Estimating the Benefit of Location-Awareness for Mobile Data Management Mechanisms

Uwe Kubach and Kurt Rothermel

Institute of Parallel and Distributed High-Performance Systems (IPVR),
University of Stuttgart, Breitwiesenstr. 20-22, 70565 Stuttgart, Germany
Uwe.Kubach@informatik.uni-stuttgart.de

Abstract With the increasing popularity of mobile computing devices, the need to access information in mobile environments has also grown rapidly. In order to support such mobile information accesses, location-based services and mobile information systems often rely on location-aware data management mechanisms like location-aware caching, data dissemination or prefetching. As we explain in this paper, the location-awareness of such mechanisms is only useful, if the accessed information is location-dependent, i.e. if the probability with that a certain information object is accessed depends on the user's location.
Although the location-dependency of the accessed information is crucial for the efficiency of location-aware data management mechanisms and the benefit they can get out of their location-awareness, no metric to measure the location-dependency of information has been proposed so far. In this paper, we describe such a metric together with a second one for a further important characteristic of mobile information accesses, the so-called focus.

1 Introduction

Location-based services provide their users with local information depending on their current geographic position. For example, a user can ask for nearby restaurants or the shopping centers in his proximity. An important requirement for such location-based services to be beneficial is that the offered information is location-dependent, i.e. that the relevance of each information object for the user depends on his location.

Such a location-dependency is not only exploited for the pre-selection of information in location-based services but also in many other data management mechanisms supporting mobile information systems. Since they consider a user's location, these mechanisms are called location-aware. Examples are location-aware caching [11,5], dissemination [6], prefetching [4], and hoarding mechanisms [8].

So far, the location-dependency has been mostly considered as a binary characteristic of an information system, i.e. an information system respectively the information offered by the system was said to be location-dependent or not. However, we claim that there is a complete spectrum of mobile information systems, which differ in the degree of their location-dependency. This spectrum ranges from inherently location-dependent systems to location-independent systems with many nuances in between.

In inherently location-dependent systems each information object belongs to a fixed location and is only accessed when the user is located there. An example for such a system is a map application showing the user a map of his environment [13]. A browser-based mobile tourist guide is an example of a system that is neither inherently location-dependent nor location-independent. Although the user can potentially access all information objects available in the information space from any location, he will in most cases preferably access information about his current environment. In contrast, wireless web browsers [3] will mostly be used in a location-independent manner. For example, if a user looks for his stock quotes, his location will usually have no influence on which quotes he requests.

The location-dependency of the information available in a mobile information system strongly influences the amount of data a user needs at a certain location and with it the amount of data that has to be cached, disseminated or prefetched for a certain location. Therefore, the location-dependency is crucial for the efficiency of location-aware data management mechanisms. In this paper, we propose a metric that allows to quantify the location-dependency of a single information object as well as that of a whole information space. Thereby we provide a means to better understand what location-dependency is and to decide whether using a location-aware data management mechanism is suitable or not.

In addition to the location-dependency, we discuss a further characteristic of mobile information systems, which we call the focus. It also has a strong effect on the efficiency of location-aware data management mechanisms.

Finally, we use a hoarding mechanism [8], which we developed for a platform supporting location-aware applications [7], as an example to illustrate how the discussed characteristics influence the efficiency of a location-aware mobile data management mechanism.

The remainder of this paper is structured as follows: in the following section, we discuss the related work, before we introduce two metrics to measure the disparity of frequency distributions in Section 3. Next, in Section 4 and Section 5, we describe our metrics for the location-dependency and the focus. Afterwards, we evaluate the metrics and give an example for their application in Section 6. Finally, Section 7 concludes our paper.

2 Related Work

To our knowledge, no measure for the location-dependency of an information space has been proposed so far. For the location-dependency it is important, how information requests are distributed over the area, in which the information system is accessible. Such a spatial distribution also plays an important role in wireless ad hoc sensor networks. There, the distribution of the sensors strongly influences the coverage of the sensor network.

However, most work in this area is focused on the question of how to arrange a set of sensors to observe an area, room, or building as good as possible. Only a few work has been done on measuring the coverage [9,10]. In these articles, coverage is measured in terms of how well an object that crosses the area covered by the sensor network can be observed. The inequalities in the spatial distribution of the sensors are not considered. In contrast, the inequalities in the spatial distribution of the information requests have a strong impact on the location-dependency and must therefore be considered when measuring this dependency.

Inequalities in spatial distributions are often considered in social sciences, e.g. income or wealth inequalities [1,12]. There, the Herfindahl coefficient and the Gini coefficient are widely used to measure the inequalities. These two coefficients are also the basis of our metric. In the following section, they are discussed in detail.

3 Preliminaries

Basically, our idea to measure the location-dependency of an information object is to consider how the requests for the object are distributed over the area in which the information system can be accessed. If many requests are concentrated on a few locations, the object is obviously only relevant at certain locations, i.e. it is strongly location-dependent. However, if the requests are equally distributed over the plane, the object's relevance for the users does not depend on their location, i.e. the object is location-independent.

In spatial statistics two coefficients, namely the Herfindahl and the Gini coefficient, are commonly used to describe such inequalities in frequency distributions. Since we also use these coefficients for our metrics, we describe them in this section.

A frequency distribution is a function $f : O \rightarrow \mathbb{N}$ that assigns to each unit of observation $o \in O$ the absolute frequency with that it occurred during an observation. An example for a frequency distribution is a function that assigns to each information object within an information space the number of requests that occurred for the object during a certain period of time. In the following, $N = |O|$ denotes the number of different units of observation that might occur. We assume that the units of observation are consecutively numbered from 1 to N. Hence, $f(i)$ denotes the frequency with that the unit of observation i occurred.

3.1 Herfindahl Coefficient

The Herfindahl coefficient HC is defined as the sum of the squares of the relative frequencies with which the units of observation occur:

$$HC = \sum_{i=1}^{N} \left(\frac{f(i)}{\sum_{j=1}^{N} f(j)} \right)^2$$

For the Herfindahl coefficient values between $\frac{1}{N}$ and 1 are possible. It is minimal, when the frequencies are equally distributed, whereas the maximum value is reached, when only one certain unit of observation occurs at all.

3.2 Gini Coefficient

The Gini coefficient is based on the Lorenz curve, which can be used to depict inequalities in frequency distributions. To construct a Lorenz curve, the frequencies $f(i)$ first have to be sorted in an increasing order, i.e. we have to ensure that

$$f(1) \leq f(2) \leq \ldots \leq f(N-1) \leq f(N).$$

The Lorenz curve is then constructed as polygon through the points

$$(0,0),(u_1,v_1),(u_2,v_2),\ldots,(u_{N-1},v_{N-1}),(u_N,v_N).$$

The coordinates of each point (u_j, v_j) are calculated as follows:

$$u_j = \frac{j}{N} \qquad \text{and} \qquad v_j = \sum_{i=1}^{j} \frac{f(i)}{\sum_{k=1}^{N} f(k)}$$

This means that the Lorenz curve is obtained by plotting the cumulative relative frequencies with that the units of observation occur during the observation period against the share of considered units of observation. If the considered frequencies are equally distributed, the Lorenz curve is equivalent to the 45 degree line. The bigger the area between the Lorenz curve and the 45 degree line is, the higher is the inequality in the corresponding frequency distribution.

In Table 1 an example for a frequency distribution is given. It shows how 100 information requests might be distributed over an information space containing 10 information objects, i.e. $N = 10$. Besides the absolute request frequencies $f(i)$, the share of the considered units of observation, and the cumulative relative frequencies are included in the table. Figure 1 shows the corresponding Lorenz curve together with the 45 degree line.

Object ID	Number of Requests	Share	cumulative rel. frequencies
1	4	0.1	0.04
2	5	0.2	0.09
3	5	0.3	0.14
4	6	0.4	0.20
5	7	0.5	0.27
6	8	0.6	0.35
7	9	0.7	0.44
8	12	0.8	0.56
9	16	0.9	0.72
10	28	1.0	1.0

Table 1. A frequency distribution.

The Gini coefficient GC is defined as the ratio of the area enclosed by the Lorenz curve and the 45 degree line to the area enclosed by the 45 degree line and the x-axis. It can be calculated as follows:

$$GC = \frac{2 \cdot \sum_{i=1}^{N} i \cdot f(i)}{N \cdot \sum_{i=1}^{N} f(i)} - \frac{N+1}{N}$$

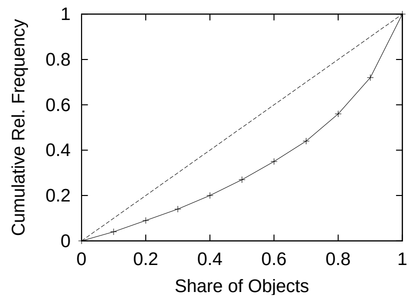

Figure 1. Lorenz curve.

Like the Herfindahl coefficient, the Gini coefficient is minimal ($GC = 0$) when the frequencies are equally distributed, and maximal ($GC = \frac{N-1}{N}$), when only one certain unit of observation occurs at all. Often the normalized Gini coefficient $GC^* = \frac{N}{N-1} \cdot GC$ is used in order to get values ranging from 0 to 1.

3.3 Comparison

The difference between the two coefficients is that the Gini coefficient solely reflects relative concentrations, whereas the Herfindahl coefficient also expresses the absolute concentration of a frequency distribution.

Thereby, absolute concentration means that the observed incidences are distributed over a small set O of possible units of observation, e.g. if an information space only contains a few information objects. Thus, the absolute concentration can be high, although the frequencies are equally distributed. Whereas relative concentration means, that a big part of the observed incidences is distributed over a small part of all possible units of observation. This implies an unequal frequency distribution.

4 Location-Dependency

In this section, we describe our metrics for the location-dependency of a single information object and that of a complete information space. In both metrics we assume that the coverage area of the information system, i.e. the area where the system can be accessed, is separated into a set S of equally sized squares, which do not overlap each other.

4.1 Definitions

As explained above, the location-dependency of an information object should describe, how strong the requests for this object are concentrated on certain locations. Thus, we define the location-dependency of a single information object as follows:

Definition 1. *Let i be an information object and $f : S \to \mathbb{N}$ a frequency distribution that assigns to each square $s \in S$ the frequency with that the information object i has* *been accessed in the square s. Then, the normalized Gini coefficient of the distribution f is called the location-dependency of the information object i.*

We chose the Gini coefficient for the definition of the location-dependency because we want our metric to reflect solely the distribution of the information accesses over the grid. The Herfindahl coefficient would also reflect the number of considered squares.

In order to make estimations about the benefit of location-awareness for a mobile data management mechanism, it is not enough to analyze only one single information object. Therefore the whole information space has to be considered. Thus, we also define the location-dependency of an information space:

Definition 2. *Let IS be an information space and $f : IS \to \mathbb{N}$ a frequency distribution that assigns to each object in the information space the number of requests that occurred for this object during the observation period. Let furthermore $l : IS \to [0, 1]$ be a function that assigns to each object in the information space its location-dependency. Then, the location-dependency L of the information space is defined as the weighted average of the location-dependencies of the objects in the information space:*

$$L = \sum_{i \in IS} \frac{f(i)}{\sum_{j \in IS} f(j)} \cdot l(i)$$

According to our definition, the location-dependency of an information space lies between 0 and 1, where 0 indicates a location-independent information space and 1 indicates an inherently location-dependent information space.

4.2 Discussion

Unfortunately, a high location-dependency is only a necessary condition for a beneficial effect of a data management mechanism's location-awareness. That it is not a sufficient condition can be seen from the access matrix given in Figure 2. The matrix shows how often each of the five considered information objects $i_1, \ldots, i_5$ are accessed in each of the five considered squares $s_1, \ldots, s_5$. The value given in row m and column n states how often the information object i_n is accessed in square s_m.

Since each information object is only accessed in one square, we have a strongly location-dependent information access ($L = 1$). However, a location-aware mobile data management mechanism can not profit from this location-dependency, because all objects are accessed in the same square s_3. For example, a location-aware filter mechanism could not filter out any information objects, if the user is located in square s_3. A location-aware caching strategy also can not profit from knowledge about the user's position, since in all squares each information object is of the same value. Hence, a strong location-dependency is only a hint that exploiting location-awareness might be beneficial.

So far we have not made any assumption about the size of the squares in the set S. However, it has to be reasonably chosen. If the squares are too big, i.e. the resolution of the grid is too small, inequalities in the frequency distributions are hardly recognized. The worst case is, if the whole coverage area of the information system consists of just

	i_1	i_2	i_3	i_4	i_5
s_1	0	0	0	0	0
s_2	0	0	0	0	0
s_3	10	10	10	10	10
s_4	0	0	0	0	0
s_5	0	0	0	0	0

Figure 2. Strongly location-dependent access matrix.

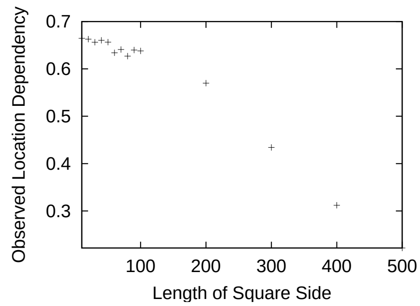

Figure 3. Location-dependencies observed with different grid resolutions.

one square. Then, inequalities can not be observed at all. Because, if only one square is considered, the access frequencies are trivially the same in all squares.

In Figure 3 this effect is illustrated. It shows the observed location-dependencies for a certain information object, which we got using different grid sizes. Although we always used the same, quite location-dependent, access pattern, we got different values for the observed location-dependency. As expected, the observed location-dependency decreases with a decreasing grid resolution, i.e. an increasing length of the squares' sides. For a length of a square side of up to $100m$, what corresponds to 100 squares in the considered $1000m$ x $1000m$ coverage area, we get acceptable values for the location-dependency. For bigger squares the accuracy becomes poor and the observed location-dependency is noticeably lower than the actual one.

However, the square size should not be too small either because smaller squares obviously imply higher storage costs and computational costs for the determination of the location-dependency. Moreover, if the squares are smaller, more requests for an information object have to be observed to make a general statement on how the requests to the object are distributed over the squares. This means that the observation periods have to be longer.

5 Focus

In this section, we introduce the focus of a mobile information system as a second characteristic that has a strong impact on the efficiency of location-aware mobile data management mechanisms.

5.1 Definitions

As shown in the previous section, a high location-dependency alone is not enough to guarantee a high benefit from considering a user's location. In addition, it is also important that at each location, only a specific part of all available information is accessed. This means, that we also have to consider, how the information requests originating from each location are distributed over the information space. Therefore, we define the focus as follows:

Definition 3. *Let $s \in S$ be a square and let $f : IS \rightarrow \mathbb{N}$ be a frequency distribution that assigns to each information object in IS the frequency with that it has been accessed within the square s during the observation period. Then, the focus within the square s is the Herfindahl coefficient of the distribution f.*

This time we chose the Herfindahl coefficient for our definition, since now the absolute number of information objects requested in a square is also important, not only the inequality in the frequency distribution. For example, the number of objects that should be cached or prefetched for a certain square grows with the absolute number of objects that the average user requests there.

For a good estimation of a location-aware mechanism's performance, all squares have to be considered. The contribution of each square should correlate with the number of requests originating from this square, because the focus in a square with only a few requests will have less impact on the performance of a location-aware mechanism than that of a square with many requests will have. Thus, we define the the focus of an information system as the weighted average of the focuses observed in each square $s \in S$:

Definition 4. *Let $fr : S \rightarrow \mathbb{N}$ be a frequency distribution that assigns to each square $s \in S$ the number of information requests that occurred there during the observation period. Furthermore, let $fo : S \rightarrow [\frac{1}{N}, 1]$ be a function that assigns to each square $s \in S$ the focus observed there. Then, the focus F of the according information system is defined as follows:*

$$F = \sum_{s \in S} \frac{fr(s)}{\sum_{i \in S} fr(i)} \cdot fo(s)$$

With this definition values between $\frac{1}{N}$, where N is the number of objects in the information space, and 1 are possible for the focus of an information system. The value $\frac{1}{N}$ indicates that within each considered square all objects of the information space are accessed with the same frequency. A value of 1 for the focus indicates that within each square only one information object is requested.

5.2 Discussion

Similar to the location-dependency, a high focus alone is also no guarantee for a beneficial use of a location-aware mechanism. For an example, consider the access matrix in Figure 4. With this access pattern, we will get a high focus, since in each square only one information object is accessed. However, a location-aware mechanism can not profit from the knowledge of a user's position, since information object i_3 is preferred at any location, what can easily be observed without any location-awareness.

Note, that the location-dependency is 0 for the access pattern given in Figure 4. In fact, if both the location-dependency and the focus are high (see Figure 5) the benefit from using location-awareness will also be high. Because then, at each location many requests are concentrated on a small part of the available information objects (high focus) and the preferred information objects will differ between different locations (high

	i_1	i_2	i_3	i_4	i_5
s_1	0	0	10	0	0
s_2	0	0	10	0	0
s_3	0	0	10	0	0
s_4	0	0	10	0	0
s_5	0	0	10	0	0

Figure 4. Highly focused access matrix.

	i_1	i_2	i_3	i_4	i_5
s_1	10	0	0	0	0
s_2	0	10	0	0	0
s_3	0	0	10	0	0
s_4	0	0	0	10	0
s_5	0	0	0	0	10

Figure 5. Highly focused and strongly location-dependent access matrix.

location-dependency). Thus, location-awareness can then be used to identify the objects which are specifically preferred at each location. Without considering a user's location, all information objects could seem to be of equal popularity. This would be, for example, the case with the access pattern given in Figure 5.

6 Evaluation

In this section, we determine the location-dependency and the focus of some typical access patterns in order to give an idea about the meaning of concrete values obtained with the metrics. We also analyse how important parameters of the access patterns influence these values. Finally, we use a location-aware hoarding mechanism as an example to show how the location-dependency and the focus may effect the efficiency of a location-aware mobile data management mechanism.

6.1 Location-Dependency

For the analysis of the location-dependency, we assume that the considered information system covers a quadratic area of $1000m$ x $1000m$. For the determination of the location-dependency the coverage area was separated into 1600 squares measuring $25m$ x $25m$. We always considered the requests for one single information object. In order to specify the origin of an information request we used a coordinate system with a resolution of 1 meter, with $(0, 0)$ representing the upper left corner of the coverage area and $(1000, 1000)$ the lower right corner.

Location-Independent Access For the simulation of a location-independent access, we assumed that the coordinates from which the requests for the considered information object originate are both equally distributed over the range from 0 to 1000. In Figure 6(a) a graphical representation of this access pattern is given. It shows the number of requests for the considered information object that occurred in each square during the observation period. The brighter a square is depicted the higher is the number of requests originating from this square. Since in all squares approximately the same number of requests occurred, the brightness of the squares does not differ much. And, as expected, we get a very low value for the location-dependency: 0.023.

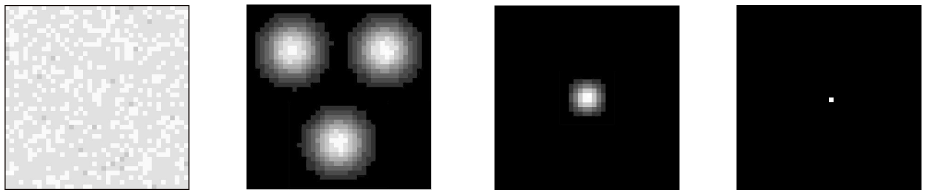

(a) (b) (c) (d)

Figure 6. The considered access patterns.

Hot Spot Access For the simulation of the wide spectrum of location-dependent information accesses, we used an access pattern, which we call the hot spot access pattern. In this pattern, the requests for the considered information object primarily originate from one or more hot spots. If there is more than one hot spot, we first randomly select one of the hot spots as the origin of a request. Each hot spot is selected with the same probability. Afterwards, the coordinates, where the information object was requested, are determined. Therefore, we assume that the x- and y-coordinate are distributed according to a gaussian distribution around the coordinates of the selected hot spot.

Thus, this pattern has two important parameters, which allow to simulate information accesses with quite different location-dependencies. These two parameters are the number of hot spots and the standard deviation σ of the gaussian distribution. Figures 6(b) and 6(c) show examples for the hot spot access pattern, with three hot spots and a standard deviation of 100, and with one hot spot and a standard deviation of 50. The plots in Figure 7 show the location-dependencies we get with one and three hot spots for different values of σ.

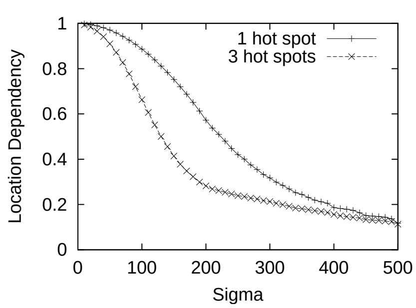

Figure 7. Location-dependencies with different standard deviations.

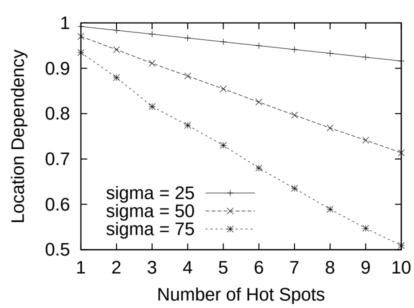

Figure 8. Effect of the number of hot spots on the location-dependency.

In Figure 8 the location-dependencies are shown, which we got for different numbers of hot spots. The hot spots were evenly distributed over the information system's coverage area. The three plots show the results for standard deviations of $\sigma = 25$, $\sigma = 50$, and $\sigma = 75$.

Inherently Location-Dependent Access With an inherently location-dependent information access, all requests for a certain information object are done at the same location, thus they all originate from the same square (see Figure 6(d)) and, not surprisingly, the location-dependency is 1.

6.2 Focus

In the previous section, we had to consider the whole coverage area of the information system but only one information object to make statements about the location-dependency. In contrast, we now have to consider only one square of the coverage area, but all information objects requested there. Analogously to our examinations of the location-dependency, we now determine the focus of typical access patterns and analyze the effect of the patterns' parameters on the focus.

Uniform Distribution As mentioned above, the focus does not only reflect inequalities in the frequency distributions of the information requests but also the absolute number of locally accessed information objects. Hence, we can get high focuses, although the local access frequencies are equally distributed, if only a small number N of different objects is locally requested at all. For a uniform distribution the focus F will be $F = \frac{1}{N}$.

Zipf-like Distribution In [2] it has been shown that the distribution of the requests over single web pages follows a Zipf-like distribution. In such a distribution the relative probability of a request for the i'th most popular page is proportional to $\frac{1}{i^\alpha}$. The observed value of α varies between the different considered traces, ranging from 0.64 to 0.83.

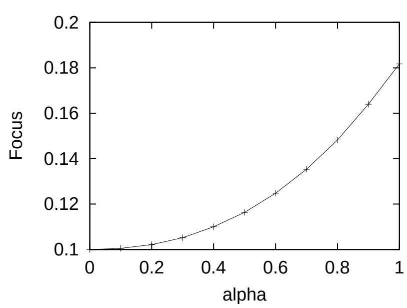

Figure 9. Focus depending on the parameter α of a Zipf-like distribution.

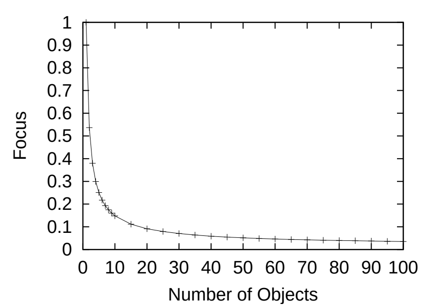

Figure 10. Focus depending on the total number of different objects accessed locally.

We assumed such a distribution for the information requests observed within the considered square. We experimented with different values of the parameter α. A higher value for α means a stronger concentration on the most popular information objects.

Figure 9 shows the focus depending on the parameter α. We assumed that a total of 10 different objects is accessed in the considered square.

As mentioned above, the focus also reflects an absolute concentration, i.e. the total number of different objects that are accessed in a square is reflected in our metric. Figure 10 shows how this number influences the focus. To get these results, we simulated a Zipf-like distributed information access and varied the total number of locally requested objects. The parameter α of the Zipf-like distribution was this time set to a fixed value of 0.8.

6.3 An Example Application

In this section, we illustrate the effects that the location-dependency and the focus might have on location-aware mobile data management mechanisms using a hoarding mechanism, we proposed in [8], as an example.

This mechanism aims to allow the users of a mobile information system to access information during disconnections, i.e. when no network is available. Therefore, the mechanism tries to predict the information objects a user will probably access during a future disconnection and transfers them to the user's device as long as there is a connection. Thus, the information is already locally available, when the user accesses it during the disconnection. In order to predict the objects that a certain user will probably need, the mechanism uses knowledge about the objects' popularities at each location and predictions of the user's future movement (for details see [8]).

A measure often used to rate the efficiency of hoarding and caching mechanisms is the hit ratio. This is the ratio between the number of requests that can be answered with information objects that are stored locally and the total number of requests a user makes.

If there is no location-dependency and no focus, i.e. if all objects are requested everywhere with the same probability, the hoarding mechanism will not be able to benefit from its location-awareness. In this worst case scenario, the user will only get an average hit ratio of $\frac{m}{n}$, where m is number of information objects that can be stored locally on the user's device and n is the total number of objects in the whole information space. Usually, m will be much smaller than n. If m is bigger than n or equal to n, we will trivially get a hit ratio of one.

Finally, let us consider an example with a location-dependency of 1 and a high focus. In this example, at each square s a specific set of information objects O_s is requested, where $O_i \cap O_j = \emptyset$, if $i \neq j$ and $|O_s| = \frac{n}{|S|}$. n is again the total number of objects in the information space and S the set of all squares within the information systems' coverage area. Furthermore, let us assume that a user located in square s accesses all objects in O_s with the same probability. Then, we will get an average hit ratio of $\min(\frac{m \cdot |S|}{n \cdot k}, 1)$, where k is the number of squares the user will visit during the disconnection. Figure 11 shows the hit ratios for this and the previous example depending on the number of squares a user visits during a disconnection. The parameters were chosen as follows: $|S| = 100$, $m = 100$, $n = 5000$.

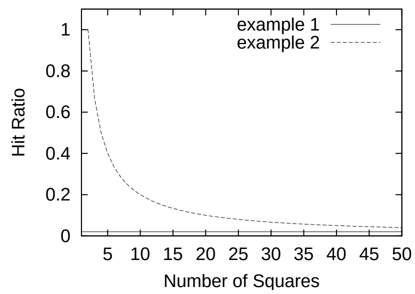

Figure 11. Hit ratios of a location-aware hoarding mechanism.

7 Conclusion

In this paper, we claimed that there are not only inherently location-dependent and location-independent information systems, but that there is a complete spectrum of systems differing in their degree of location-dependency. We supported this claim by giving examples of information systems with different location-dependencies.

Furthermore, we proposed a metric, which can be used to measure the location-dependency of an information system and which helps to better understand what location-dependent information is. We explained why considering only the location-dependency of an information system is not enough in order to estimate the benefit a location-aware data management mechanism can get from its location-awareness. Additionally, we proposed a second metric, the focus, which together with the metric for the location-dependency finally provides a good means for this estimation.

In the evaluation part of the paper, we analyzed the location-dependency and the focus of various access patterns to give an idea of the values that can be expected with typical access patterns. Finally, we used a location-aware hoarding mechanism to show how the location-dependency and the focus can effect the efficiency of a mobile location-aware data management mechanism.

For the near future, we plan to combine the two metrics into only one metric completely characterizing an information space. To do so, we first have to further analyze the proposed metrics with a simulation tool which allows to directly manipulate the location-dependency and the focus of a considered information space. With such a tool we could more exactly analyze how the proposed metrics effect different typical location-aware mechanisms. This will also help us to further refine our metrics such that they also distinguish access patterns, which are, with our current metrics, considered as equivalent.

Finally, we plan to adjust our metrics to other fields of application, e.g. the measurement of coverage in a wireless ad-hoc sensor network (see Section 2). In contrast to existing metrics, our metric would measure the expected coverage for an average object, not the maximal or minimal possible coverage. Another field of application for our metrics are mobile ad-hoc networks. There, inequalities in the average spatial distribution of the network nodes are crucial for the delay and the throughput of the network.

References

1. A. B. Atkinson. *Wealth, Income and Inequality*. Oxford University Press, Incorporated, second edition, 1981.
2. L. Breslau, P. Cao, F. Li, G. Philips, and S. Shenker. Web caching and zipf-like distributions: Evidence and implications. In *Proceedings of the IEEE INFOCOM '99*, pages 126–134, New York City, NY, USA, March 1999.
3. H. Chang, C. Tait, N. Cohen, M. Shapiro, S. Mastrianni, R. Floyd, B. Housel, and D. Lindquist. Web browsing in a wireless environment: Disconnected and asynchronous operation in artour web express. In *Proceedings of the Third Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '97)*, pages 260–269, Budapest, Hungary, September 1997.
4. J. Chim, M. Green, R. Lau, H. Leong, and A. Si. On caching and prefetching of virtual objects in distributed virtual environments. In *Proceedings of the Sixth ACM International Conference on Multimedia (ACM MULTIMEDIA '98)*, pages 171–180, Bristol, UK, 1998. `http://www.acm.org/sigmm/MM98/electronic_proceedings/chim/index.html`.
5. S. Dar, M. Franklin, B. Jónsson, D. Srivastava, and M. Tan. Semantic data caching and replacement. In *Proceedings of the 22nd International Conference on Very Large Data Bases*, Bombay, India, 1996.
6. N. Davies, K. Cheverst, K. Mitchell, and A. Friday. Caches in the air: Disseminating information in the guide system. In *Proceedings of the Second IEEE Workshop on Mobile Computing Systems and Applications (WMCSA '99)*, pages 11–19, New Orleans, LA, USA, February 1999.
7. F. Hohl, U. Kubach, A. Leonhardi, K. Rothermel, and M. Schwehm. Next century challenges: Nexus - an open global infrastructure for spatial-aware applications. In *Proceedings of the Fifth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '99)*, pages 249–255, Seattle, WA, USA, August 1999.
8. U. Kubach and K. Rothermel. Exploiting location information for infostation-based hoarding. In *Proceedings of the Seventh Annual ACM SIGMOBILE International Conference on Mobile Computing and Networking (MobiCom '01)*, pages 15–27, Rome, Italy, July 2001.
9. S. Meguerdichian, F. Koushanfar, M. Potkonjak, and M. B. Srivastava. Coverage problems in wireless ad-hoc sensor networks. In *Proceedings of Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies (Infocom 2001)*, volume 3, pages 1380–1387, Anchorage, AK, USA, April 2001.
10. S. Meguerdichian, F. Koushanfar, G. Qu, and M. Potkonjak. Exposure in wireless ad-hoc sensor networks. In *Proceedings of the Seventh Annual ACM SIGMOBILE International Conference on Mobile Computing and Networking (MobiCom '01)*, pages 139–150, Rome, Italy, July 2001.
11. Q. Ren and M. Dunham. Using semantic caching to manage location dependent data in mobile computing. In *Proceedings of the Sixth Annual ACM SIGMOBILE International Conference on Mobile Computing and Networking (MobiCom 2000)*, pages 210–221, Boston, MA, USA, August 2000.
12. A. Sen. *On Economic Inequality*. Clarendon Paperbacks, second edition, 1993.
13. T. Ye, H.-A. Jacobsen, and R. Katz. Mobile awareness in a wide area wireless network of info-stations. In *Proceedings of the Fourth Annual ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom '98)*, pages 109–120, Dallas, TX, USA, 1998.