# Yelp: Case Study

*Systems Analysis & Design*

# Learning Objectives

By the end of this session, you will have acquired the following information:

- Yelp System Design
- Geo Index

# Proximity Service

- A proximity service is utilized to identify nearby locations, including restaurants, hotels, theaters, museums, and more.
- This service is a fundamental component that enables features such as locating the top-rated restaurants in the vicinity on platforms like Yelp.

Yelp › Restaurants

# Top 10 Best Restaurants Near San Francisco, California

Sort: **Recommended**

**All "Restaurants" results in San Francisco, California**

## Filters

$ | $$ | $$$ | $$$$

### Suggested

- [ ] Open Now 4:00 PM
- [ ] Offers Delivery
- [ ] Reservations
- [ ] Offers Takeout
- [ ] Good for Dinner
- [ ] Hot and New

### Features

- [ ] Good for Kids
- [ ] Has TV
- [ ] Outdoor Seating
- [ ] Gender-neutral restrooms

See all

### Neighborhoods

- [ ] Alamo Square
- [ ] Anza Vista
- [ ] Ashbury Heights
- [ ] Balboa Terrace

See all

### Distance

- ( ) Bird's-eye View
- (•) Driving (5 mi.)
- ( ) Biking (2 mi.)
- ( ) Walking (1 mi.)
- ( ) Within 4 blocks

### 1. Bottega

4.3 (965 reviews)

Italian · Pasta Shops · Pizza · $$ · Mission

**Closed** until 8:30 AM

Outdoor seating · Locally owned & operated

Waitlist opens at 11:00 am

"Went to Bottega for a date night and let me tell you it was amazing. Had a reservation so did not have to wait long, once seated the waiter attended to us..." more

✓ Delivery ✓ Takeout

Start Order

### 2. Memento SF

4.8 (29 reviews)

New American · Desserts · Cocktail Bars · Noe Valley

**Closed** until 2:30 PM

"My bestfriend and I went to this restaurant before a show and it was such a great choice! Our server Christine truly went above and beyond, making sure we..." more

✗ Outdoor seating ✗ Delivery

Find a Table

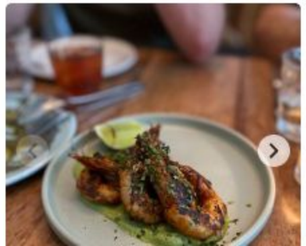

### 3. Pearl

4.4 (472 reviews)

Italian · Seafood · Mediterranean · $$$ · Outer Richmond

**Closed** until 2:30 PM

"This place has very good alfredo here and the customers they gave me was awesome, I will def give them a try if your out in the bay area. My wife thought I was..." more

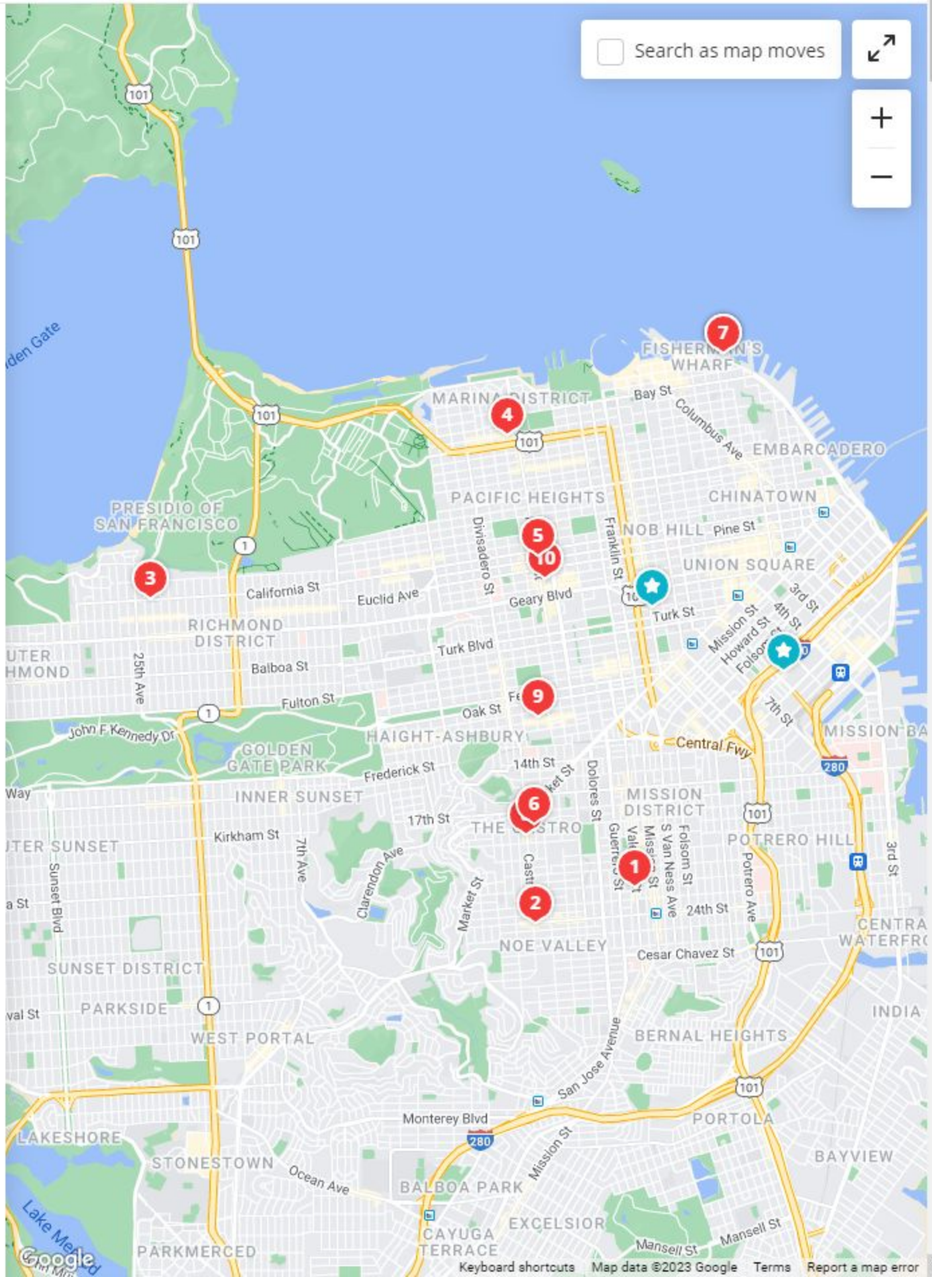

# Functional Requirements

- Retrieve all businesses based on a user’s location, defined by a pair of latitude and longitude coordinates, and a specified radius.
- Business owners have the ability to add, delete, or update their business information. However, these changes do not need to be reflected in real-time.
- Customers have the capability to view comprehensive information about a business.

# Non-Functional Requirements

- Low Latency: Users should have the ability to view nearby businesses swiftly.
- Data Privacy: Location information is sensitive data. When designing a location-based service (LBS), user privacy should always be a primary consideration.
- High Availability and Scalability: We should ensure that our system is capable of handling traffic spikes during peak hours in densely populated areas.

# API Design

- `GET /v1/search/nearby`

| Field | Description | Type |
| --- | --- | --- |
| latitude | Latitude of a given location | float |
| longitude | Longitude of a given location | float |
| radius | Optional. Default is 5km | int |

# API Design

| API | Detail |
| --- | --- |
| `GET /v1/businesses/:id` | Return detailed information about a business |
| `POST /v1/businesses` | Add a business |
| `PUT /v1/businesses/:id` | Update details of a business |
| `DELETE /v1/businesses/:id` | Delete a business |

# Read/Write Ratio

- The read volume is high due to the frequent use of the following two features:
  - Searching for nearby businesses
  - Viewing detailed information about a business
- The write volume is low because operations such as adding, removing, and editing business information are infrequent.

# Business Table

- A relational database, such as MySQL, could be a suitable choice.
- A geo index table is used for the efficient processing of spatial operations.

| Business |
| --- |
| business_id (primary key) |
| address |
| city |
| state |
| country |
| latitude |
| longitude |

# High-Level Design

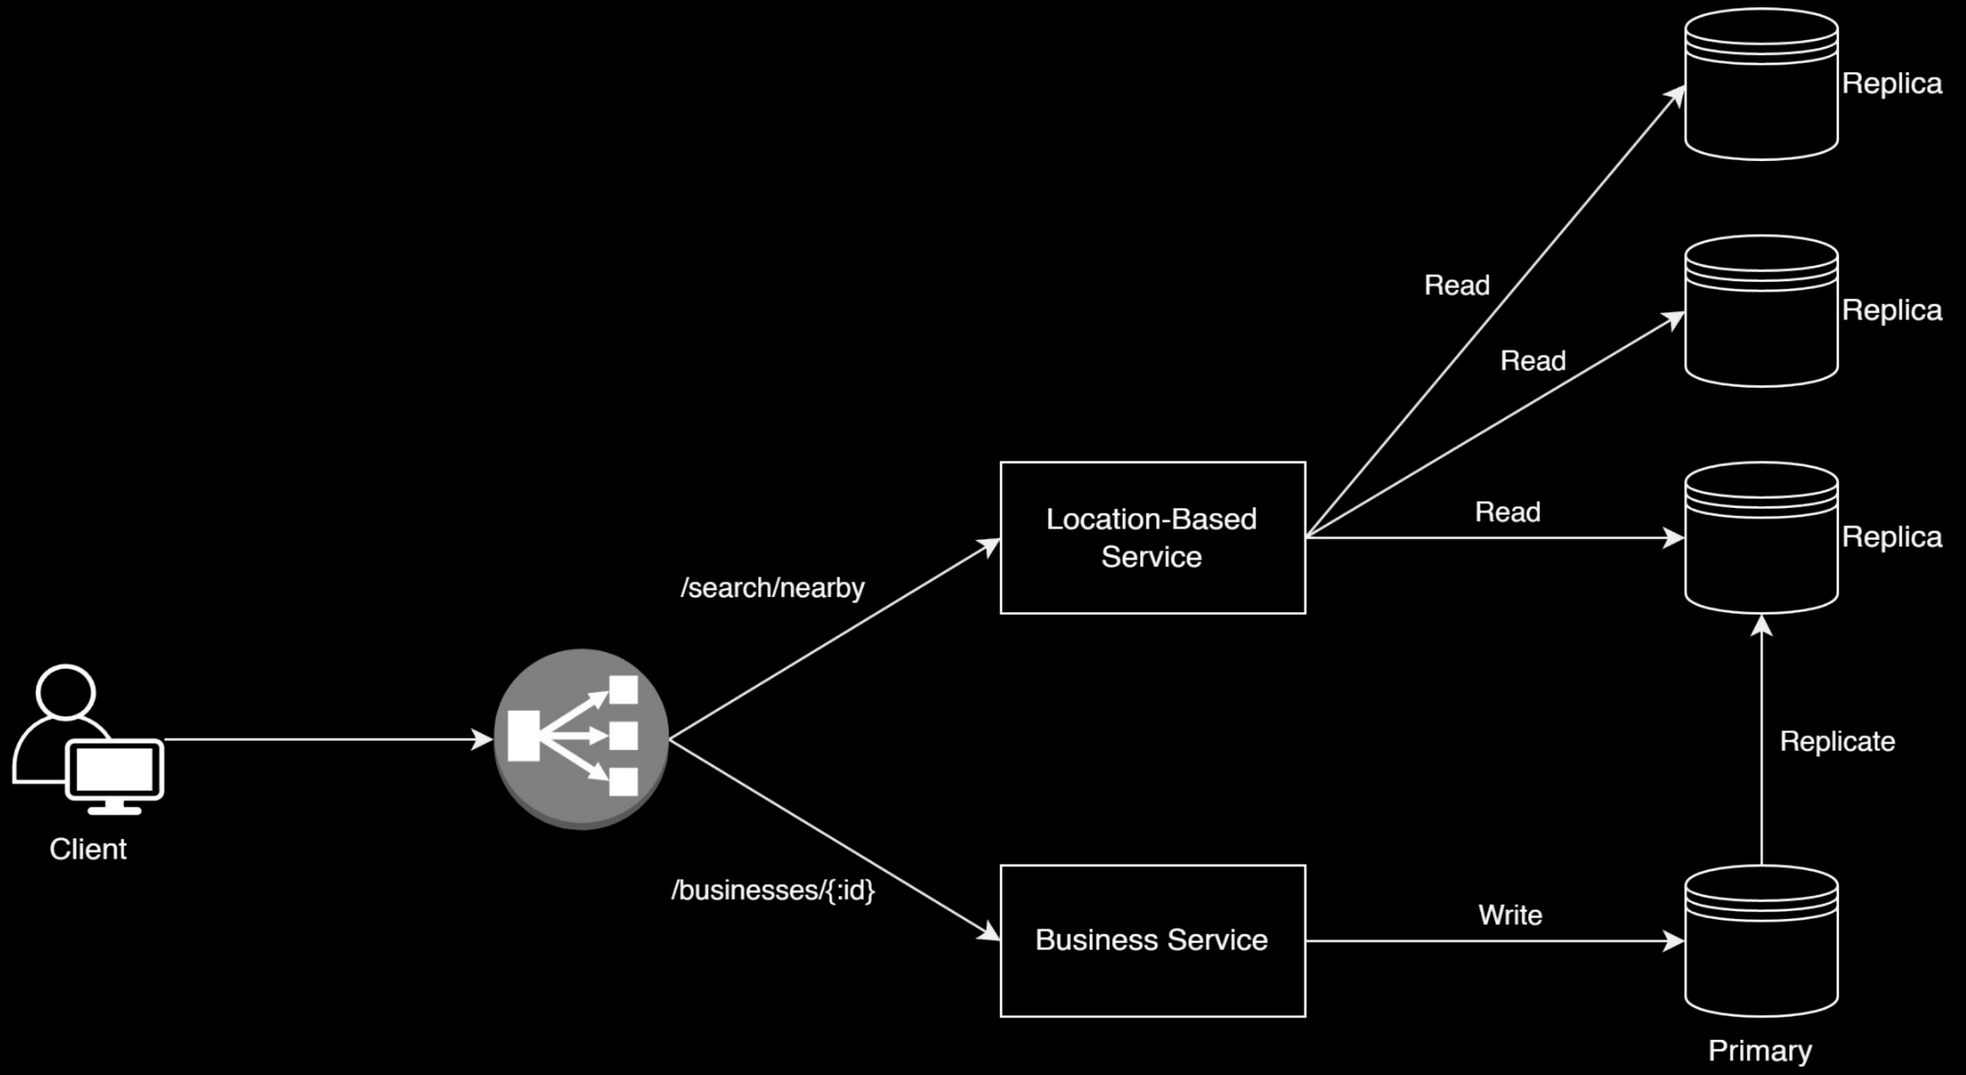

# Fetch Nearby Businesses

1. Two-Dimensional Search
2. Evenly Divided Grid
3. Geohash
4. Quadtree
5. Google S2 (Hilbert Curve)

# Two-Dimensional Search

The naive way to find nearby businesses is to draw a circle with a predefined radius and find all the businesses within that circle.

```
SELECT business_id, latitude, longitude
FROM business
WHERE
    (latitude BETWEEN {:my_lat} - radius AND {:my_lat} + radius)
AND
    (longitude BETWEEN {:my_long} - radius AND {:my_long} + radius)
```

This query is not efficient because we need to perform an intersection operation.

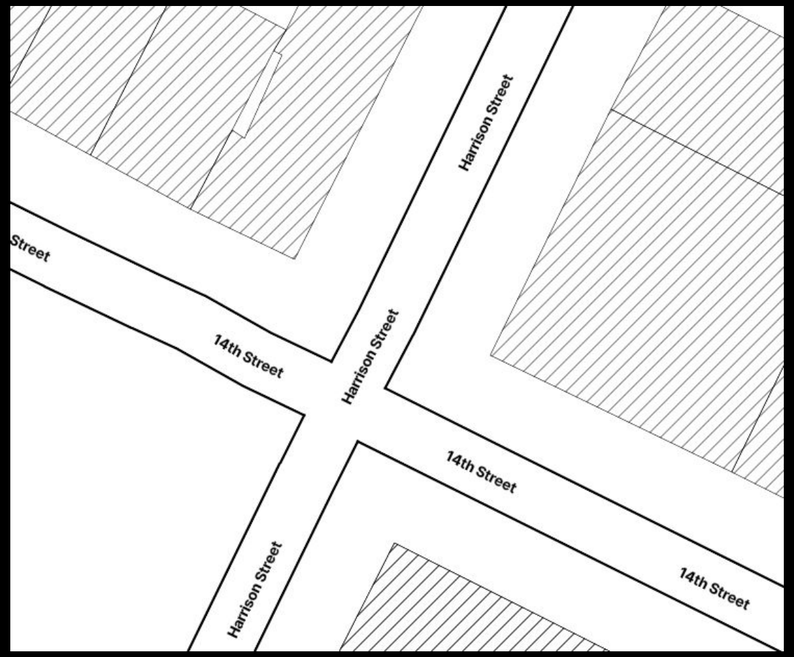

# Evenly Divided Grid

- One simple approach is to evenly divide the world into small grids.
- A single grid could contain multiple businesses, and each business on the map belongs to one grid.
- The distribution of businesses is not even.
- Ideally, we would want to use more granular grids for dense areas and larger grids for sparse areas.

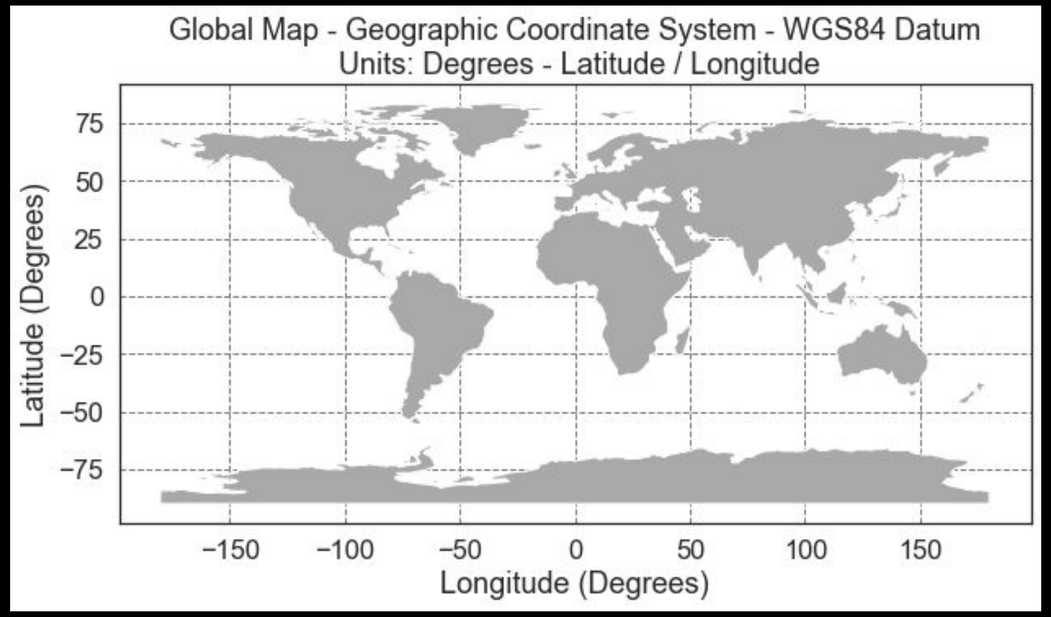

# Geohash

- Geohash is superior to the evenly divided grid option.
- It operates by converting the two-dimensional longitude and latitude data into a one-dimensional string of letters and digits.

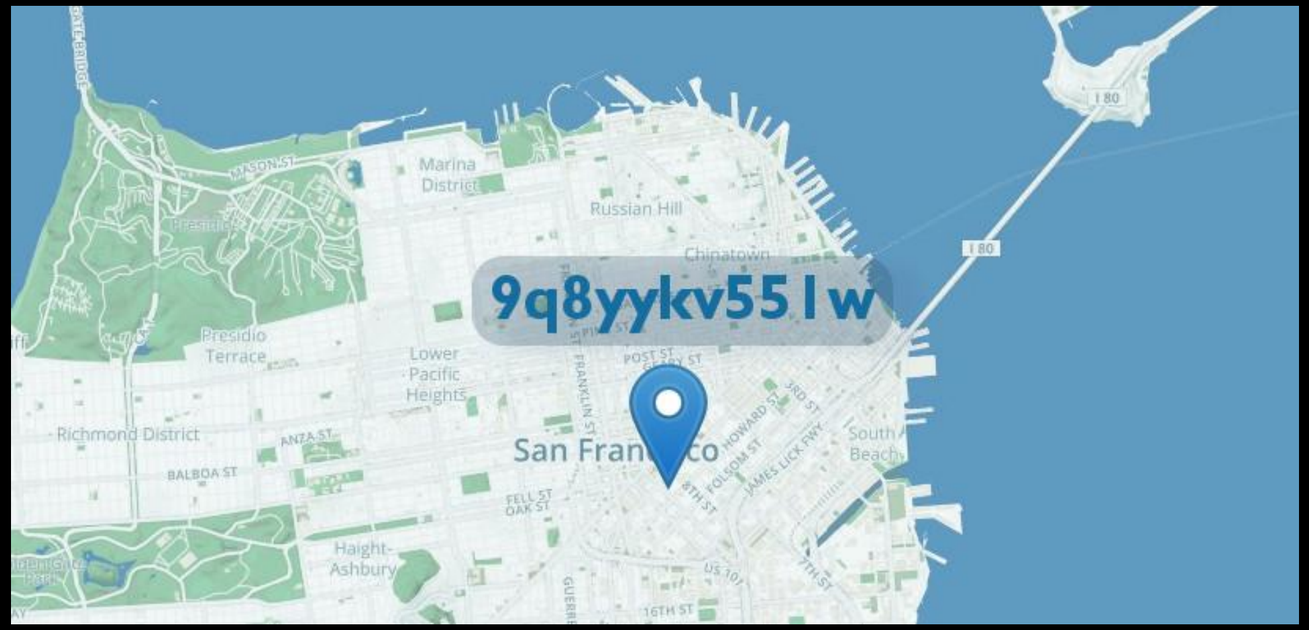

# Geohash

- The world is recursively divided into smaller and smaller grids with each additional bit.
- The longitude range [-180,0] is represented by 0, and the longitude range [0, 180] is represented by 1.

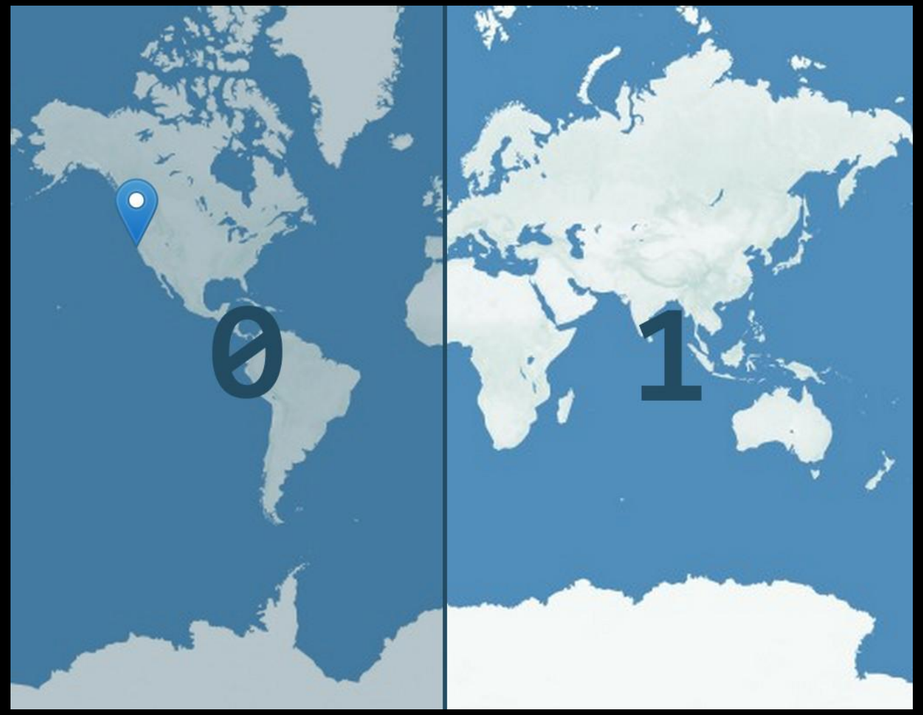

# Geohash

- The world is recursively divided into smaller and smaller grids with each additional bit.
- The latitude range [-90,0] is represented by 0, and the longitude range [0, 90] is represented by 1.
- Repeat this subdivision, alternating between longitude and latitude, until the remaining area is within the desired precision.
- Geohash usually uses base32 representation.

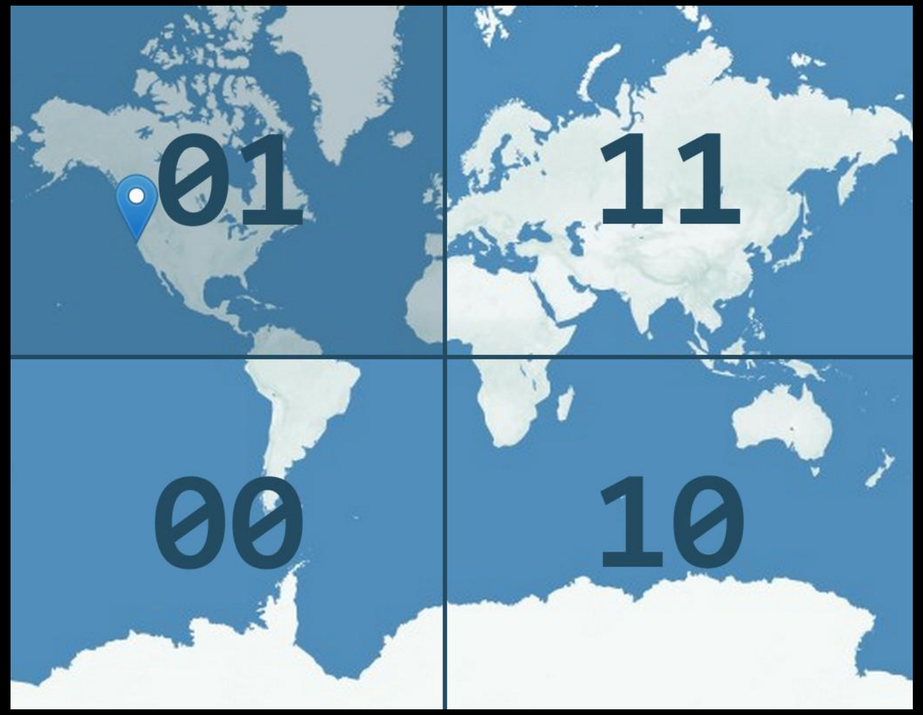

# Geohash

- For example, the longitude and latitude coordinates (37.77564, -122.41365) result in the binary sequence “0100110110010001111011110” and produce the geohash “9q8yy”.

| Binary | 01001 | 10110 | 01000 | 11110 | 11110 |
|---|---|---|---|---|---|
| Decimal | 9 | 22 | 8 | 30 | 30 |
| Base 32 | 9 | q | 8 | y | y |

# Geohash

- We can decode a geohash into its corresponding latitude and longitude.

| Base 32 | 9 | q | 8 | y | y |
|---|---|---|---|---|---|
| Decimal | 9 | 22 | 8 | 30 | 30 |
| Binary | 01001 | 10110 | 01000 | 11110 | 11110 |
| Longitude | 0-0-1 | -0-1- | 0-0-0 | -1-1- | 1-1-0 |
| Latitude | -1-0- | 1-1-0 | -1-0- | 0-1-0 | -1-1- |

# Geohash

| Geohash length | Grid (width, height) |
|---|---|
| 1 | (5009.4km, 4992.6km) |
| 2 | (1252.3km, 624.1km) |
| 3 | (156.5km, 156km) |
| 4 | (39.1km, 19.5km) |
| 5 | (4.9km, 4.9km) |
| 6 | (1.2km, 609.4m) |
| 7 | (152.9m, 152.4m) |
| 8 | (38.2m, 19m) |
| 9 | (4.8m, 4.8m) |
| 10 | (1.2m, 59.5cm) |
| 11 | (14.9cm, 14.9cm) |
| 12 | (3.7cm, 1.9cm) |

# Geohash

- Geohash has 12 levels of precision.
- The precision factor determines the size of the grid.
- We aim to find the minimal geohash length that encompasses the entire circle drawn by the user-defined radius.
- In Yelp, we are only interested in geohashes with lengths between 4 and 6.

# Geohash

- Geohashing guarantees that the longer the shared prefix is between two geohashes, the closer they are.
- However, the reverse is not necessarily true: two locations can be very close but may not share a prefix at all.

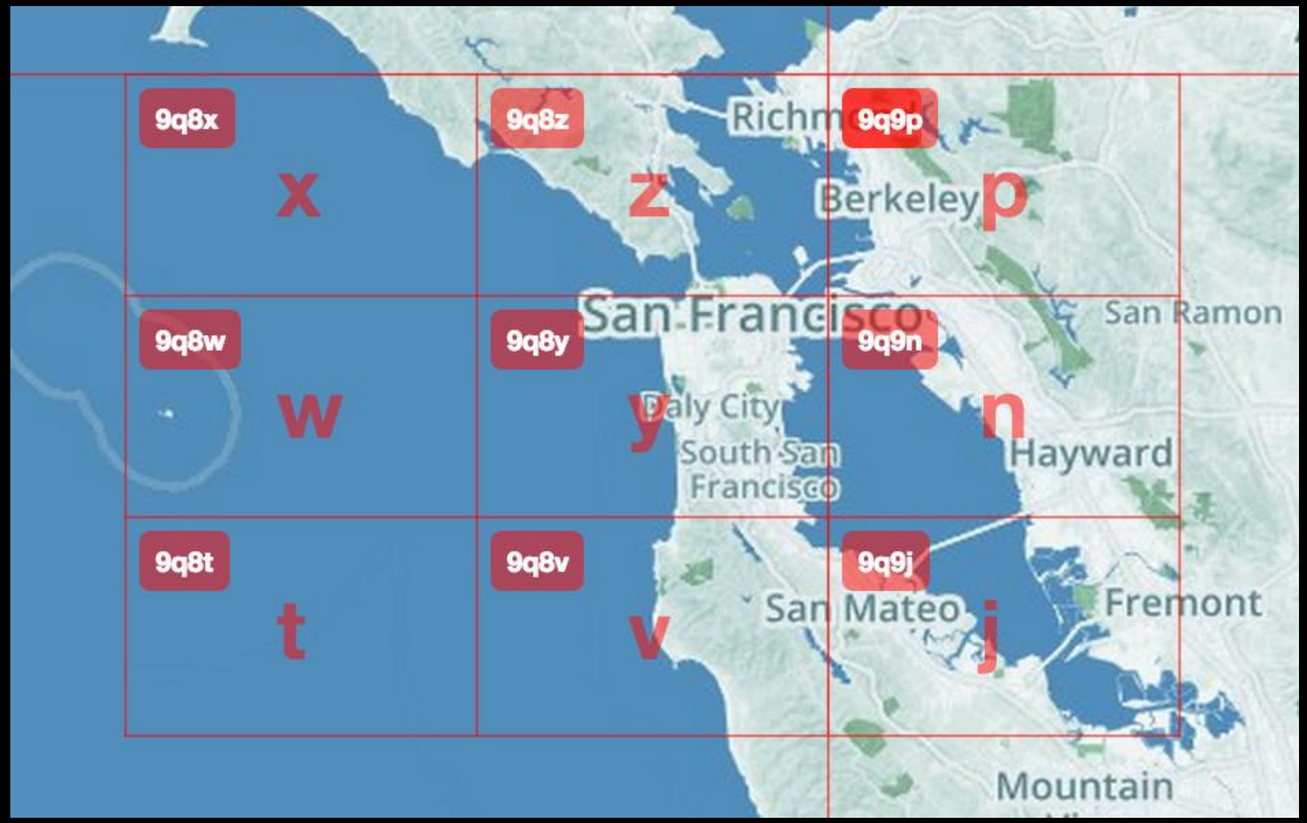

# Geohash

- Two positions can have a long shared prefix, but they may belong to different geohashes.
- A common solution is to fetch all businesses not only within the current grid but also from its neighboring grids.
- The geohashes of neighbors can be calculated in constant time.
- What should we do if there are not enough businesses returned from the current grid and all its neighboring grids combined?
  - We can remove the last digit of the geohash and use the new geohash to fetch nearby businesses.
  - If there are still not enough businesses, we can continue to expand the scope by removing another digit.

# Quadtree

- A quadtree is a data structure that is commonly used to partition a two-dimensional space by recursively subdividing it into four quadrants until the contents of the grids meet certain criteria.
- With a quadtree, we build an in-memory tree structure to answer queries. It is not a database solution.

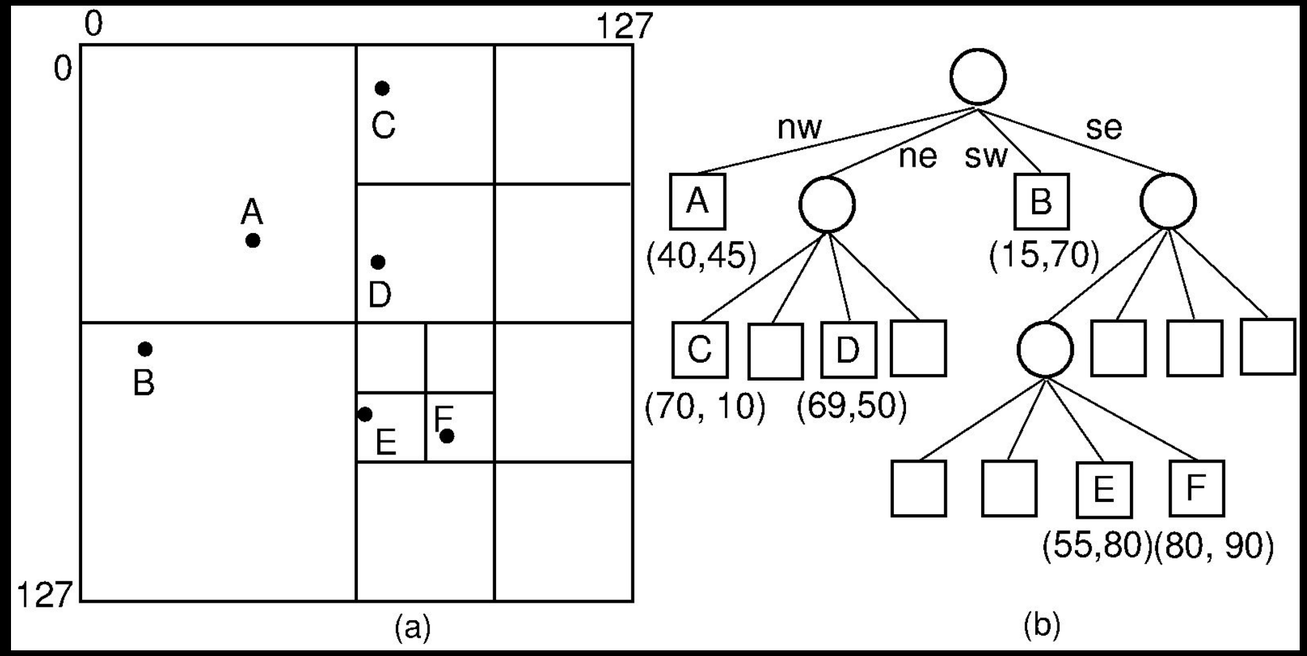

# Google S2 (Hilbert Curve)

- It is an in-memory solution.
- It maps a sphere to a one-dimensional index based on the Hilbert curve.
- Two points that are close to each other on the Hilbert curve are also close in one-dimensional space.
- http://bit-player.org/extras/hilbert/

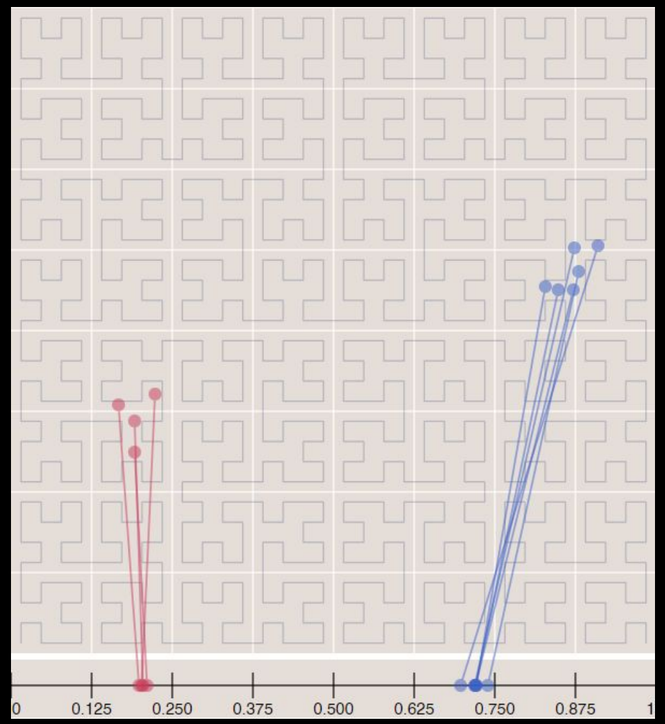

# Comparison

| Geo Index | Companies |
| --- | --- |
| geohash | Bing map, Redis, MongoDB, Lyft |
| quadtree | Yext |
| Both geohash and quadtree | Elasticsearch |
| S2 | Google Maps, Tinder |

# Further Resources

- System Design Interview — An Insider's Guide — volume 2  (pages: 9 - 42)