

Design and Optimization of Spatial Organizations For Context Exchange and Surveillance

Ayis Ziotopoulos, Gustavo de Veciana

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

Abstract

Modern environments 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. 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.

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. 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.

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 e.g. a point of sale in a retail store, 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., [1], [2]. 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 notifying 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, see e.g., [3]. 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 an established idea, see e.g., [4]. Context aggregation is widely recognized as an efficient policy for coping with the scalability issues challenging ubiquitous systems, see e.g., [5] and [6].

Contributions and organization. This paper does not target the CoMoRea09 challenge. Instead, our work focuses on issues relevant to context scaling modeling and context management. 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 a random Voronoi tessellation, for a spatial organization of context exchanges to/from mobile users/terminals. In order to argue quantitatively about the relative merits of different architectures, we also propose a taxonomy for how context content scales with the area of a cell for various applications. In Section III 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. Our main result is that 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 IV we consider hierarchical organizations for context exchange. 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 V we elaborate

on different mechanisms to surveil a space’s cells and consider their contribution to the energy costs. The paper concludes with Section VI.

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, [7], 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., [1], [2]. 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. When 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 [7]. 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. The definition of a homogeneous Poisson process Π with intensity λ and the corresponding tessellation $V(\Pi)$ can be found in [7].

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.

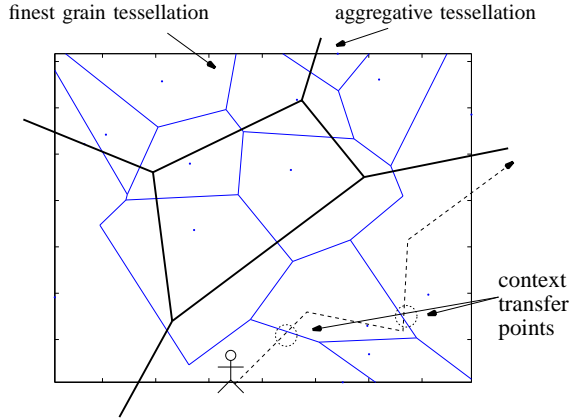


Fig. 1. ‘Finest grain’ and ‘aggregative’ Voronoi tessellations modeling contextual spaces.

B. Model for context content of a cell

Each cell of a tessellation is associated with a certain amount of context to be exchanged. 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.

Definition II.1. (Context content function) The context content function $c : B \rightarrow \mathbb{R}^+$ where B 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 $A \in B$, 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 $A_1, \dots, A_n \in B$ of A . Note that since $c(\cdot)$ 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.

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. Our idealized models for the context content function capture only some basic characteristics of such systems. In Section III we try to address the implications of this fact. 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

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

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 $\lambda_f > \lambda_a$.

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 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 α . The special case $\alpha = 1$ corresponds to scaling with no spatial correlation.

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. [8]. 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., [1].

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*

$$\frac{4 * v}{\pi} \sqrt{\lambda} \text{ crossings/unit time.} \quad (1)$$

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*

$$h + O * d. \quad (2)$$

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

The parameter O 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 $O = 1$. The effect of $O \neq 1$ can be evaluated by scaling the context content function in Assumption II.3. The parameter h can model fixed protocol overheads.

III. Analysis of Aggregative Tessellations

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 a 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.

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

$$\sqrt{\lambda_a} * (h + (\frac{\lambda_f}{\lambda_a})^\alpha * c(V_f)) < \sqrt{\lambda_f}(h + c(V_f)). \quad (3)$$

The following result, which is derived in [9], summarizes when aggregation is indeed beneficial.

Theorem III.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

$$\sqrt{\lambda_a}(h + (\frac{\lambda_f}{\lambda_a})^\alpha * c(V_f)) = \sqrt{\lambda_f}(h + c(V_f)) \quad (4)$$

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

$$\left| 1 - \frac{\sqrt{\lambda_{a,opt}}(h + c(V_a))}{\sqrt{\lambda_f}(h + c(V_f))} \right| = 1 - \frac{2\alpha}{2\alpha - 1} \frac{x}{1 + x} \left(\frac{2\alpha - 1}{x} \right)^{\frac{1}{2\alpha}}. \quad (5)$$

Note that when $\alpha < \frac{1}{2}$ the context content function scales sub-linearly in the area, which is slow enough that aggregation always helps. As shown on the top in Fig. 2 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 $h, \alpha, c(V_f)$, 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 $\alpha > \frac{1}{2}$, the context content grows quickly so more care needs to be taken in using aggregative cells. 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 bottom in Fig. 2. 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. 2 bottom 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 III.1. The relative cost reduction achieved by aggregation for the case shown on the bottom in Fig. 2 is 25.7%. A special case worth mentioning is $\alpha = 1$ i.e. the context content of a cell grows linearly with respect to its area, a plausible model for e.g. acquiring context from sensors whose context content is uncorrelated. For more details see [9].

Up to now we have assumed that a mobile is interested in the entire context content of each cell it visits. However, this need not hold in practice. Indeed, a user/application interacting with a fine grain organization could select exactly in which 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. Below, we present the following fact stating the conditions under

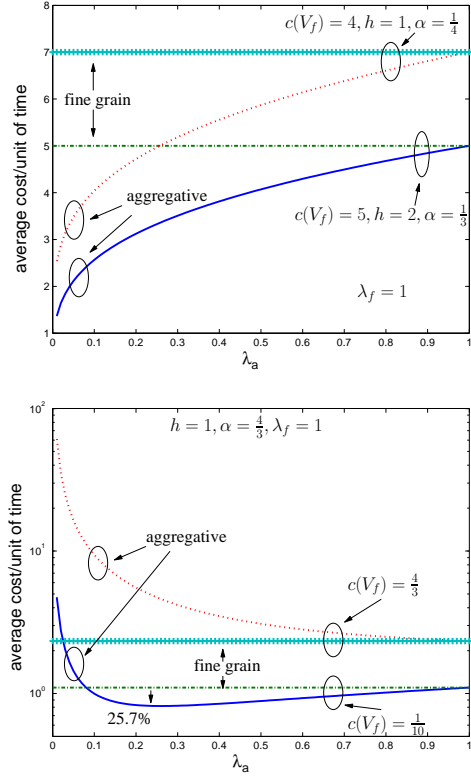


Fig. 2. Context aggregation using aggregative tessellations when $\alpha < \frac{1}{2}$ (top) and $\alpha > \frac{1}{2}$ (bottom).

which aggregation can be useful if mobiles download context selectively.

Fact III.2. A mobile interested in a fraction $p, 0 < p < 1$ of the context content of each cell it visits can benefit from aggregation if

$$p > \frac{2\sqrt{x}}{1 + x}$$

where x is as defined in Theorem III.1.

For a proof and a more complete treatment of this case see [9].

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

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, demonstrated in [9], 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.

IV. 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 it 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 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. A typical cell from the aggregative tessellation $V(\Pi_a)$ has an average context content $c(V_a)$, area $1/\lambda_a$, and will on average cover λ_f/λ_a fine grain cells. The cells covered by an

aggregative cell, have 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. A cell of the 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:

$$s = \frac{\lambda_f c(V_f) - \lambda_a c(V_a)}{\lambda_f - \lambda_a}. \quad (6)$$

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 result, where again we have relegated the derivations to [9].

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

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

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

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 $\lim_{r \rightarrow \infty} g(r) = \infty$ so the condition in Theorem IV.1 is 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 on the top in Fig. 3 and middle in Fig. 3 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.

As α reduces to $\frac{1}{2}$, there will be more redundancy in the contextual information across fine grain cells. Thus, the hierarchical approach can reduce costs by employing

larger values for λ_a , i.e. the $V(\Pi_a)$ cells do not need to be extremely coarse. However, in this case the performance of the approach based on aggregation also improves – recall from Theorem III.1 that for $\alpha < \frac{1}{2}$ aggregation always produces savings. Note that if in practice the coarseness of aggregation one can achieve is constrained, then aggregation may be more cost effective than a hierarchical organization, see top Fig. 3.




Fig. 3. Hierarchical vs. aggregative when $\alpha > \frac{1}{2}, c(V_f) = 1$ (top), $\alpha > \frac{1}{2}, c(V_f) = 10$ (middle) and $\alpha < \frac{1}{2}$ (bottom).

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 III.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 IV.1 goes to 0 as $r \rightarrow \infty$. If the average amount of context for a fine grain cell of the $V(\Pi_f)$ tessellation is extremely low or equivalently the overhead h is high then, x has a high value and there is the possibility that the hierarchical approach will not produce any cost savings. But for cases of interest where the overhead h is low or equivalently $c(V_f)$ is high, i.e., there is a lot of context per cell, the hierarchical approach can produce savings over the aggregative approach, see bottom Fig 3. Thus in this case, a designer should be careful enough to evaluate both approaches before deciding which one is better. On the bottom in Fig. 3 we observe that the aggregative approach results in 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 on the bottom in Fig. 3 outperform the aggregative approach.

V. Assessing the Cost of Surveillance in Ubiquitous Environments

Throughout this paper we have implicitly assumed the existence of a surveillance mechanism informing mobiles about cell boundary crossings. 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.
- 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.

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 $f_d = 1 \text{ Hz}$. Motivated by 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 V.1. *We assume that the average power for a direct surveillance mechanism per unit of area is given by*

$$2 * \sqrt{\lambda} * K_d * f_d * E_d. \quad (7)$$

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

$$2 * \sqrt{\lambda_a} * K_d * f_d * E_d + \lambda_0 * \frac{4 * v}{\pi} \sqrt{\lambda_a} c(V_a). \quad (8)$$

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

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

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

$$\lambda_0 * f_i * E_i. \quad (9)$$

Observe that the overall system 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. Lower bounds on the frequencies for both approaches are given in [9].

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

Fact V.4. *The direct surveillance is more efficient than the indirect surveillance if*

$$2 * \sqrt{\lambda_a} * K_d * f_d * E_d < \lambda_0 * f_i * E_i \quad (10)$$

For services offered on a 'personalized' scale i.e. $\lambda_f \sim \Theta(\lambda_0)$, $\sqrt{\lambda_a} \ll \lambda_0$ for an aggregative tessellation and the

leverage of the shared infrastructure by a direct surveillance mechanism provides significant gains.

VI. 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 gross simplifications that it would be of interest to relax, and are part of our future work.

VII. Acknowledgments

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

References

- [1] 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.
- [2] 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.
- [3] A. Schmidt, "Ubiquitous computing - computing in context," Ph.D. dissertation, University of Lancaster, 2002.
- [4] B. Brumitt, J. Krumm, B. Meyers, and S. Shafer, "Ubiquitous computing and the role of geometry," in *IEEE Personal Communications*, 2000, pp. 41–43.
- [5] 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.
- [6] 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.
- [7] D. Stoyan, W. S. Kendall, and J. Mecke, *Stochastic Geometry and its Applications*. Chichester: John Wiley and Sons, 1987.
- [8] I. Rhee, M. Shin, S. Hong, K. Lee, and S. Chonga, "On the levy-walk nature of human mobility," in *Infocom*, 2008.
- [9] A. Ziotopoulos and G. de Veciana, "Design and optimization of spatial organizations for context exchange and surveillance," ECE Department, University of Texas at Austin, Tech. Rep., 2008. [Online]. Available: <http://www.ece.utexas.edu/~gustavo/contextStochasticGeometry.pdf>