

ABSTRACT

There has been a tremendous growth in the field of Mobile Geographic Information Systems (Mobile GIS) and the trend is continuing. Mobile GIS is fundamentally changing the way geography is utilized and data is handled in mobile environments. It is bridging the gap between working in the office and working on the move. Users of Mobile GIS can now retrieve, manipulate and update enterprise data right from their PDA's anytime, from anywhere in the world. The resulting benefit is consistently improved productivity. Examples include real time access and tracking of shipments, emergency services, car navigation services, real time stock quotes, field services. With the proliferation of Mobile GIS applications, the emphasis is now laid on spatial data; data that have a characteristic location attribute. Issues of storage, retrieval and efficient management of spatial data in these applications are gaining prominence.

User queries can be answered faster if the requested information is in the cache rather than sending the query request to the original source, which may be remotely located. Caching is especially important in today's mobile environments where limited bandwidth and errant wireless connections make disconnected mode of operation both pertinent and indispensable. In this new application domain, traditional caching schemes like page caching and tuple caching do not scale well and often times fail badly due their inherent assumption of spatial locality of reference, inability to support content based reasoning or sometimes simply due to the dynamics of spatial data and mobile environments.

We need an altogether different approach to manage and efficiently utilize cache while dealing with spatial data in mobile GIS applications. Spatial queries in these applications tend to exhibit semantic locality due to the inherent location attribute associated with spatial data. We exploit this fact for caching purpose and this is where semantic caching comes into picture. We use semantic cache descriptors to determine and manipulate the contents of the cache. In this thesis we propose an infrastructure that supports semantic caching of data in mobile environments. We build a prototype system as a "proof of concept" that captures the essential elements of semantic caching and demonstrates its working. Our system uses the existing relational database framework to efficiently manage semantic cache. We devise and implement a new cache replacement algorithm that takes advantage of schema knowledge in determining the replacement victims. We formulate the problem as an integer linear program and then solve a relaxation to obtain fractional optimal solution. We also build a solver that, for restricted cases, determines if a query can be answered from local cache. Finally we propose a variety of techniques for semantic query processing and cache management.