In Proc. of 12th Open European Summer School (EUNICE), Stuttgart, Germany, September 2006

Investigation of Geographical Routing Enhancements for Location Based Push Services

S. D. Hermann and A. Wolisz
Telecommunication Networks Group
Technische Universität Berlin
Einsteinufer 25, 10587 Berlin, Germany
Email: [hermann, wolisz]@tkn.tu-berlin.de

***Abstract*— Location based push services will offer customers the possibility to select geographical regions for their services or to distribute information to them.**

A convenient way to disseminate information in a geographical region is the utilization of geographical routing. Messages can be forwarded via geometrical operations through a carrier access network to the mobile devices in a selected region.

Geographical forwarding requires non-negligible processing power in the intermediate nodes. For its reduction, we propose the utilization of a storage and computational enhanced node (GEN), filtering similar and identical target regions for the reutilization of recent forwarding decisions.

In this paper, the required efficient geometrical filtering mechanism for this approach are investigated. We show that the computational effort can be significantly reduced by appropriate geometrical approximations of the target regions and deployment of caches.

I. Introduction

By location based services (LBS) [1], useful information can be provided to users of transportable wireless devices such as cellular phones, personal digital assistants and notebooks, based on their current geographical positions.

Generally, two classes of LBS can be distinguished, the location based *pull* and *push* services. In the first case, users actively request a service from a provider via their mobile devices. For instance, by sending a SMS message with the respective keyword to a provider, users can request the location of the closest automated teller machine, doctor or restaurant.

Contrary, location based push services are to be provided without a prior interaction between the initiator of the service and the users. They offer customers the possibility to select certain spatial areas as target regions for services and information. The services can be triggered if a users changes its spatial relation to a region (see [2]) or information can be directly sent to the users in it. These services yield a wide variety of new applications. For instance, warnings can be sent to traffic participants who approach an accident. Additionally, there exist many commercial applications like the sending of localized advertisements and coupons or the dissemination of video streams to users being in certain regions, e.g. with additional information for attendees of a sight-seeing tour.

One of the main differences between the pull and push services is the required availability of the mobile device positions. For pull services, the position of a device is needless as long as no request is issued by it, and can be determined immediately after a provider received a request. The provisioning of push services requires the maintenance of the current locations of all mobile devices, e.g. in a suited database. Information could then be sent to the devices in a certain region by a spatial query in this database and the emission of the information to the returned addresses of the devices.

Obviously, the provisioning of the push services by such a database approach would involve a huge amount of signaling information, since the positions of all mobile devices have to be continuously transmitted to the database server and information has to be send via unicast to the devices.

A more convenient mechanism for the dissemination of information in the target regions is the utilization of geographical routing [3]. Messages contain the coordinates of the geographical target regions and are forwarded by means of geometrical operations. But geographical forwarding decisions require more complex computations than e.g. the relatively simple binary operations for IP routing, which leads to non-negligible processing overhead in the intermediate nodes.

For the reduction of the processing effort, caches can be deployed for forwarding decisions of messages to recurring, identical target regions. After making the forwarding decision based on geometrical computations for a message, each intermediate node caches the result and uses the description of the region or other message header parts as key. If subsequent messages have to be forwarded, the nodes perform lookups in their caches. If a forwarding decision already exists, it will be reused; otherwise the routing will be conducted with the geometrical computations. This approach introduces a certain overhead for the caches and lookups and should be used in an efficient manner.

Two main challenges thus arise when the push services should be provided by means of geographical routing:

- The additional processing load of the intermediate nodes due to the geometrical computations must be reduced, which is critical especially for computational restricted intermediate nodes
- Existing forwarding decisions should be efficiently reused without introducing large overhead to the intermediate nodes

To overcome the aforementioned problems, we propose the delegation of the geometrical operations to computational and storage enhanced nodes in the network, possessing mechanisms for the identification of similar requested target regions.

The rest of this paper is organized as follows. A model for the considered system is described in the next section. In section III, we present how location based push services can be realized by means of geographical routing and introduce our approach for the reduction of the computational effort. The focus in this paper is on the investigation of similar target region filtering and the reduction of the processing effort for the similarity checks, which will be described in section IV. In section V, we present some numerical results. Related work is presented in section VI. Section VII concludes the paper and gives an outlook to our ongoing work.

II. System Model

In the sequel, we consider a large rectangular section of a geographical territory G with edge lengths G_x and G_y. The territory is covered by a wireless carrier access network with a large number of base stations. Each of the base stations covers a small part of the whole territory, providing network access to the users of mobile devices who are located in that part. The base stations are connected via intermediate nodes to a gateway of the carrier, connecting the access network with the rest of the Internet.

In the territory, customers choose regions (i.e. the mobile devices in them) as targets for their messages. The target regions are modeled according to the model that has been proposed by Navas and Imielinski in [4] for the evaluation of geographical routing mechanisms. We investigate a mixture of polygonal target regions with a number of vertices v, being chosen from an uniform distribution on a bounded interval $[v_{min}; v_{max}]$. For the construction of a polygonal region, a center is defined first. Then, in equally spaced angular sectors of size $2\pi/v$, the vertices are generated by selecting the distance to the chosen center from an uniform distribution on the interval $[d_{min}; d_{max}]$. Afterwards, the polygon is rotated by an angle $\alpha \in [0; 2\pi[$.

For the spatial distribution of the regions being selected in the geographical territory section, we consider two different cases. In the first case, the centers of the regions are uniformly distributed over the territory, i.e. the probability density function of the x coordinate (and the y coordinate, respectively) is

$$f(x) = \begin{cases} \frac{1}{G_x} & \text{if } 0 \leq x \leq G_x, \\ 0 & \text{else} \end{cases} \quad (1)$$

In a real scenario, the target regions may have a certain concentration at points of interests in the territory. Thus, the x and y coordinate of the centers are chosen from a Beta distribution with probability density function

$$f(x) = \begin{cases} \frac{x^{\mu_1 - 1}(G_x - x)^{\mu_2 - 1}}{B(\mu_1, \mu_2) \cdot G_x^{\mu_1 + \mu_2 - 1}} & \text{if } 0 \leq x \leq G_x, \\ 0 & \text{else} \end{cases} \quad (2)$$

In both cases, regions are allowed to overlap the borders of the territory section without being truncated.

III. Providing Location Based Push Services via Geographical Routing

In this section, we present the provisioning of location based push services using geographical routing. After a description of the existing mechanisms, we present our new approach for the reduction of the computational effort of the intermediate nodes in the network, relying on geocast enhanced nodes (GEN).

A. Approach Based on Current Geographical Routing

Providers of location based push services offer them via portal servers in the Internet to customers. If an initiator of a service wants to send information to a certain region in the geographical territory, he contacts one of the portals. The portal of the provider possesses the knowledge about the appropriate carrier of an access network that covers the target region and sends a message containing the description of the requested service through the Internet to the gateway of the carrier.

After the message arrived at the gateway, an appropriate geographical addressed message is generated for its content. The message will then be forwarded by means of geographical routing through the carrier network to the target region.

In the recent years, several geographical addressing and routing mechanisms for the dissemination of messages in target regions (also known as *geocast*) have been proposed, whereas many of them are intended for the usage in ad hoc networks. An approach being suited for the deployment in the considered cellular networks has been proposed by T. Imielinski and J. C. Navas in [5], which will be briefly described in the following paragraphs.

The approach relies on three different components, GeoRouters, GeoNodes and GeoHosts. GeoRouters use the description of the geographical target region (e.g. in form of a closed polygon with coordinates as its vertices) to forward geographically addressed messages. The GeoRouters can be considered as standard network routers with additional geographical functions. Since not every router in a network is assumed to possess this functionality, messages can be forwarded in an overlay manner through tunnels between the GeoRouters.

The messages are forwarded by the routers until they arrive at the GeoNodes. These nodes are responsible for the dissemination of messages in a certain part of the geographical territory (i.e. a cell or subnet) that has been assigned to them. Each part of the territory has at most one GeoNode. The task of the node is to store the received messages and to periodically send them to the part of the territory as long as their lifetimes are not expired. In the case of missing GeoNodes, messages are disseminated only once in the territory.

The territory part being assigned to a GeoNode may be larger than a target region, requiring an additional filtering of the messages. This is performed by a GeoHost software running on the mobile devices. Two different variants for the filtering have been proposed. In the first variant, the GeoNode sends a list of available messages, their destination regions and assigned multicast addresses to the mobile devices. After the GeoHost detected that the current position of the device is inside a destination region, software clients can receive the message by associating with the corresponding multicast address. In the second variant, the original messages are directly sent to the devices, and the filtering is performed by the running GeoHost applications.

For the routing of the messages, appropriate mechanisms are required. The GeoRouters in the network form a hierarchical tree structure, with the GeoNodes as leaves. The part of the territory a GeoNode is responsible for is called its service area. Each GeoNode sends its service area (or a polygonal approximation) to the next GeoRouter in the hierarchy. After receiving the service areas from all attached GeoNodes, the GeoRouter will be able to forward a geographically addressed message to them by performing geometrical intersection checks with the target region in the message and the service areas of the GeoNodes. The message will be forwarded to all GeoNodes whose checks lead to a positive result.

If there is another GeoRouter in the hierarchy above the first GeoRouter, it has to send the service areas to this upper router. Instead of sending the single areas which would result in a huge amount of areas in the upper routers, it computes and sends the geometrical union of them. The upper router considers the merged areas as the service area of the first router and performs the routing in the same way.

B. Load Reduction by Forwarding Decision Reutilization

One of the problems of geographical routing are the required geometrical computations in the intermediate nodes. In earlier studies, it has been shown that the duration of the forwarding decisions is more than $4\,000$ times larger than that for IP (see [6]).

Thus, we propose the reutilization of existing forwarding decisions in the intermediate nodes. The approach is based on a computational and storage enhanced node with geocast functionalities (GEN) near the gateway of a carrier access network. Again, the gateway interprets the request message and generates an appropriate geocast message for its content. The message is not sent directly to the target region but to the GEN before, which performs a check if another message has been recently sent to an identical region. An example of a carrier access network is displayed in figure 1.

Usually, neither the coordinates of a target region nor the positions of the mobile devices can be exactly determined. Therefore, the GEN checks not only for completely identical regions, but also for those being similar to the new region.

For this mechanism, a definition of the similarity has to be introduced. Assuming that all parts of the considered geographical territory are equally important, the areas of the target regions are a suited basis for the determination of similarity. Considering the intersection of two target regions, we have to deal with three different areas, the area a_t of the new target region t, the area a_c of the cached region c and the intersection area of both regions a_i.

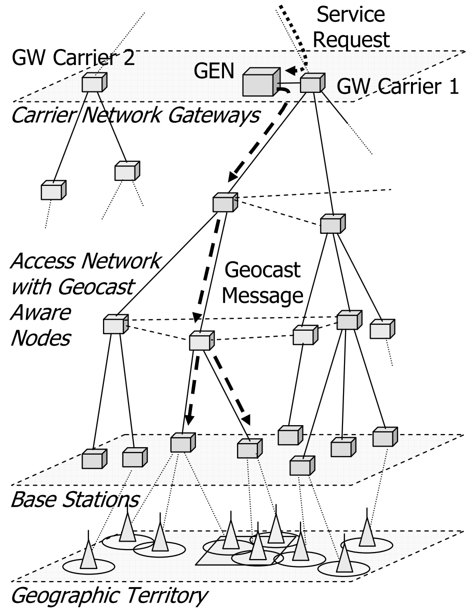

Fig. 1. Example of a carrier access network with geocast enhanced node (GEN) near the gateway

Several definitions of similarity exist which contain these three areas, an overview is e.g. presented in the paper [7]. Following the definition of similarity form Walker et al., we chose

$$\mathrm{similarity}(t,c) := \min\left(\frac{a_i}{a_t}, \frac{a_i}{a_c}\right) \tag{3}$$

The range of the similarity function is $[0,1]$, whereas a similarity of 0 means that the regions do not intersect and 1 means that the regions are identical.

If no recently requested identical region exists, the check in the GEN will return the region with the highest similarity. Depending on policies if the similarity satisfies a minimum degree, the new message can be forwarded to that region by reusing the formerly established forwarding decisions in the intermediate nodes.

The policies for the similarities have to be determined. They might depend on the concrete application, i.e. the requirements for localized advertisements may be lower than that for safety warnings. If no appropriate recently selected target region exists, the message will be sent with standard geocast operations through the network.

In order to perform the similarity checks with the regions that have been recently selected as targets, the GEN possesses two basic components. The first one is a spatial database caching the target regions of previous messages (in the sequel, we will refer to this database as *cache*). If a new message arrives at the GEN, a cache lookup is performed which returns a superset of candidates consisting of all cached target regions that are probably similar to the new region. The task of the second component of the GEN is to filter the region out of the candidate superset that is identical with (or possesses the highest similarity to) the new target region. This component will be called *filter* in the sequel.

One of the problems is that the similarity checks by the filter consume a non-negligible amount of computing power. In order to avoid that the GEN becomes a performance bottleneck resulting in congestion if the number of messages exceeds a certain rate, the computational effort must be kept small.

IV. Caching and Filtering of Similar Target Regions

The design goals of the caching and filtering mechanisms are the efficient utilization of the available computing resources and the determination of a previous target region being highly similar to a new one.

The computational effort mainly depends on the number and the complexity of the similarity checks that need to be performed for each target region. A common approach for the reduction of the computational effort for the similarity checks is the utilization of polygonal approximations of the target regions, but these approximations introduce deviations from the originally requested region and the ones being used for the computations.

In the remainder of this paper, we present an analysis of the deployment of filters using various polygonal approximations of target regions and the effective utilization of different combinations of filter mechanism with caches.

A. Investigated Filter and Cache Combinations

In the sequel, we investigate three different combinations of filters and caches and their suitability for the derivation of similar target areas. The first variant is the combination of a cache with a filter performing exact computations of the similarity. If a new target area arrives, a superset of candidates is fetched from the cache. Then, the similarity of each superset member to the new target region is computed. Finally, the region with the best result will be returned.

In the second variant, we analyze a cache which is combined with a two stage filter architecture. The candidates being returned by a lookup in the cache are first checked with a filter using approximations of the region shapes in order to reduce the size of the candidate set. Afterwards, exact similarity checks are performed in the second stage with the members of the reduced set and the new target region, and the region with the best result is returned. The filtering in the first stage is intended for the detection of regions in the superset that would only lead to insufficient results in the second stage. It can be performed in two different ways. In the first one, decisions are based on *intersection* checks. The filter checks if the approximation of the new target region intersects the approximations of the regions being returned by the cache lookup, and the exact similarity checks in the second stage are only performed with the candidates where this check has been positive. The second method for the reduction of the candidate sets is the computation of the *similarity* of the approximated regions. Only those candidates whose approximations have a certain similarity to the approximation of the new target region are investigated in the second stage.

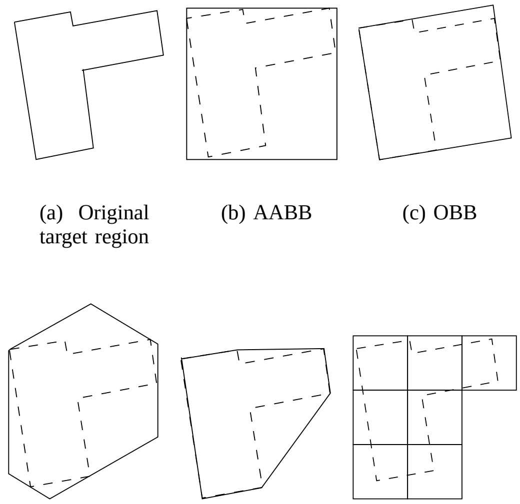

(a) Original target region (b) AABB (c) OBB

(d) 3-DOP (e) Convex Hull (f) Raster

Fig. 2. Different target region approximations

The third variant utilizes a cache in combination with a filter performing similarity checks solely based on approximated shapes of the target regions. The member of the superset whose approximation possesses the highest similarity to the approximation of the new target region will be returned.

B. Deployed Approximation Techniques

The running time of algorithms for the computation of intersection regions being required for the similarity checks depends on the number of edges of the new target region n_t and the cached region n_c. The intersection region can be computed in time

$$O((n_t + n_c)\log(n_t + n_c) + k), \tag{4}$$

with k being the number of intersection points of the edges of the two regions [8]. One objective of an approximation of a region is therefore the reduction of the number of edges.

In this section, we give a brief overview on the approximations that have been investigated for the filters mechanisms, together with the computational upper bounds for the construction of the polygons. In Figure 2, approximations for an exemplary shape are being displayed.

An approximation should enable a good compliance of the areas of the original target region and the polygon as well as computational inexpensive construction.

Axis Aligned Bounding Box (AABB)

The axis aligned bounding box of a region is given by the smallest rectangle that contains all points of the region. The rectangle is oriented to two axes of a fixed coordinate system. It can be described by four parameters, the minimum and maximum spatial extend of the region according to each of the two axes.

The minimum bounding rectangle of a region can be computed in linear time $O(n)$, whereas n denotes the number of the edges of the original region.

Oriented Bounding Box (OBB)

The oriented bounding box of a region is the smallest rectangle in terms of its area containing all points of the region without requiring the rectangle to be aligned to the axes. Thus, it can be described by five parameters, the four known from the AABB plus an angle describing the rotation of the rectangle. An OBB can be computed in time $O(n \log(n))$. The advantage to the AABB is the better compliance with the area of the region.

Discrete Oriented Polytop with k-Faces (k-DOP)

A discrete oriented polytop can be considered as an enhanced AABB. In 2-dimensional space, a k-DOP is the smallest polygon that contains all points of a target region and is similarly constructed as the AABB. Instead of using 2 coordinate axes, it is computed by determining the minimum and maximum spatial extend of the region regarding k fixed axes in the plane. For $k = 2$, a k-DOP is identical to an AABB.

Each k-DOP can be described by $2k$ parameters and can be computed in linear time $O(n)$.

Convex Hull

A convex hull of a region is defined as the smallest convex polygon that contains all points of the originally selected target region. Several algorithms for the efficient computation of convex hulls have been proposed. The recent algorithms compute the convex hulls in time $O(n \log(n))$. A convex hull is described by its edges, thus, the number of required parameters is not fixed.

Raster Approximation

Another mechanism for the approximation of a target region is the utilization of a fixed raster in the geographical territory. By the raster, the territory is partitioned into tiles. Each tile being intersected by the region that should be approximated is marked, leading to a binary pattern for the region. The advantage of this mechanism is the possibility to perform the computations for the similarity checks with binary operations which can easily be integrated into hardware.

C. Caching of Recent Target Regions

The regions that have recently been the targets of messages are stored in a spatial database. For efficiency reasons, the database uses the AABBs of the regions as their keys being organized in a tree structure. A query in the database is performed based on the AABB of the new target region and returns a candidate set consisting of all cached regions whose AABBs intersect the AABB of the new target.

D. Evaluation of the Mechanisms

The average similarity that can be achieved for a new target region to the best matching cached one is relevant for the deployment of a GEN. The probability to find a region with high similarity depends on the cache size. Expedient similarity check results are to be achieved for appropriate cache sizes.

TABLE I

PARAMETERS FOR THE MODEL OF THE TARGET REGIONS

Parameter	Value
v_{min}	4
v_{max}	8
d_{min}	200m
d_{max}	500m

Regarding the cache and filter combinations, efficient similarity checks have to be realized for a proper utilization of the available computing power and storage. The introduction of the approximating filters in the two stage approach should significantly reduce the size of the candidate set being returned by the cache in order to keep the total processing effort of the two stages low.

The filter decisions based on the approximated regions in the third filter and cache combination have to be as effective as possible. Due to the simplifications of the region shapes and the resulting different areas of the approximated and original regions, a filter using approximations may return a higher or lower similarity compared to the one which would be returned by a filter computing the similarity for the original regions. Therefore, the best match being returned by an approximating filter may differ from the best match that would have been returned for the original regions.

V. Numerical Results and Discussion

In this section, we present some numerical results that have been derived for the mechanisms.

For the evaluation of the approach, the filter mechanisms have been implemented and integrated in the discrete event simulation tool OMNeT++ [9]. We investigated a section of a geographical territory with a size of $5\,000\text{m} \times 5\,000\text{m}$. The model for the regions that are selected as targets in the territory has been parameterized according to the values in table I. As mentioned in section II, we investigated two different spatial distributions of the centers of the regions, the uniform and Beta distribution with parameters $\mu_1 = 2$ and $\mu_2 = 2$.

For each of the distributions, we performed 15 simulation runs. In each run, 300 different regions have been selected. The intervals being displayed in the graphs denote the 95% confidence intervals.

First, we investigated the effect of the different spatial distributions of the target regions on the similarity of a new target region to a cached one. Obviously, the probability to find a region being similar to a newly selected one is heavily influenced by one parameter, the size of the cache for the previous targets.

In figure 3, the average similarity of the best match of a stored target region to a new request is depicted. The results have been obtained with the first variant of the filter mechanisms, namely the combination of a cache with the filter

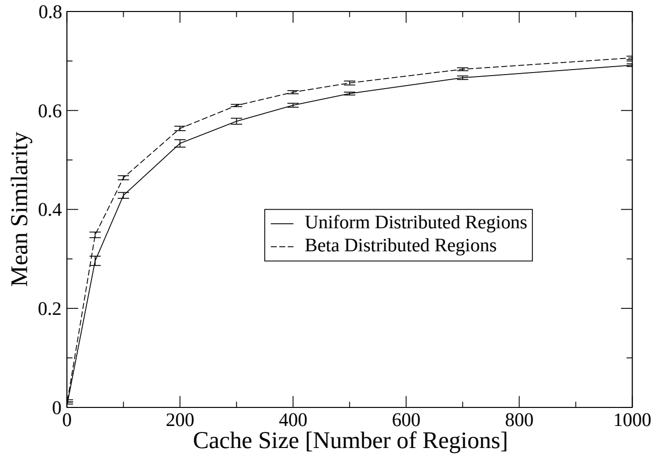

Fig. 3. Mean similarity of the best matching region being cached depending on the amount of regions in the cache

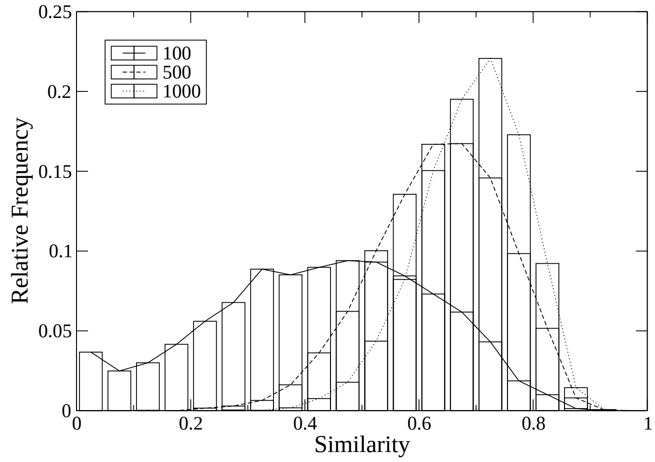

(a) Uniform distributed target regions

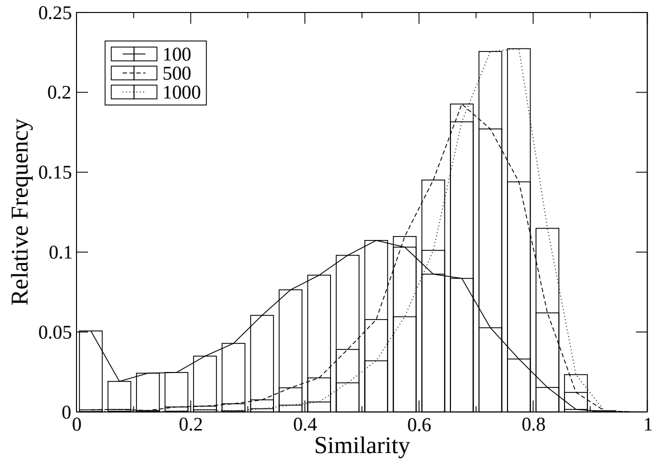

(b) Beta distributed target regions

Fig. 4. The histograms show the distributions of the best similarity check results for three different cache sizes of 100, 500 and $1\,000$ regions

performing the exact similarity checks when a message should be sent to a new geographical region. As the graph shows, the difference between the average similarity is only marginal for the two distributions of the centers. The influence of the cache size on the average similarity decreases with increasing cache size. Before each investigation, the cache was filled with the maximum number of target regions, which were generated according to the same model being used for the target regions.

The frequencies of the maximum similarities of new target regions to cached ones are presented in the histograms 4(a) for the uniform and figure 4(b) for the Beta distribution. In the histograms, the results for three different cache sizes are presented, i.e. caches for 100, 500 and $1\,000$ target regions. If a Beta distribution is assumed in order to reflect the spatial concentration of target regions, the values in the histograms are shifted towards the direction of a higher similarity compared to the uniform distributed case. Considering the distributions, a cache size of $1\,000$ regions leads to a similarity of 70% or more for over 59% of the new target regions and to a similarity of 80% or more for about 14% of the Beta distributed regions. The bucket/cell size being chosen in the histograms is 0.05, and the amount of sample regions is $4\,500$ for each curve.

In the second filter combination, the approximating filters are intended for a reduction of the candidate sets being returned by a cache lookup. This can be achieved in two different ways, by intersection or similarity checks. The efficiency of simple intersection checks is presented in figure 5 for different approximation techniques. It leads to an average reduction of about 28% of the candidate set size, if convex hull approximations are utilized. The same result is achieved by performing the checks with the original polygons. Discrete oriented polytops lead to a reduction of about 17%, raster approximation to a reduction of about 13%, and oriented bounding boxes achieve only about 7%. The parameters for the approximations are $k = 3$ for the k-DOP filter and the raster filter uses a lattice spacing of about 110 meters in x and y direction. In all cases, the ratio of additionally filtered regions is nearly independent from the cache size (except for cache sizes smaller than 100 regions). In consideration of the fact that only small reductions of the candidate set size can be achieved by this mechanism, the total computational effort will most likely be even higher than that of the first filter combination.

The second method for the reduction of the candidate sets is the computation of the *similarity* of the approximated regions. The best match being returned by the approximating filter does not necessarily coincide with that of the original regions. Therefore, we investigated the ratio of the best candidates being returned by the approximating filter that has to be used for further testing. The outcome is presented in figure 6. The graph shows the average proportion of the reduced candidate set from the approximating filter containing the region that will lead to the best result when the exact similarity checks

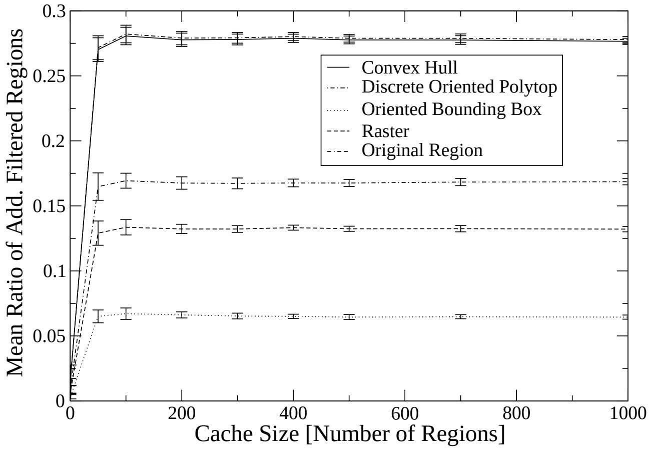

Fig. 5. By using approximating filters, the amount of regions being returned by the cache can additionally be reduced with boolean intersection checks. The graph shows the ratio of additional reduced regions depending on the cache size for Beta distributed regions

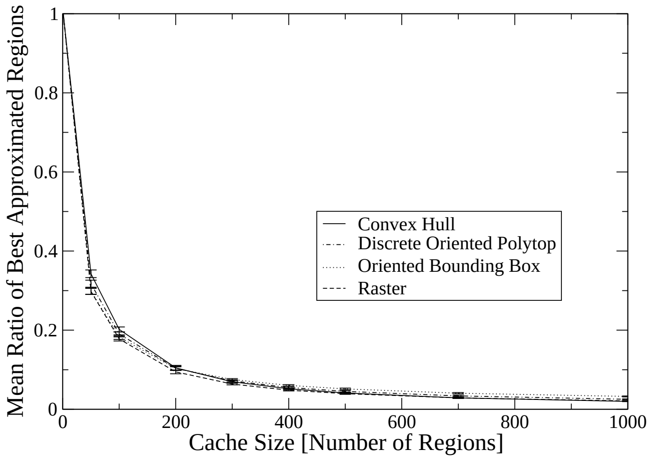

Fig. 6. Mean best ratio of the approximating filter results containing the region that will lead to the best match when the similarity checks are performed with the original regions. The target regions are Beta distributed

are performed. For instance, for a cache size of $1\,000$ regions, the original region that will lead to the best match is on average within the best 5% of the reduced candidates being returned by the approximating filter. The results of the different approximation techniques are hardly distinguishable from each other. Due to possible the large reduction of the candidate set, this filter might be more efficient than the first cache and filter combination.

For the third filter combination, the characteristics of the approximating filters regarding the determination of best matches that differ from those that would have been returned by a filter using the original regions have been analyzed. For the investigation, we took the best matching region being returned by an approximating filter for a new target region and computed the similarity that would have been determined for the two original regions. Afterwards, the difference between the similarity of the best matching original polygons has been computed.

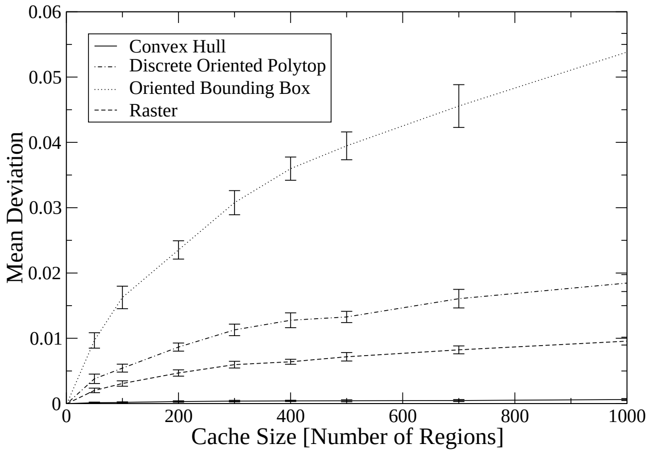

Fig. 7. The region being returned as best match from an approximating filter may not coincide with the best match for the original regions. The figure shows the mean deviation of the similarity of returned regions due to approximations from the original regions.

The results are displayed in figure 7. For all approximation techniques, the deviation between approximated and original regions increases with the size of the cache. The smallest difference between the best match being determined based on approximated regions and the original regions is achieved by the filter using convex hulls of the regions. The average deviation stays below 0.1% for all tested cache sizes. Quite acceptable results are also obtained by the raster and DOP filter with an average deviation below 1% and 2%, respectively. The performances of the OBB filter is about a factor of 5 worser than that of the raster one. If the effort for the computation of the approximating shapes of the regions is being neglected, the third filter combination is the fastest of the presented ones.

VI. Related Work

The utilization of geocast for the provisioning of location based services has already been proposed in the context of various application. It became a fixed part of environments for LBS, e.g. in the NEXUS platform [10]. Geocast is also utilized for the location based service discovery in networks. For instance, in [11], control messages are sent based on geocast both for the discovery and advertisement of services. The deployment of geocast in cellular networks has been addressed in the paper [12]. The authors deal with the occurring problems of geocast zone formulation, routing, group construction and maintenance in the networks.

The caching of geographical forwarding decisions in intermediate systems has been proposed in [5]. A mechanism for the avoidance of geometrical operations in intermediate systems is the GPS-Multicast approach presented in [13]. Nodes which are responsible for parts of the geographical territory being covered by a network are grouped to so called atoms and partitions (i.e. several atoms), which are mapped to multicast addresses. If a message should be sent to a region, the smallest partition or atom is determined which contains the region. Then, the message is sent to the respective multicast address. Due to the partitioning, the approach tends to the erroneous sending of messages to nodes which have to filter and discard them afterwards.

Mechanisms for the simplification of geometrical computations for geographical routing have been analyzed in [4]. By utilizing polygonal approximations of geographical regions, the computational expense of forwarding relevant intersection checks of target regions and service areas (i.e. the geographical regions being covered by nodes in the forwarding direction) is reduced. The work focuses only on the required intersection checks without considering similarity checks of target regions.

Considering the area of computer science, another application of geometric object approximations can be found in geographical information systems. These systems perform intensive operations on spatial databases. Queries in these databases may result in large response sets, whereas the determination of the sets involves computational intensive algorithms. In this context, the deployment of a two stage filtering for the reduction of the response sets based on intersection checks has been introduced in [14].

The geometrical algorithms that are mentioned in this paper belong to the state of the art in the area of computational geometry. An algorithm for the intersection of polygons is presented in [8], some of the algorithms for the construction of approximations can be found e.g. in [15].

VII. Conclusion and Outlook

In this paper, we presented a novel approach for the provisioning of location based push services which relies on geographical routing mechanisms. By the introduction of a central computational and storage enhanced node (GEN) close to the gateway of a carrier access network, a reduction of the required processing effort for the routing of messages to geographical target regions can be achieved.

The basic element for this approach is a mechanisms in the enhanced node being able to detect messages which are sent to identical or similar target regions. If the similarity of a new target region to a previous one is acceptable, cached geometrical forwarding decisions in the intermediate nodes are reused.

The focus of this paper was on the investigation of the component for the detection of similar regions. By utilizing a model for target regions being based on that from Navas and Imielinski, we analyzed the distributions of the similarities of the regions. We have shown that a big ratio of target regions possesses a large similarity to cached ones.

Additionally, different mechanisms for the simplification of the similarity computations based on approximations of target regions have been analyzed. We discovered that by the deployment of a two stage filter approach, significant reductions of candidate sets and thus of the required processing power can only be achieved if similarity instead of intersection checks are utilized. Further results have shown that filters solely using approximations of target regions for the determination of the similarity also achieve good results, i.e. the best matches being returned by them differ only little from those that would have been returned after an exact computation. For instance, they differ less than 1% for filters using raster approximation.

Regarding the next steps of our ongoing work, we will investigate the effect of the deployment of this approach on the overall processing effort in a carrier access network. Different mechanisms for the reuse of forwarding decisions in the network will be analyzed, for example the deployment of dynamic multicast groups.

Acknowledgment

The authors would like to thank Oliver Belaifa and Zakaria Kasmi who implemented and tested the presented filter mechanisms and integrated them in the simulation tool. Additionally, they implemented a very helpful tool for the visualization and validation of the similarity checks.

References

[1] J. Schiller and A. Voisard, *Location-Based Services*, ser. Morgan Kaufmann Series in Data Management Systems. Morgan Kaufmann, 2004.

[2] S. D. Hermann, G. Schäfer, A. Wolisz, and M. Lipka, "Concept for the hierarchical and distributed processing of area based triggers," in *Fourth Annual IEEE International Conference on Pervasive Computing and Communications (PerCom)*, 2006.

[3] C. Maihöfer, "A survey of geocast routing protocols," *IEEE Communications Surveys & Tutorials*, vol. 6, no. 2, pp. 32–42, 2004, http://www.comsoc.org/livepubs/surveys/public/2004/apr/maihofer.html.

[4] J. C. Navas and T. Imielinski, "On reducing the computational cost of geographic routing," Siemens TTB center, Computer Science Dep. Rutgers University, Tech. Rep. DCS-TR-408, Jan. 2000.

[5] T. Imielinski and J. C. Navas, "Gps-based geographic addressing, routing, and resource discovery," *Communications of the ACM*, vol. 42, no. 4, pp. 86–92, 1999.

[6] J. C. Navas and T. Imielinski, "Geocast - geographic addressing and routing," in *Proc. of the Third ACM/IEEE Intl. Conf. on Mobile Computing and Networking (MobiCom'97)*, Budapest, Hungary, Sept. 1997.

[7] R. R. Larson and P. Frontiera, "Spatial ranking methods for geographic information retrieval (gir) in digital libraries," in *Research and Advanced Technology for Digital Libraries: 8th European Conference, ECDL 2004*, Sep 2004.

[8] I. J. Balaban, "An optimal algorithm for finding segments intersections," in *11th annual symposium on computational geometry (SCG '95)*. New York, NY, USA: ACM Press, 1995, pp. 211–219.

[9] A. Varga, "OMNeT++, discrete event simulation system, http://www.omnetpp.org."

[10] D. Fritsch and S. Volz, "NEXUS - the mobile GIS-Environment," in *First Joint Workshop on Mobile Future and Symposium on Trends in Communications*, Bratislava, Slovakia, October 2003, pp. 1–5.

[11] N. Klimin, W. Enkelmann, H. Karl, and A. Wolisz, "A hybrid approach for location-based service discovery in vehicular ad hoc networks," in *Proc. of 1st Intl. Workshop on Intelligent Transportation (WIT)*, Hamburg, Germany, Mar. 2004.

[12] B. An, S. Tekinay, S. Papavassiliou, and A. Akansu, "A cellular architecture for supporting geocast services," in *Vehicular Technology Conference, 2000. IEEE VTS-Fall VTC 2000. 52nd*, vol. 3, Boston, MA, USA, Sep 2000, pp. 1452–1459.

[13] T. Imielinski and J. C. Navas, "Gps-based addressing and routing," Rutgers University Computer Science, Tech. Rep. LCSR-TR-262, March 1996, updated August.

[14] T. Brinkhoff, H.-P. Kriegel, and R. Schneider, "Comparison of approximations of complex objects used for approximation-based query processing in spatial database systems," in *Ninth International Conference on Data Engineering*. Washington, DC, USA: IEEE Computer Society, 1993, pp. 40–49.

[15] M. de Berg, M. van Kreveld, M. Overmars, and O. Schwarzkopf, *Computational Geometry. Algorithms and Applications*, 2nd ed. Springer, 2000.