be acquired drops. Indeed, when the there is an excessive amount of overhead it dominates the cost and the amount of context transfers does not play a key role. As a result for large amounts of overhead, all the curves regardless of the value of p converge. Note, that for low values of overhead it might not be feasible to achieve savings through an aggregative tessellation if the probability of interaction with the fine-grained tessellation is too low, see the curves above the $q = 1$ line in the invalid region. Additionally, observe that for $p = 1$ we can always achieve savings through an aggregative tessellation.

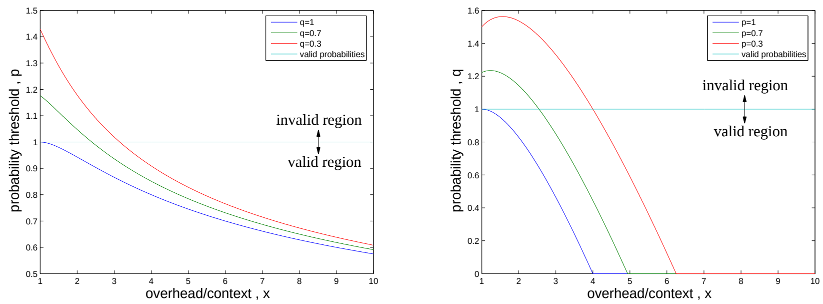

Fig. 3. Minimum probability vs. overhead, to achieve savings using aggregation. Cell interaction probability p (left), context fraction q (right).

B. Dynamic Context

Aggregation pays off in terms of bandwidth/energy consumption but this comes at the expense of the accuracy of highly dynamic data. As aggregative cells become larger and larger, e.g. the readings from highly dynamic sensors provided to a mobile at the time of crossing an aggregative cell will be invalid at the time the mobile reaches the sensors. In practice, a wide class of contextual information is static e.g. a map of the current floor of the mall, or slowly varying e.g. readings from a temperature sensor. A designer trying to decide on the appropriate level of aggregation has to consider the nature of the contextual information as well as the average sojourn time of a mobile through a typical cell. The following fact, serves as a rule of thumb for deciding when aggregation is acceptable for dynamic data.

Fact III.3. *Aggregation is meaningful for acquiring data from sensors that change with frequency f if*

$$f = O(\sqrt{\lambda}v)$$

where v is the average speed of the mobiles.

The proof for this fact relies on the observation, demonstrated in the Appendix/proof of Fact VI.4, that the mean sojourn time of a mobile is proportional to $\frac{1}{v\sqrt{\lambda_a}}$.

IV. Analysis of Aggregative Tessellations under Generalized Context Scaling

In the previous section we assumed that the context content function was additive. However, as discussed in Section II one can consider applications with different scaling characteristics. To capture the more general case we consider context scaling of the form posited in Assumption II.3. In this case, by analogy with Eq. 4 one finds that coarse grained context regions will be beneficial if

$$\sqrt{\lambda_a} * (h + (\frac{\lambda_f}{\lambda_a})^{\alpha} * c(V_f)) < \sqrt{\lambda_f}(h + c(V_f)). \tag{9}$$

The results below are shown in the Appendix and summarize the characteristics of minimum cost organizations for context exchange.

Theorem IV.1. *Under Assumptions II.3, II.4 and II.6, an organization for context exchanges based on an aggregative tessellation with intensity* λ_a *is beneficial if*

- $\alpha < \frac{1}{2}$ *and* $\lambda_a \in (0, \lambda_f)$*. In this case the cost is strictly increasing in* λ_a *thus, the optimal intensity should be as small as possible.*
- $\alpha > \frac{1}{2}$, $\frac{c(V_f)}{h} < \frac{1}{2\alpha - 1}$ *and* $\lambda_a \in (\hat{\lambda}_a, \lambda_f)$*, where* $\hat{\lambda}_a$ *is the maximum solution to the equation*

$$\sqrt{\lambda_a}(h + (\frac{\lambda_f}{\lambda_a})^{\alpha} * c(V_f)) = \sqrt{\lambda_f}(h + c(V_f)) \tag{10}$$

  such that $\hat{\lambda}_a < \lambda_f$*. In this case the optimal intensity for the aggregative tessellation is*

$$\lambda_{a,opt} = (\frac{2\alpha - 1}{x})^{\frac{1}{\alpha}} * \lambda_f$$

Let x *denote the overhead ratio* $x \triangleq \frac{h}{c(V_f)}$*, then the maximum relative cost reduction of acquiring context from aggregative versus the finest grain organization is given by*

$$|1 - \frac{\sqrt{\lambda_{a,opt}}(h + c(V_a))}{\sqrt{\lambda_f}(h + c(V_f))}| = 1 - \frac{2\alpha}{2\alpha - 1}\frac{x}{1 + x}(\frac{2\alpha - 1}{x})^{\frac{1}{2\alpha}}. \tag{11}$$

Note that when $\alpha < \frac{1}{2}$ the context content function of a cell from the 'finest grain' tessellation scales sub-linearly in the area, which is slow enough that aggregation always helps. As shown on the left in Fig. 4 in this case any value for λ_a less than λ_f achieves a cost savings. Note that this is true irrespective of the values of $h, \alpha, c(V_f)$. Of course, the higher the values of $h, c(V_f), \alpha$, the higher the amount of context exchanged, but asymptotically, the amount of cost per unit time for a typical user goes to 0 as λ_a decreases.

When $\alpha > \frac{1}{2}$, context content grows quickly so more care needs to be taken in using aggregative cells, this was already observed in the special case of Theorem III.1. In particular, the interval for the intensity of the aggregative tessellation to be beneficial is now bounded from below, e.g., compare the lower curves shown on the right in Fig. 4. versus the upper curves where aggregation does not pay off. Thus, aggregation is beneficial only for a certain range of values for $c(V_f)$ that depends on the overhead h and α. These conditions guarantee the existence of an optimal intensity for the aggregative tessellation.

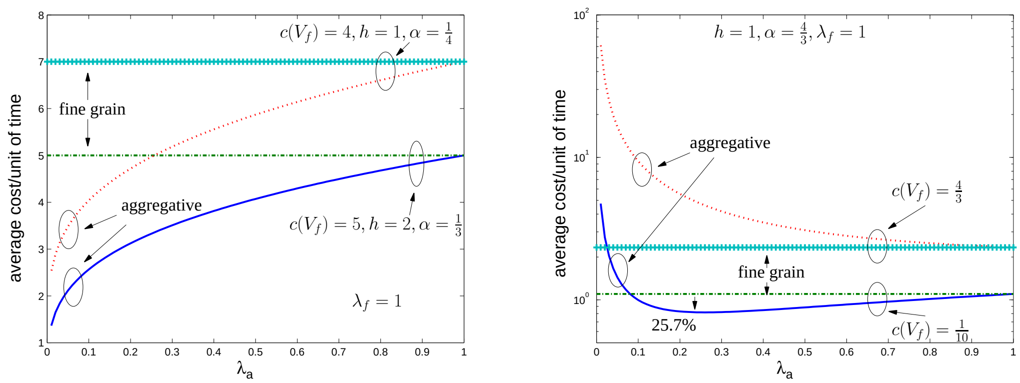

Fig. 4. Context aggregation using aggregative tessellations when $\alpha < \frac{1}{2}$ (left) and $\alpha > \frac{1}{2}$ (right).

Fig. 4 right exhibits a case where $\frac{c(V_f)}{h} > \frac{1}{2\alpha - 1}$, e.g., $c(V_f) = \frac{4}{3}, \alpha = \frac{4}{3}, h = 1$. As can be seen the average cost associated with the aggregative tessellation always exceeds that of the finest grain organization. Thus, aggregation does not pay off. By contrast, when $\frac{c(V_f)}{h} < \frac{1}{2\alpha - 1}$, e.g. $c(V_f) = \frac{1}{10}, \alpha = \frac{4}{3}, h = 1$, there is an interval for λ_a in which aggregation is beneficial. The left and right boundaries of the interval as well as the location of the optimal rate of the aggregate tessellation are those predicted by Theorem IV.1. The relative cost reduction achieved by aggregation for the case shown on the right in Fig. 4 is 25.7%. Recall that for $\alpha = 1$, Fig. 2, we had a higher relative savings of 42.5%. Indeed for $\alpha = \frac{4}{3}$ the benefit of aggregation diminishes because the context content for aggregative cells grows super-linearly, hence the same contextual information could be acquired by exchanging less data by using smaller cells.

An important special case of our context model is when h=0, i.e. the cost is simply proportional to the amount of context exchanged. The following result, proven in the Appendix summarizes what happens in this scenario.

Theorem IV.2. *Under Assumptions II.3, II.4 and II.6 with $h = 0$, the cost associated with an* *aggregative organization associated with an aggregative tessellation of intensity λ_a results in cost reduction only if $\alpha < \frac{1}{2}$ and $\lambda_a \in (0, \lambda_f)$. The optimal intensity for the aggregative tessellation should be as small as possible, since the cost decreases as the intensity λ_a approaches zero.*

One can interpret this result as follows. Recall that the problem with aggregating and exchanging context from coarse grain cells, is that mobile users would obtain contextual data that may not be relevant to them as they move through space. However, when context content scales slowly, i.e., slower than the square root of the area, this corresponds to an application which has quite a bit of spatial redundancy in the contextual information. In this case, coarse aggregative cells effectively compress contextual data, or alternatively the redundancy ensures that most of the context content of an aggregative cell will be relevant. Additionally, the reduced rate of paying for fixed overheads makes the aggregative approach appealing.

In this special case the overhead for acquiring contextual information from a cell is proportional to the context content of the cell, there are no fixed overheads. Additionally, the amount of redundancy in a cell's context is reduced. Therefore, by using increasing amounts of aggregation one can never match the cost of acquiring context from the finest scale tessellation.

A. Intuition: the $\alpha = \frac{1}{2}$ case

It should be obvious by looking at Theorems IV.1 and IV.2 that the case $\alpha = \frac{1}{2}$ has special significance. In this section we will offer an intuitive explanation why this is true. For a 2-dimensional homogeneous Poisson tessellation with rate λ it is known, see [20], that the average chord is proportional to $\frac{1}{\sqrt{\lambda}}$. Assuming simplistically that a mobile moves in a straight line trajectory inside a cell of the aggregative tessellation, the average number of fine-grained cells it crosses is $\sqrt{\frac{\lambda_f}{\lambda_a}}$, see Figure 5. At the same time, the ratio of the context acquired from an aggregative cell over a fine-grained cell is $(\frac{\lambda_f}{\lambda_a})^{\alpha}$. If $\alpha < \frac{1}{2}$, the rate at which context is acquired is smaller than the rate at which new fine-grained cells are crossed, amounting to a true economy. If $\alpha > \frac{1}{2}$ this does not hold anymore and h, the overhead, associated with each transfer plays a dominant role.

B. Selective Context

In Section III we examined the case of a mobile having selective interest in the context of the cells of the finest-grain organization. We will revisit the same case here for generic α. The following fact

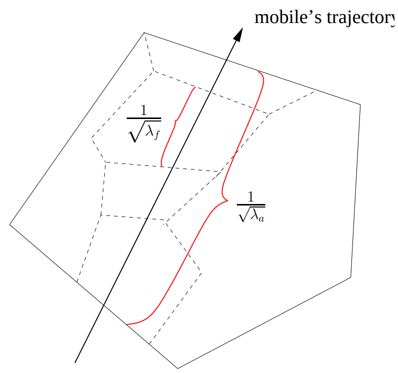

Fig. 5. Intuitive explanation for $\alpha = \frac{1}{2}$.

holds, for a proof see the Appendix.

Fact IV.3. *Under the assumptions of Fact III.2 a necessary condition for an aggregative organization to achieve savings compared to a fine-grained organization is*

- $\alpha < \frac{1}{2}$ *and* $\lambda_a \in (0, \hat{\lambda}_a)$*, where* $\hat{\lambda}_a < \lambda_f$ *is the unique solution to the equation* $r^{2\alpha} - rp(x+q) + x = 0, r \triangleq \sqrt{\frac{\lambda_f}{\lambda_a}}$.
- $\alpha > \frac{1}{2}$ *and* $\frac{p(x+q)}{2\alpha} > \frac{2\alpha}{2\alpha-1}x, x > 2\alpha - 1$*. In this case, there exist* $\lambda'_a < \lambda''_a < \lambda_f$ *s.t.* $\lambda_a \in (\lambda'_a, \lambda''_a)$.

where x *is defined as in Theorem III.1. The optimal scales of aggregation are still given by Theorem IV.1.*

The main conclusion of this fact is that the selective context transfer does not alter the structure of the problem, just the range of the beneficial levels of aggregation. Due to space limitations, we do not discuss the quantitative implications of this fact. We refer the reader to the discussion in Section III which is essentially analogous.

V. Hierarchical Organization for Context Exchange

In this section we focus on applications which have a sub-additive context content function, i.e., $\alpha < 1$ in Assumption II.3. Recall that sub-additivity likely results from spatial redundancy or shared context across fine grain cells. Intuitively, it makes sense to consider a hierarchical organization, whereby shared context is delivered via a coarser level of granularity, while context that is specific to a location is delivered via a fine grained organization. For example, for the case of a mall discussed in Section I, the part of the contextual information that is shared among all stores on the same floor, e.g., locations of emergency exit points, could be exchanged *once* a mobile enters the floor level while information specific to each store, e.g., discounts offered by a store, can be acquired once the mobile enters a store.

In this section a hierarchical organization for context exchanges involves *both* the 'aggregative' and 'finest grain' tessellations introduced earlier, but they are used in a different manner. In particular, when a mobile crosses a cell of the 'finest grain' tessellation she obtains only the context data which is unique to that cell. The *shared* context is exchanged with mobiles when they cross cells of the aggregative tessellation. The idea is to try to minimize overheads while maximizing the relevant context that is exchanged to users.

The effectiveness of our proposed hierarchical organization depends on the average amount of shared context among fine grain cells of the $V(\Pi_f)$ tessellation. We estimate the average shared context as follows. Each cell of the fine-grained tessellation has an average context content $c(V_f)$ and an average *unique* context among their peers denoted by $c(V_f|V_a) < c(V_f)$, since we operate on the $\alpha < 1$ regime and there is spatial redundancy for the context content. The unique part of each fine-grain cell's context, $c(V_f|V_a)$, is stored in the cell. The shared context among all fine-grained cells covered by an aggregate cell, $c(V_f) - c(V_f|V_a)$, is stored in the aggregate cell. A typical cell from the aggregative tessellation $V(\Pi_a)$ has a total,unique plus shared, average context content $c(V_a)$, area $1/\lambda_a$, and will on average cover λ_f/λ_a fine grain cells. Thus, the total context of the aggregative cell should satisfy

$$c(V_a) = \underbrace{\frac{\lambda_f}{\lambda_a} c(V_f|V_a)}_{\text{sum of unique}} + \underbrace{[c(V_f) - c(V_f|V_a)]}_{\text{shared}},$$

where the first term is the sum of the unique context of its constituent fine grain cells, and the second term is the context shared by the fine grain cells. Denoting the context shared by fine grain cells by $s = c(V_f) - c(V_f|V_a)$ one can solve the above equation to obtain:

$$s = \frac{\lambda_f c(V_f) - \lambda_a c(V_a)}{\lambda_f - \lambda_a}. \tag{12}$$

Analogous to the previous sections, the cost per unit of time for a typical user under this hierarchical organization is now given by

$$\sqrt{\lambda_a}(h + s) + \sqrt{\lambda_f}(h + c(V_f) - s),$$

where h is the overhead associated with each context exchange. The first term corresponds to the shared context which is exchanged from aggregative cells while the second term corresponds to the costs associated with exchanging context which is unique to the 'finest grain' cells. Under this model we can show the following results, where again we have relegated the derivations to the Appendix.

Theorem V.1. *Under Assumptions II.3, II.4 and II.6, the hierarchical organization for context exchanges achieves a cost saving over the fine-grained organization if:*

$$\alpha < 1 \text{ , } \lambda_a \in (0, \lambda_f) \textit{ and } x \triangleq \frac{h}{c(V_f)} < \frac{r^2 - r^{2\alpha}}{r+1} \triangleq f(r)$$

where $r \triangleq \sqrt{\frac{\lambda_f}{\lambda_a}}$.

The interpretation of Theorem V.1 is that when context content scales sub-linearly, for increasing levels of aggregation the hierarchical organization will always produce savings compared to acquiring context from a fine-grained tessellation.

Theorem V.2. *Under Assumptions II.3, II.4 and II.6, the hierarchical organization for context exchanges achieves a cost saving over the aggregative organization if:*

$$\alpha < 1 \text{ , } \lambda_a \in (0, \lambda_f) \textit{ and } x \triangleq \frac{h}{c(V_f)} < \frac{r^{2\alpha} - 1}{r+1} \triangleq g(r)$$

where $r \triangleq \sqrt{\frac{\lambda_f}{\lambda_a}}$.

Note that if $\alpha > \frac{1}{2}$ exchanging context from a hierarchical organization can lead to a cost savings. Indeed if $\alpha > \frac{1}{2}$, then $\lim_{r \to \infty} g(r) = \infty$ so the condition in Theorem V.2 is eventually satisfied irrespective of the value of $c(V_f), h$. So it suffices to employ a sufficiently coarse granularity (λ_a small enough) for the hierarchical approach to result in cost savings. Observe in Fig. 6 and in Fig. 7 that the hierarchical approach can reduce the cost for any value of $c(V_f)$ while the aggregative approach has at best a certain range over which it can achieve cost reduction.

By contrast, if $\alpha < \frac{1}{2}$ exchanging context from a hierarchical organization may or may not be preferable to an organization based on aggregation, depending on the coarseness of aggregate cells one can practically achieve. From Theorem IV.1 we know that for $\alpha < \frac{1}{2}$ an aggregative approach always results in cost savings. In this case the limit of the upper bound $g(r)$ of Theorem V.2 goes to 0 as $r \to \infty$. Fig. 8. Thus in this case, a designer should be careful enough to evaluate both approaches before deciding which one is better. In Fig. 8 we observe that the aggregative approach results in

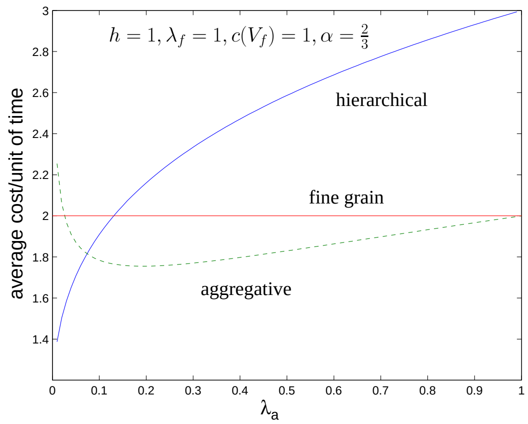

Fig. 6. Hierarchical vs. aggregative when $\alpha > \frac{1}{2}, c(V_f) = 1$**.**

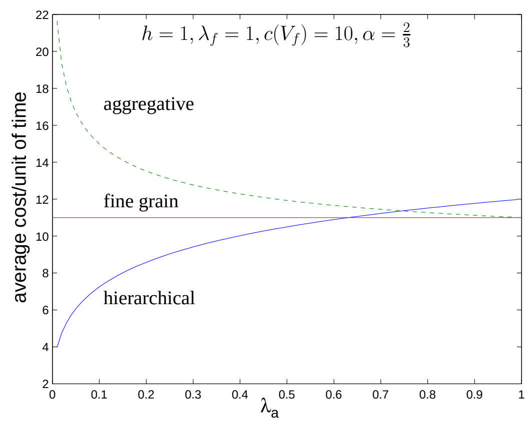

Fig. 7. Hierarchical vs. aggregative when $\alpha > \frac{1}{2}, c(V_f) = 10$**.**

cost savings for all allowable values of λ_a, while the hierarchical approach needs cells to be coarse enough to do so. Once cells are coarse enough, the hierarchical approach produces savings that for the specific values chosen for the graph in Fig. 8 outperform the aggregative approach for all practically achievable scales of aggregation.

A comparison summary of the two organizations is shown in Table V.

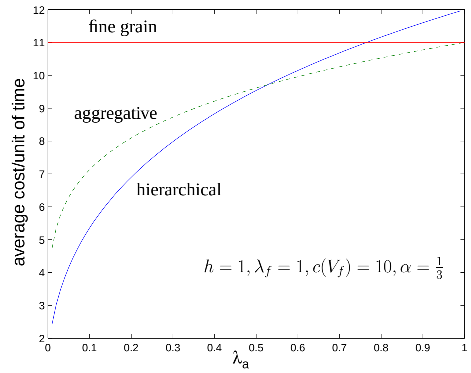

Fig. 8. Hierarchical vs. aggregative $\alpha < \frac{1}{2}$.

$\alpha > 1$,'super-linear'	Aggregative organizations *may* help depending on the scale of aggregation
	Hierarchical organizations *cannot* help
$1 > \alpha > \frac{1}{2}$,'spatially correlated'	Aggregative organizations *may* help depending on the scale of aggregation
	Hierarchical organizations *eventually* help
	Hierarchical organizations *eventually* outperform aggregative ones
$\alpha < \frac{1}{2}$,'highly spatially correlated'	Aggregative organizations *always* help
	Hierarchical organizations *eventually* help
	Aggregative organizations *eventually* outperform hierarchical ones

TABLE I. Aggregative vs. Hierarchical Organizations.

VI. Assessing the Cost of Surveillance in Ubiquitous Environments

Throughout this paper we have implicitly assumed the existence of a surveillance mechanism that is part of a space's infrastructure and allows mobiles to detect when they cross cell boundaries. We envisage two generic types of surveillance mechanisms.

- A *direct* mechanism that is part of a space's infrastructure monitoring each cell's boundary. An airport or shopping mall with RFID readers installed on the doors exciting the RFID tags on the mobiles passing through, would be an example of such a mechanism. Such a mechanism is assumed in [21].
- An *indirect* mechanism that detects boundary crossings by comparing each mobile's location to the location of the cell boundaries. A tracking service that is part of the infrastructure or self-positioning by each mobile device can be used to calculate location. We assume that self-

positioning mobile nodes detect boundary crossings using an a-priori downloaded map of the cell boundaries. For a well known location system based on this approach, see e.g., [22].

Let us first consider direct surveillance mechanisms. We abstract the underlying mechanism by assuming that the cost to detect a boundary crossing is E_d units of energy/device, e.g., the energy in an RFID reader's pulse to read a potential tag. Additionally we let f_d denote the frequency with which the mechanism checks for boundary crossings, e.g. an RFID reader on a door sends a pulse every second to detect mobiles, we say that $f_d = 1\ Hz$. Clearly, there will be a trade-off between the surveillance frequency and the timeliness with which contextual data is delivered. Finally, it is known that the average cell boundary per unit of area for a homogeneous Poisson Voronoi tessellation with rate λ is $2\sqrt{\lambda}$ [18]. Motivated by the use-cases presented in [21] and practical considerations we note that mobiles moving from cell to cell pass through designated points e.g., doors and detectors will have a certain coverage range so only a fraction, K_d, of the total cell boundary has to be surveilled directly. The following assumption captures these elements.

Assumption VI.1. *We assume that the average power for a direct surveillance mechanism per unit of area is given by*

$$2 * \sqrt{\lambda} * K_d * f_d * E_d. \tag{13}$$

With this additional assumption the power expended for surveilling cell boundaries and exchanging context using an aggregative tessellation with intensity λ_a is given by

$$2 * \sqrt{\lambda_a} * K_d * f_d * E_d + \lambda_0 * \frac{4 * v}{\pi} \sqrt{\lambda_a}(c(V_a) + h). \tag{14}$$

This in turn can be simplified as $\sqrt{\lambda_a} * (\hat{h} + \hat{K} c(V_a))$ where $\hat{h}$ and $\hat{K}$ are appropriate constants. This cost function has the same form as that considered in Section II, This cost function has the same form as that considered in Section II, which leads to the following corollary.

Corollary VI.2. *Theorem IV.1 can be applied for optimizing the intensity of an aggregative tessellation for context exchange using direct surveillance. The overhead ratio is given by* $x \triangleq \frac{\hat{h}}{\hat{K} c(V_f)}$.

Note that for the case $\alpha > \frac{1}{2}$ the frequency at which the shared infrastructure surveils boundary crossings, f_d, plays a key role. An increased value of f_d is beneficial in two ways: it increases the timeliness of mobile detection and increases the parameter overhead ratio x above the $2\alpha - 1$ threshold required for the aggregative tessellation to produce savings, see Theorem IV.1. Also note that the added cost of surveillance does not change the underlying structure of the problem i.e. the existence of optimal values for the scale of aggregation still depends on the scaling (α) of the context content function. Thus, the results derived in the previous sections provide a designer with the tools to roughly evaluate how to optimize the aggregative organizations.

For the indirect surveillance mechanism we define the frequency with which a mobile acquires location information f_i and the corresponding energy expended E_i in a similar way as in the direct case.

Assumption VI.3. *Under Assumption II.4 the average power per unit area expended by an indirect surveillance mechanism to track boundary crossings is*

$$\lambda_0 * f_i * E_i. \tag{15}$$

Observe that the power expended increases linearly with the intensity of the mobiles. Such an approach would face scalability problems if the number of mobiles increases significantly as expected in ubiquitous computing scenarios.

Note that the frequencies f_d, f_i, must be high enough to ensure that a mobile does not 'miss' acquiring context from a cell in a timely manner. Intuitively, the higher the intensity of the aggregative process, λ_a, the higher the cost for exchanging context. The following fact, formally proved in the Appendix, provides lower bounds on f_d, f_i.

Fact VI.4. *Under Assumption II.4 and surveillance frequencies f_d, f_i chosen to ensure that a mobile transitioning through an average sized cell of an aggregative organization of rate λ_a is not missed, the frequency f_d must satisfy $f_d > K_d^f$ and the frequency f_i should satisfy $f_i > K_i^f * \sqrt{\lambda_a}$.*

K_d and K_i are constants depending on the average velocity of the mobiles,the range of the devices used to perform the surveillance and the average velocity of the mobiles respectively. Note that the previous fact is predicated on not missing on *average* sized cell, so in practice the variability in cell sizes would require surveillance frequencies to be higher.

For a given aggregative organization, i.e., fixed λ_a a designer can consider which surveillance mechanism is more energy efficient.

Fact VI.5. *The direct surveillance is more efficient than the indirect surveillance if*

$$2 * \sqrt{\lambda_a} * K_d * f_d * E_d < \lambda_0 * f_i * E_i \tag{16}$$

For services offered on a 'personalized' scale i.e. $\lambda_f \sim \Theta(\lambda_0)$, $\sqrt{\lambda_a} << \lambda_0$ *for an aggregative tessellation and the leverage of the shared infrastructure by a direct surveillance mechanism provides significant gains.*

VII. Conclusions and Future Work

This paper is a first attempt at studying the fundamental characteristics of context exchange and surveillance organizations for ubiquitous applications. To allow for quantitative arguments we propose a simple stochastic geometric model that naturally represents the main characteristics of such systems. The key results show how the effectiveness of optimal aggregative versus hierarchical organizations depend on the manner in which context content scales with area. We also consider how energy costs for direct and indirect surveillance mechanisms would vary under such organizations. Clearly, our model has several simplifications that it would be of interest to relax, and are part of our future work. Among other issues, it would be of interest to capture how limited caching of contextual data might enable mobiles to reduce their energy expenditures while maintaining updated contextual information.

VIII. Acknowledgments

This work was partially supported by NSF Award CNS-0509355 and an AFOSR Award FA9550-07-1-0428.

References

[1] D. M. et. al., "Electronic textiles: A platform for pervasive computing," Proceedings of the IEEE, No. 12., Vol. 91, 2003.

[2] R. Want, "An introduction to RFID technology," *Pervasive Computing*, January 2006.

[3] F. Baccelli and S. Zuyev, "Poisson-Voronoi spanning trees with applications to the optimization of communication networks," *Oper. Res.*, vol. 47, no. 4, pp. 619–631, 1999.

[4] F. Baccelli, M. Klein, M. Lebourges, and S. Zuyev, "Stochastic geometry and architecture of communication networks," *Communications journal*, 1996, select Proceedings of the third INFORMS telecommunications conference.

[5] A. Schmidt, "Ubiquitous computing - computing in context," Ph.D. dissertation, University of Lancaster, 2002.

[6] A. K. Dey, "Providing architectural support for building context-aware applications," Ph.D. dissertation, Georgia Institute of Technology, Atlanta, GA, USA, 2000, director-Gregory D. Abowd.

[7] C. P. John Loughney, Madjid Nakhjiri and R. Koodli, "Context transfer protocol, RFC 4067," July 2005. [Online]. Available: ftp://ftp.rfc-editor.org/in-notes/rfc4067.txt

[8] B. Brumitt, J. Krumm, B. Meyers, and S. Shafer, "Ubiquitous computing and the role of geometry," in *IEEE Personal Communications*, 2000, pp. 41–43.

[9] G. Chen and D. Kotz, "Context aggregation and dissemination in ubiquitous computing systems," in *WMCSA '02: Proceedings of the Fourth IEEE Workshop on Mobile Computing Systems and Applications*, 2002, p. 105.

[10] N. H. Cohen, A. Purakayastha, J. Turek, L. Wong, and D. Yeh, "Challenges in flexible aggregation of pervasive data," IBM Research Division, Thomas J. Watson Research Center., Tech. Rep., 2001.

[11] H. Ryu, I. Park, S. J. Hyun, and D. Lee, "A task decomposition scheme for context aggregation in personal smart space," in *SEUS*, 2007, pp. 20–29.

[12] T. Zimmer, "Qoc: Quality of context - improving the performance of context-aware applications," in *Advances in Pervasive Computing*, T. P. et al., Ed., vol. 207, May 2006.

[13] M. Beigl, A. Krohn, T. Zimmer, C. Decker, and P. Robinson, "Awarecon: Situation aware context communication," in *Proceedings of Ubicomp*, Seattle, October 2003, pp. 12–15.

[14] T. Yamabe, A. Takagi, and T. Nakajima, "Citron: A context information acquisition framework for personal devices," in *RTCSA '05: Proceedings of the 11th IEEE International Conference on Embedded and Real-Time Computing Systems and Applications*. Washington, DC, USA: IEEE Computer Society, 2005, pp. 489–495.

[15] E. J. Y. Wei and A. T. S. Chan, "Towards context-awareness in ubiquitous computing," in *EUC*, 2007, pp. 706–717.

[16] S. J. H. Yang, A. F. M. Huang, R. Chen, S.-S. Tseng, and Y.-S. Shen, "Context model and context acquisition for ubiquitous content access in ulearning environments," in *SUTC '06: Proceedings of the IEEE International Conference on Sensor Networks, Ubiquitous, and Trustworthy Computing - Vol 2 - Workshops*. Washington , DC , USA: IEEE Computer Society, 2006, pp. 78–83.

[17] M. Khedr and A. Karmouch, "Negotiating context information in context-aware systems," *IEEE Intelligent Systems*, vol. 19, no. 6, pp. 21–29, 2004.

[18] D. Stoyan, W. S. Kendall, and J. Mecke, *Stochastic Geometry and its Applications*. Chichester: John Wiley and Sons, 1987.

[19] I. Rhee, M. Shin, S. Hong, K. Lee, and S. Chonga, "On the levy-walk nature of human mobility," in *Infocom*, 2008.

[20] L. Muche and D. Stoyan, "Contact and chord length distributions of the poisson voronoi tessellation," *Journal of Applied Probability*, vol. 29, no. 2, pp. 467–471, 1992.

[21] A. Ziotopoulos, M. Jacome, and G. de Veciana, "An RFID-based platform supporting context-aware computing in complex spaces," in *Second International Workshop on Managing Context Information and Semantics in Mobile Environments (MCISME), in conjunction with the 8th International Conference on Mobile Data Management (MDM'07)*, May 2007.

[22] P. Bahl and V. N. Padmanabhan, "RADAR: An in-building RF-based user location and tracking system," in *INFOCOM (2)*, 2000, pp. 775–784.

Appendix

Proof of Theorem III.1: Starting from Eq. 4 we get

$$\sqrt{\lambda_a} * (h + \frac{\lambda_f}{\lambda_a} c(V_f)) < \sqrt{\lambda_f} * (h + c(V_f)) \Rightarrow h + \frac{\lambda_f}{\lambda_a} c(V_f) < \sqrt{\frac{\lambda_f}{\lambda_a}}(h + c(V_f)).$$

By substituting $r \triangleq \sqrt{\frac{\lambda_f}{\lambda_a}}$ and $x \triangleq \frac{h}{c(V_f)}$ we get

$$x + r^2 < r(x+1) \Leftrightarrow r^2 - r(x+1) + x < 0.$$

Observe that this inequality has a solution since $\Delta = (x+1)^2 - 4x = (x-1)^2 \geq 0$. The inequality can be equivalently rewritten

$$(r - \frac{(1+x)+(x-1)}{2})(r - \frac{(1+x)-(x-1)}{2}) < 0 \Rightarrow (r-x)(r-1) < 0.$$

The solution to the previous inequality is either

$$(r < x) \wedge (r > 1) \Leftrightarrow 1 < r < x \Leftrightarrow 1 < \sqrt{\frac{\lambda_f}{\lambda_a}} < \frac{h}{c(V_f)} \Leftrightarrow \lambda_a \in (\lambda_f (\frac{c(V_f)}{h})^2, \lambda_f),$$

or

$$(r > x) \wedge (r < 1) \Leftrightarrow x < r < 1 \Leftrightarrow \frac{h}{c(V_f)} < \sqrt{\frac{\lambda_f}{\lambda_a}} < 1 \Leftrightarrow \lambda_a \in (\lambda_f, \lambda_f (\frac{c(V_f)}{h})^2).$$

The second solution is rejected as it does not correspond to aggregation, i.e., $\lambda_a > \lambda_f$.

To calculate the optimal rate for the aggregative tessellation we will minimize the amount of context exchanged in that case. Taking the derivative with respect to λ_a for the left hand-side of Eq. 4 and setting it equal to zero we get

$$\frac{h}{2\sqrt{\lambda_a}} - \frac{\lambda_f c(V_f)}{2\lambda_a^{\frac{3}{2}}} = 0 \Leftrightarrow \lambda_{a,opt} = \frac{c(V_f)}{h}\lambda_f.$$

This solution is guaranteed to lie in the interval $(\lambda_f (\frac{c(V_f)}{h})^2, \lambda_f)$ since $(\frac{c(V_f)}{h})^2 < \frac{c(V_f)}{h} < 1$ for $\frac{c(V_f)}{h} < 1$, therefore, it is a true minimum.

The maximum relative cost reduction compared to acquiring context from the finest grain organization $V(\Pi_f)$ is

$$1 - \frac{\sqrt{\lambda_{a,opt}}(h + c(V_a))}{\sqrt{\lambda_f}(h + c(V_f))}.$$

By plugging in the value for $\lambda_{a,opt}$ we get $1 - \frac{\sqrt{hc(V_f)}}{\frac{h+c(V_f)}{2}}$, which can be simplified to $1 - \frac{2\sqrt{x}}{1+x}$.

Proof of Theorem IV.1: Starting from Eq. 9 we get

$$\sqrt{\lambda_a} * (h + (\frac{\lambda_f}{\lambda_a})^\alpha * c(V_f)) < \sqrt{\lambda_f}(h + c(V_f)).$$

Substituting $x \triangleq \frac{h}{c(V_f)}$ and $r \triangleq \sqrt{\frac{\lambda_f}{\lambda_a}}$ we get

$$x + r^{2\alpha} < r(x+1) \Leftrightarrow r^{2\alpha} - r < x(r-1). \tag{17}$$

We care about solutions that correspond to aggregation, i.e., $r > 1$. First assume that $\alpha < \frac{1}{2}$. In this case $x(r-1) > 0$ and $r^{2\alpha} - r < 0$, so by default we can see that the inequality is satisfied for every $r > 1$. That means that the allowable range for λ_a is $(0, \lambda_f)$.

Now, we will examine the case $\alpha > \frac{1}{2}$. Observe that in this regime,asymptotically, the term $r^{2\alpha}$ will grow bigger than the term $r(x+1)$. This means that the inequality has at best an interval over which it can be satisfied. We will identify the conditions under which this interval does exist. Manipulating Eq. 17 we get

$$f(r) \triangleq r^{2\alpha} - r(x+1) + x < 0, r > 1 \tag{18}$$

Observe that $f(1) = 1-(x+1)+x = 0$. To ensure the interval in question exists we impose a decreasing condition on $f()$ at $r = 1$, $f'(1) = 2\alpha r^{2\alpha-1} - (x+1) = 2\alpha - (x+1) < 0 \Rightarrow \frac{x+1}{2\alpha} > 1, \alpha > \frac{1}{2}$. The minimum for $f()$ is attained at r_0 s.t. $r_0^{2a-1} = \frac{x+1}{2\alpha}$. The condition for decreasing behavior imposed on $f()$ guarantees that $r_0 > 1$ i.e. we are in the aggregation regime. The value at r_0 is indeed a minimum for $f()$ since $f''(r_0) = 2\alpha(2\alpha-1)r_0^{2\alpha-2} > 0$. Additionally, we also need to prove that the minimum value for $f()$ is negative, thus satisfies Eq. 18. As noted previously $f(1) = 0$, therefore the minimum $f(r_0) < 0$.

All the previous properties proven for $f()$ taken collectively ensure that the interval $(\hat{\lambda_a}, \lambda_f)$ where $\hat{\lambda_a}$ is the maximum solution to the equation 10 is the interval where aggregation is a win. The existence of $\hat{\lambda_a}$ is guaranteed by the fact that f is continuous, has a negative minimum for $r_0 > 1$ and asymptotically goes to infinity for $r \to \infty$.

To find the value where the amount of context transfered by an aggregative organization is minimized we define the following auxiliary function.

$$g(\lambda_a) \triangleq \sqrt{\lambda_a}(h + \frac{\lambda_f^\alpha}{\lambda_a^\alpha}c(V_f)) - \sqrt{\lambda_f}(h + c(V_f)).$$

We take the derivative of $g()$ and set it equal to 0.

$$\frac{h}{2\sqrt{\lambda_a}} + (\frac{1}{2} - \alpha)\lambda_f^\alpha c(V_f)\frac{1}{\sqrt{\lambda_a}\lambda_a^\alpha} = 0,$$

which gives us the value for $\lambda_{a,opt}$

$$\lambda_{a,opt} = (\frac{(2\alpha-1)c(V_f)}{h})^{\frac{1}{\alpha}}\lambda_f.$$

This expression is valid since $\alpha > \frac{1}{2}$. The optimality is verified by checking the sign of the second derivative of $f()$ at $\lambda_{a,opt}$.

$$f''(\lambda_\alpha) = -\frac{h}{4\lambda_\alpha^{\frac{3}{2}}} - \frac{1-2\alpha}{2}c(V_f)\lambda_f^\alpha(\alpha + \frac{1}{2}).\frac{1}{\lambda_a^{\alpha+\frac{3}{2}}}.$$

Manipulating the previous expression we get

$$f''(\lambda_a) = \frac{1}{4\lambda_\alpha^{\frac{3}{2}}}(\frac{(2\alpha+1)(2\alpha-1)c(V_f)\lambda_f^\alpha}{\lambda_a^\alpha} - h).$$

Substituting the expression for $\lambda_{a,opt}$ we get

$$f''(\lambda_{a,opt}) = \frac{\alpha h}{2\lambda_{a,opt}^{\frac{3}{2}}} > 0.$$

therefore, we have achieved a minimum.

The optimal rate for the aggregative tessellation corresponds to aggregation if

$$\frac{(2\alpha-1)c(V_f)}{h} < 1 \Leftrightarrow \frac{c(V_f)}{h} < \frac{1}{2\alpha-1}.$$

which is the condition we have already identified above.

Due to space limitations we will not provide a proof for the expression of the maximum relative cost reduction. The steps are essentially the same as with the case $\alpha = 1$.

Proof of Theorem IV.2: For the case that the cost is proportional to the amount of context exchanged, a exchange of d amount of context through a, e.g., wireless network with maximum packet size equal to p_{max}, will require on average $\frac{d}{p_{max}-o}$ packets to be sent. If each packet incurs an overhead of o then the total amount of data exchanged is $\frac{d}{p_{max}-o} * o + d = (\frac{o}{p_{max}-o} + 1) * d$. For the case of tessellations $V(\Pi_a)$ and $V(\Pi_f)$ substituting $\frac{\lambda_f^\alpha c(V_f)}{\lambda_a^\alpha}$ and $c(V_f)$ to the previous equation we get:

$$\sqrt{\lambda_a}(\frac{o}{p_{max}-o}+1)\frac{\lambda_f^\alpha c(V_f)}{\lambda_a^\alpha} < \sqrt{\lambda_f}(\frac{o}{p_{max}-o}+1)c(V_f) \Rightarrow \lambda_a^{\frac{1}{2}-\alpha} < \lambda_f^{\frac{1}{2}-\alpha}.$$

If $\frac{1}{2} - \alpha > 0 \Leftrightarrow \alpha \leq \frac{1}{2}$ aggregation is always a win. If $\alpha > \frac{1}{2}$ aggregation never pays off.

Proof of Theorem V.1: The condition for a hierarchical organization to obtain savings compared to acquiring context from the fine-grained tessellation is

$$\sqrt{\lambda_a}(h+s) + \sqrt{\lambda_f}(h + c(V_f) - s) < \sqrt{\lambda_f}(h + c(V_f)) \quad (19)$$

where $r \triangleq \sqrt{\frac{\lambda_f}{\lambda_a}} > 1, s \triangleq \frac{r^2 - r^{2\alpha}}{r^2-1}c(V_f)$ is the amount of shared context content between the cells of a fine-grained tessellation that are covered by a cell of the aggregative tessellation. Manipulating the previous equation we get

$$(h+s) = r(h + c(V_f) - s) < r(h + c(V_f)) \Leftrightarrow$$

$$(x + \frac{s}{c(V_f)}) + r(x + 1 - \frac{s}{c(V_f)}) < r(x+1), x \triangleq \frac{h}{c(V_f)}$$

Plugging-in the expression for s we get

$$x < r\frac{s}{c(V_f)} - \frac{s}{c(V_f)} = \frac{s}{c(V_f)}(r-1) \Leftrightarrow x < \frac{r^2 - r^{2\alpha}}{r^2-1}(r-1) = \frac{r^2 - r^{2\alpha}}{r+1}$$

If $a \geq 1 \Rightarrow r^2 < r^{2\alpha}$, for $r \in [1, \infty)$. So the previous inequality cannot be satisfied. If $a < 1 \Rightarrow r^2 > r^{2\alpha}$, for $r \in [1, \infty)$. So there exists a level of aggregation i.e. r_0 s.t. for $r > r_0$ the hierarchical organization wins.

Proof of Theorem V.2: The region of allowable values for λ_a and α can be defined by the sub-additive growth constraint. We want the following relation to hold

$$\frac{c(V_f)\lambda_f^\alpha}{{\lambda_a}^\alpha} < \frac{c(V_f)\lambda_f}{\lambda_a} \Leftrightarrow (\frac{\lambda_f}{\lambda_a})^\alpha < \frac{\lambda_f}{\lambda_a}. \tag{20}$$

Additionally, we have to ensure that the amount of context found in a typical cell of a $V(\Pi_f)$ tessellation, the lowest level of the hierarchy, is less than the amount of context in a cell of the, higher, $V(\Pi_a)$ tessellation. This is an extra constraint relating the rates of the $V(\Pi_f)$ tessellation with the rate of the $V(\Pi_a)$ tessellation. The following equation expresses the previous equation:

$$\frac{c(V_f)\lambda_f^\alpha}{{\lambda_a}^\alpha} > c(V_f) \Leftrightarrow (\frac{\lambda_f}{\lambda_a})^\alpha > 1. \tag{21}$$

Combining Eqs. 20 and 21 we get the following allowable regions for λ_a with respect to λ_f and α:

$$1 < (\frac{\lambda_f}{\lambda_a})^\alpha < \frac{\lambda_f}{\lambda_a} \Leftrightarrow 0 < \alpha < 1 \text{ and } \lambda_a \in (0, \lambda_f).$$

In order for the hierarchical approach to have reduced cost compared to the aggregative approach we would like the following inequality to hold:

$$\sqrt{\lambda_a}(h+s) + \sqrt{\lambda_f}(h + c(V_f) - s) < \sqrt{\lambda_a}(h + c(V_a)).$$

where $s \triangleq \frac{r^2 - r^{2\alpha}}{r^2-1}c(V_f)$ is the amount of shared context content between the cells of a fine-grained tessellation that are covered by a cell of the aggregative tessellation.

Substituting the expression for $c(V_a)$ from Section V and setting $r \triangleq \sqrt{\frac{\lambda_f}{\lambda_a}} > 1$ we get

$$(h+s) + r(h + c(V_f) - s) < h + r^2(c(V_f) - s) + s) \Leftrightarrow$$

$$x + \frac{s}{c(V_f)} + r(x + 1 - \frac{s}{c(V_f)}) < x + r^2(1 - \frac{s}{c(V_f)}) + \frac{s}{c(V_f)}, x \triangleq \frac{h}{c(V_f)}$$

Substituting the expression for s we get

$$x + \frac{r^2 - r^{2\alpha}}{r^2-1} + r(x + 1 - \frac{r^2 - r^{2\alpha}}{r^2-1}) < x + r^2(1 - \frac{r^2 - r^{2\alpha}}{r^2-1}) + \frac{r^2 - r^{2\alpha}}{r^2-1}$$

Cancelling common terms we get the expression of Theorem V.2.

$$rx < r^2(1 - \frac{r^2 - r^{2\alpha}}{r^2 - 1}) - r(1 - \frac{r^2 - r^{2\alpha}}{r^2 - 1}) \Leftrightarrow$$
$$rx < r(r-1)\frac{r^2 - 1 - r^2 + r^{2\alpha}}{(r-1)(r+1)} \Leftrightarrow$$
$$x < \frac{r^{2\alpha} - 1}{r+1}$$

Proof of Fact VI.4: For the direct case, note that the detection of a boundary crossing event must happen in an area of radius R around the device performing the detection, provided that R is smaller than the mean cell size. This assumption currently holds in practice e.g. passive state-of-the-art RFID readers have a range $R \sim 50\ cm$ while the mean diameter of a e.g. store in a mall is at least $5\ m$. Thus, the frequency f_d has to be greater than $\frac{\bar{v}}{R}$ where $\bar{v}$ is the average speed of mobiles. For the indirect case, since cells are almost surely bounded, a mobile entering a typical cell will eventually leave the cell. Consider a typical cell, by Little's theorem, the average sojourn time in the cell T_{soj} must be the ratio of the mean number of mobiles in the cell to the average rate of users into the cell. Assuming that the initial distribution of mobiles is Poisson with intensity λ_0 and under our the mobility model in Assumption II.4 one can show that the mean number of mobiles in a typical cell is λ_0/λ_a while the mean rate of mobiles entering the cell is proportional to their intensity, mean velocity and typical cells perimeter, i.e., $\frac{\lambda_0 v}{\sqrt{\lambda_a}}$. Thus, the mean sojourn time is proportional to $\frac{1}{v\sqrt{\lambda_a}}$ giving the desired result.

Proof of Fact III.2: The necessary condition to achieve savings is

$$\sqrt{\lambda_a}(h + \frac{\lambda_f}{\lambda_a}c(V_f)) < \sqrt{\lambda_f}(h + qc(V_f))p$$

Performing the substitutions $r \triangleq \sqrt{\frac{\lambda_f}{\lambda_a}}, \quad x \triangleq \frac{h}{c(V_f)}$, the previous inequality now becomes

$$r^2 - rp(x+q) + x < 0$$

We define the function $f(r) \triangleq r^2 - rp(x+q) + x$ and study its behavior in the interval $r > 0$. In order to achieve aggregation we need the minimum of f to be achieved in the $r > 1$ interval. So,

$$f'(r) = 0 \Leftrightarrow r_{opt} = \frac{p(x+q)}{2} > 1 \Leftrightarrow p(x+q) > 2$$

The minimum value achieved should be negative

$$f(r_{opt}) = \frac{p^2(x+q)^2}{4} - \frac{p^2(x+q)^2}{2} + x = x - \frac{p^2(x+q)^2}{4} < 0 \Leftrightarrow p(x+q) > 2\sqrt{x}$$

Combining the previous two equations we get

$$p(x+q) > 2\max(1, \sqrt{x})$$

If $x < 1$, the maximum value of the left hand-side of the previous equation is $x+1$ and corresponds to the case $p = q = 1$. Also $\max(1, \sqrt{x}) = 1$ and as a result $x + 1 > 2 \Leftrightarrow x > 1$ which contradicts our assumption. Therefore, we have $x > 1$ and the necessary condition is

$$p(x+q) > 2\sqrt{x}$$

Proof of Fact IV.3: We will follow the steps of the previous proof but now the definition of $f()$ is $f(r) = r^{2\alpha} - rp(x+q) + x$. For $\alpha < \frac{1}{2}$, observe that $f(1) = 1 - p(x+q) + x = (1-pq) + (1-p)x > 0$, but r grows faster than $r^{2\alpha}$ as $r \to \infty$. Therefore, there exists a point s.t. for $r > \hat{r_a}$ the aggregative organization is a win, see Fig.9 (left).

For $\alpha > \frac{1}{2}$, $f(1) > 0$ and $r^{2\alpha}$ grows faster than r as $r \to \infty$. The aggregative organization can be a win only if there exists an interval where $f()$ is decreasing and its minimum is negative. Therefore, $f'(r) = 2\alpha r^{2\alpha-1} - p(x+q) \Rightarrow f'(1) = 2\alpha - p(x+q) < 0 \Rightarrow 2\alpha < p(x+q)$. The minimum is achieved at r_0 s.t. $f'(r_0) = 0 \Leftrightarrow 2\alpha r_0^{2\alpha-1} - p(x+q) = 0 \Leftrightarrow r_0^{2\alpha-1} = \frac{p(x+q)}{2\alpha} > 1$. The previous result combined with the observation that we are operating in the $\alpha > \frac{1}{2}$ regime leads to the conclusion that $r_0 > 1$.

We will identify the necessary conditions so that the minimum $f(r_0)$ is negative. $f(r_0) = r_0^{2\alpha} - r_0 p(x+q) + x$, using the value of r_0 the previous expression becomes $r_0 \frac{p(x+q)}{2\alpha} - r_0 p(x+q) + x$. Since $r_0 > 1$ as proved earlier an upper bound for the minimum is $r_0 \frac{p(x+q)}{2\alpha} - r_0 p(x+q) + r_0 x = r_0(\frac{p(x+q)}{2\alpha} - p(x+q) + x)$. A necessary condition for the miminum to be negative is $\frac{p(x+q)}{2\alpha} - p(x+q) + x < 0 \Leftrightarrow p(x+q) > 2\alpha \frac{x}{2\alpha-1}$. Combining this result with the result for the decreasing behavior of $f()$ we get $p(x+q) > 2\alpha \max\{1, \frac{x}{2\alpha-1}\}$. Assume $x < 2\alpha - 1$, then $p(x+q) > 2\alpha$. But, at the same time $p(x+q) < x+1 < 2\alpha - 1 + 1 = 2\alpha$, which is a contradiction. Therefore, $p(x+q) > \frac{2\alpha}{2\alpha-1} x, x > 2\alpha - 1$.

Since $f(1) > 0, f(r) \to \infty$ as $r \to \infty$ and $f()$ has a negative minimum, there exist r'_a, r''_a s.t. for $r \in (r'_a, r''_a)$ an aggregative organization is a win, see Fig. 9 (right).

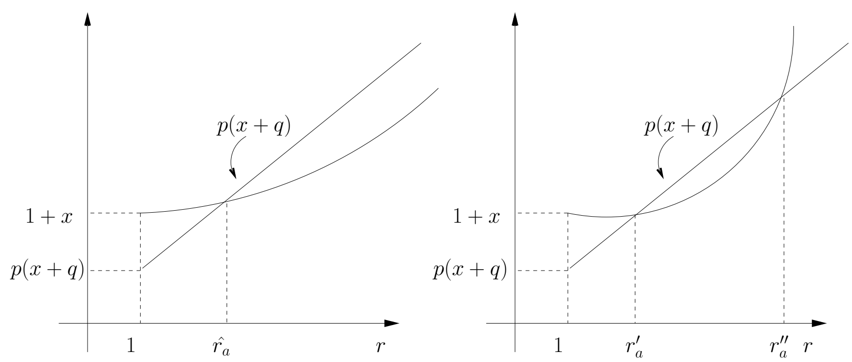

Fig. 9. Behavior of selective context transfer for aggregative organization for $\alpha < \frac{1}{2}$ (left) and $\alpha > \frac{1}{2}$ (right).