Geodetic Distance Queries on R-Trees for Indexing Geographic Data

Erich Schubert, Arthur Zimek, and Hans-Peter Kriegel

Ludwig-Maximilians-Universitt Mnchen
Oettingenstr. 67, 80538 Mnchen, Germany
{schube,zimek,kriegel}@dbs.ifi.lmu.de
http://www.dbs.ifi.lmu.de

Abstract. Geographic data have become abundantly available in the recent years due to the widespread deployment of GPS devices for example in mobile phones. At the same time, the data covered are no longer restricted to the local area of a single application, but often span the whole world. However, we do still use very rough approximations when indexing these data, which are usually stored and indexed using an equirectangular projection. When distances are measured using Euclidean distance in this projection, a non-neglibile error may occur. Databases are lacking good support for accelerated nearest neighbor queries and range queries in such datasets for the much more appropriate geodetic (great-circle) distance. In this article, we will show two approaches how a widely known spatial index structure – the R-tree – can be easily used for nearest neighbor queries with the geodetic distance, with no changes to the actual index structure. This allows existing database indexes immediately to be used with low distortion and highly efficient nearest neighbor queries and radius queries as well as window queries.

1 Introduction

Nowadays, we are much more used to seeing maps than using a globe. But once we look at maps of the world, all map projections have some error: they cannot preserve the three components of geographical information, distance, area, and angles, equally well. The projections that we are most used to are the Mercator projection and the even simpler equi-rectangular projection. Google maps, for example, is based upon the Mercator projection. Figure 1 shows three different projections used for maps. Figure 1a is a variation of the usual Mercator projection, which has a huge influence on our understanding of the world. Figure 1b is an alternate projection developed 1923 by John Paul Goode, which is an equal-area projection. The difference in the ability to preserve area can in particular be seen for the Antarctica, but also Greenland, Canada, Alaska and Russia show massive area distortions in the Mercator projection. The Mercator projection is obtained by wrapping a cylinder around the earth and projecting the earth surface onto this cylinder. This yields a good map where the cylinder touches or intersects the earth, but along the axis of the cylinder the distortion is infinite – the north and south poles do not project to the cylinder at

M.A. Nascimento et al. (Eds.): SSTD 2013, LNCS 8098, pp. 146–164, 2013.

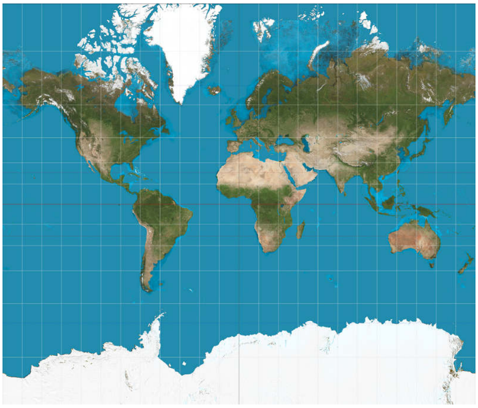

(a) Mercator projection

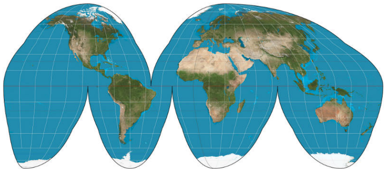

(b) Goode homolosine projection

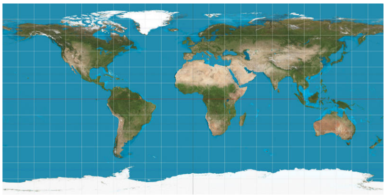

(c) Equirectangular projection

Fig. 1. Three different projections of the earth surface. Images from Wikimedia Commons, by user "Strebe"[1]

all. Therefore, Mercator maps are commonly truncated at the poles. Figure 1a for example is truncated to a latitude of $\pm 82°$. While the Mercator projection does not preserve area and distances, it preserves the shapes quite well – but even more importantly, it preserves angles, which makes it useful for navigation and this probably is the main reason for the popularity of this projection. The equirectangular projection (Figure 1c) is probably the simplest projection, where latitude and longitude translate directly into y and x. In contrast to Mercator, it does not preserve angles. However, being able to trivially translate latitude and longitude to pixels on this map projection makes it a popular choice in GIS raster applications, and this is the default projection for placing custom texture overlays on the earth surface, e.g. in Google Earth and NASA WorldWind.

There are hundreds of geodetic reference datums, some of which are historic, but many are still in use for good reasons, e.g. for land survey. We tend to assume that any geographic position is well-defined, but for example due to tectonic plate shift we have to accept the fact that even the largest mountains move with respect to each other over time. The most popular geographic coordinate system is that of geographic latitude and longitude, with respect to the World Geodetic System 1984 (WGS84) reference ellipsoid. Geographic position is then measured with three components, known as latitude, longitude, and elevation. Elevation is commonly given with respect to the reference ellipsoid's surface, so that values of 0 are approximately at sea level. In the following, we will use λ to denote latitude and ϕ to denote longitude. We will not be using elevation, since it often is not available for a data set. Furthermore, in order to fully specify position, one would also need to take the time of the measurement into account. For all these reasons, we have to live with some error in geo positioning and, thus, distance

[1] For detailed copyright information on these images (CC-BY-SA 3.0), see https://commons.wikimedia.org/wiki/Category:Images_of_map_projections