

Using Binary Trees to Speedup Address Range Lookup in Geolocation Solutions

Gennadiy Poryev, senior scientific worker,

National Technical University of Ukraine "KPI", Kiev, Ukraine

core@barvinok.net +380-44-406-85-24

I. Abstract

In this paper the specifics of the underlying algorithms for geolocation are analyzed. Concerning the storage of IP ranges the drawbacks and bottlenecks of existing implementations are determined. A new, faster method for storing and accessing the IP range databases is proposed.

Keywords: IP geolocation, binary trees, distributed networks

II. Introduction

The beginning of third millennium is marked by the networking technologies in general and Internet in particular reaching so extreme a level of integration and mobility within our everyday life that it would not be an underestimation to compare it to the importance of electric power. Today, both technologies are so crucial to the progress and the very survival of humankind that virtually no activity is done without either one.

Of particular interest is the variety of the so-called geolocation services, now widely popular among content providers. Geolocation is the set of techniques to relate client IP address to his/her **geographical location**. This kind of information helps focusing the advertisement taking into account cultural specifics and local demand, choose potentially

preferable user interface language based on the country, provide search for nearby facilities such as shopping, entertainment, transportation etc.

Since every IP address belong to the IP range registered with some LIRs (Local Internet Registries), which, in turn, are a part of larger IP range superset allocated by the superior LIR or RIR (Regional Internet Registry) it is possible to track business address of a registree entity and presume that its clientele are likely to be located in the neighboring territory.

The usual ways of doing this involve either querying RIR directly or maintain a local cache of RIR database snapshot, against which to query. The query takes a form of plain IP address and outputs the administrative and technical details, including address which is then parsed to extract country and city.

The information relevant to this simple way of geolocation is stored in two places. One being *delegated* table, which lists primary allocated IP ranges, year of allocation and ISO code of the country it's been assigned to. For the more detailed information, RIR clients either refer to RIR's WHOIS server or look up the *inetnum* database which gives more detailed business address of the entities responsible.

There are, however, other ways in which geolocation techniques can be useful. As was mentioned in our previous works [1], a special type of internetworking called “peer-to-peer” (p2p) is taking hold in the modern Internet, by some estimates already generating more than half of total Internet traffic worldwide.

Being able to efficiently manage the large and worldwide p2p network is a challenge and there were many different approaches [2], including

those concerned with optimal clustering of the p2p network. An efficient clustering is impossible without the efficient metric that we call “locality metric”, which is a method to estimate the topological distance between the arbitrary nodes of a network.

While there were many attempts to build such a metric, we believe that our approach envisioned in the series of publications [3,4] is the most suitable for p2p needs. This article, however, is not about improving the estimation method itself but rather about speeding it up, because, as of today, most reference implementations still suffer from performance issues.

On the other hand, geo-location-like methods can help advance topological metric, should we consider that the essence of both is the query for corresponding IP range for a given address.

III. Analysis of an existing solution

Unlike CARMA, current implementation of the locality metric called “SAMPAN” for **S**ervice-oriented **A**ffinity **M**etric for **P**eer-to-peer **A**rchitecture **N**etworks employs separation of the database component from the business-logic component like to widely known Model-View-Controller scheme. The database consists of three tables: *assets*, *atoms* and *linkage*.

The *assets* table is simple and contain only two columns — numeric identifier and a name for an asset in terms of RIR allocations. This is to maintain asset reference by indexes in another two tables. As of August 2012, there were over 12 thousands registered assets in the Internet.

The *atoms* table is a primary table containing both delegated ranges for IPv4 and IPv6 address space and Autonomous System Numbers (ASNs). Its declaration in terms of MySQL language is as shown in fig.1.

```
1 CREATE TABLE `atoms` (  
2   `id` INT(10) UNSIGNED NOT NULL AUTO_INCREMENT COMMENT 'Atom ID',  
3   `isocc` CHAR(2) NULL DEFAULT NULL COMMENT 'ISO country code' COLLATE 'utf8_unicode_ci',  
4   `type` ENUM('ipv4','ipv6','asn') NOT NULL COMMENT 'Atom type' COLLATE 'utf8_unicode_ci',  
5   `baseaddr` BIGINT(20) UNSIGNED NOT NULL COMMENT 'Atom base address',  
6   `volume` BIGINT(20) UNSIGNED NULL DEFAULT NULL COMMENT 'Atom range volume',  
7   `refid` INT(10) UNSIGNED NULL DEFAULT NULL COMMENT 'Reference id',  
8   PRIMARY KEY (`id`),  
9   INDEX `k_baseaddr_volume` (`baseaddr`, `volume`),  
10  INDEX `k_isocc_refid` (`refid`, `isocc`),  
11  INDEX `k_isocc` (`isocc`)  
12 )  
13 COMMENT='Contain ranges, subranges, ASes and their linkage'  
14 COLLATE='utf8_unicode_ci'  
15 ENGINE=MyISAM  
16 AUTO_INCREMENT=3975803;  
--
```

Fig.1. Declaration of the atoms table

Almost all fields in the fig.1 are self-descriptive except the *baseaddr*, *volume* and *refid*. In case of the IPv4 or IPv6 range, the *baseaddr* contain network identifier converted into the unsigned long integer, *volume* lists the total number of addresses allocated in this range (which is not always in the form of 2^n) and *refid* refers to the identifier (*id*) field of the ASN or superset IP range entry in this same table to indicate the linkage relationship. The entities are linked bottom-up, meaning that the *refid* is always pointing to the larger inclusive range or an ASN. In case of ASN, *baseaddr* contains the number of the ASN, while *volume* and *refid* are NULL. As of August 2012 there were over 3.8 million of IPv4 ranges, over 41 thousands of IPv6 ranges and over 55 thousands ASNs in the Internet.

The third *linkage* table contains information concerning the linkage of ASN into assets and assets into each other. Its declaration in terms of MySQL language is as shown in fig.2.

```

1 CREATE TABLE `linkage` (
2   `id` INT(10) UNSIGNED NOT NULL AUTO_INCREMENT COMMENT 'Link ID',
3   `type_a` ENUM('asn','asset') NOT NULL COMMENT 'Type of Endpoint A' COLLATE 'utf8_unicode_ci',
4   `node_a` INT(10) UNSIGNED NOT NULL COMMENT 'ID of Endpoint A',
5   `type_b` ENUM('asn','asset') NOT NULL COMMENT 'Type of Endpoint B' COLLATE 'utf8_unicode_ci',
6   `node_b` INT(10) UNSIGNED NOT NULL COMMENT 'ID of Endpoint B',
7   PRIMARY KEY (`id`),
8   UNIQUE INDEX `type_a_node_a_type_b_node_b` (`type_a`, `node_a`, `type_b`, `node_b`)
9 )
10 COMMENT='Contain link mesh with ASNs and ASSETS'
11 COLLATE='utf8_unicode_ci'
12 ENGINE=MyISAM
13 AUTO_INCREMENT=157636;
..

```

Fig.2. Declaration of the linkage table

The SAMPAN, like CARMA, takes the two IP addresses as input, but, unlike CARMA, uses external MySQL storage engine to perform all relevant queries. These may include the following:

- Query for *id* in the *atoms* table to determine all IP ranges the input address belongs to; check if any of the range cover both addresses;
- Among those ranges found, leave only those with non-NULL *refid* field, meaning only those ranges that are assigned to the ASN;
- For each range found and filtered in the previous query, retrieve all entries of all ASNs; check again to see if both input addresses belong to the same ASN;
- Retrieve all entries from the table *linkage* which contain any of the thus found ASN in either of the *node* fields, indicating that the ASN are a part of larger asset;

- At this point, the algorithm has two lists of assets which are related to two input addresses and those two lists are checked for common entries;
- If the previous point yielded no common asset, another query is made to determine if there exists an asset, which is linked to any two of the assets in the previous lists;
- If no common asset is found, one final query is made to verify if both ranges are registered to the same country;

The experiments undertaken on the SAMPAN reference implementation have shown the performance to be suboptimal and largely unacceptable for p2p deployment scenarios where quick estimation decisions are required. For instance, on the hardware such as four-core AMD Phenom 9550 processor with 4GB RAM and the most favorable running conditions for MySQL Server Community Edition it took almost 10 seconds to pass all calculations for all flavors for both IP addresses. Since most p2p network clients have peer lists on the order of tens, usually around 50-100, that would slow down the network significantly, requiring more than 10 minutes to just process and sort peer list.

The apparent bottleneck was determined to occur while executing the first pair of queries to find all IP ranges a given addresses are belong to. The query involved is following: **“SELECT * FROM atoms WHERE INET_ATON(ipaddr) BETWEEN baseaddr AND baseaddr+volume ORDER BY volume LIMIT 16”**.

Despite properly configured indexing (namely, *k-baseaddr-volume* index specifically for address lookup) the average query execution time is 5 seconds. Even though *EXPLAIN* modifier shows that the index is actually

being used, the number of processed rows remain large, usually one to two millions. This is believed to be caused by the use of *BETWEEN* clauses, since MySQL server is not particularly good in queries for inexact matches like this one.

Therefore, a new approach is needed, one that will be specifically tailored for IP address range lookups.

IV. Using custom binary trees to perform range lookups

Let's consider the nature of IP range lookups in this particular scenario. For any given IP address, it is very common to find more than one range it belongs to. Ranges appear to be not overlapping in the sense that every shorter range is a subrange of the larger one and every shorter range “belongs” to only one larger range, who do not overlap too.

Moreover, the ranges are finite. IPv4 address space takes no more than 2^{32} addresses, but various kinds of reserved ranges lower that value by approximately $5.8 \cdot 10^8$. Since minimal allocation slot is 4 addresses, it is reasonable to estimate the upper number of possible IPv4 ranges as $\frac{2^{32} - 5.8 \cdot 10^8}{4} \approx 9.2 \cdot 10^8$. The actual volume of allocated IPv4 ranges, as mentioned earlier, is two orders of magnitude less due to utilization of the large allocation ranges.

It seems that binary trees may be very useful to organize IPv4 ranges due to their natural adherence to power of two — a descent from the root node to the lowermost layer with the capacity to contain 4 million nodes would take about 22 steps.

Let's estimate the total memory footprint required to hold the extensive database for IPv4 ranges. For every range, the following data should be stored: a) **Start address**, which also serve as tree node identifier; for IPv4 it needs 32 bits; b) **Range volume**, number of nodes in that range; also 32 bits; c) **Reference ID**, number of ASN associated with the range; this requires 32 bits.

If the tree node is not terminating (i.e. containing the range data) but a range lookup fork, it will need two fields of 32 bits each — pointer (branch) to the left and right nodes. Therefore the byte width of the lowermost layer $W_N=12$, for all other layers $W=8$.

Hence the total memory footprint for the binary tree for all ranges consisting of N layers from the root node downward except the latest N-th layer is the following

$$M_s = \sum_{i=1}^{N-1} M_i W + M_N W_N = -(1-2^{N-1})W + 2^N W_N \quad (1)$$

For the proposed scenario, assuming $N=22$ we get $M_s=2^{26}-2$ or 64 MBytes. Again this is considerably less memory than required using SAMPAN approach (whose *atoms* table alone require 512 MBytes and above), and therefore is more affordable for mobile or embedded applications with memory constraints.

Processing range information from *delegated* and *inetnum* tables yield unsorted sequence, so to build sorted binary tree we have to devise the proper filling algorithm.

For every triplet of «range», «volume» and «refid» that enters the tree, we convert «range» from IPv4 standard dotted representation into the 32-

bit unsigned integer using standard BSD sockets function `inet_addr()` so that tree can be sorted by this «index».

The root node is exactly in the middle of all IPv4 range, which equals 2^{31} . Depending on the numeric value of the given IPv4 range, a left or right branch is initialized and the corresponding node of next level is created (if it was not initialized and created prior to that). The process repeats itself on each of subsequent nodes, creating ones and initializing branches if necessary, until the tree depth reaches preset $N=22$. At this point, a data (non-forking) node is created, holding all three values.

When the complete tree is built, the lowermost node layer will contain sorted range array, with ascending start addresses. The rough schematic layout of such tree is shown on fig.3.

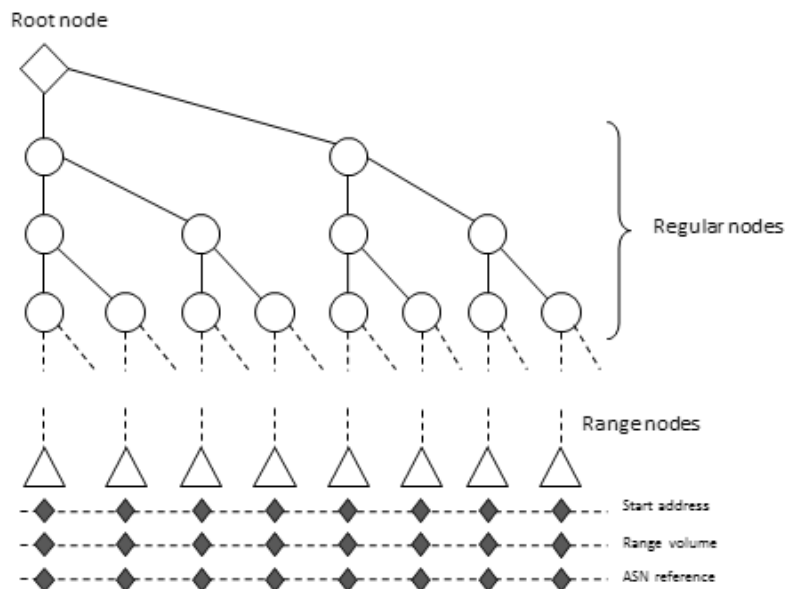


Fig.3. Schematic layout of the custom binary tree to hold IP ranges database

Range lookup method

Let's now consider how the range lookups should be done. To locate IPv4 range, converted input address is compared to the single node on each layer, starting from the root node on the topmost. If the number of input node is greater or lower than that of the sequence node and there exist corresponding branches, another node is chosen for the same lookup based on this branch. If there is no branch, the algorithm stops and returns to the previous level. If the lowermost layer is reached, the algorithm also stops and returns the corresponding node.

The tree-walking algorithm can be expressed as follows. Let $c(I,i,j)$ be a function that return a j -th node M_j on the lowermost layer, and I is the input address. This way, a stored range $M(I)$ can be calculated by the recursive function

$$M(I) = c(I,1,1) \quad (2)$$

where

$$c(I,i,j) = \begin{cases} I < N_i, \exists N_{i+1,2j-1} : c(I,i+1,2j-1) \\ I > N_i, \exists N_{i+1,2j} : c(I,i+1,2j) \\ I = N_i : M_j \end{cases} \quad (3)$$

Since every tree building iteration eventually lands some range at the lowermost layer, there should be no hanging nodes — those with neither branch initialized and not on the lowermost layer.

However, there is a chance that an input IP range would require to choose a non-initialized branch. For instance, if an IP address is allocated right before the registered range starts on the right branch, the algorithm will be stopped because there's no left branch from the previous node. In

this case the algorithm should perform a «backtrack» to find next downward branching to the existing range. This backtrack is done by repeatedly checking the previous nodes until the algorithm reaches a node with two existing branches, which it had already passed, turning to the right branch. This time the left branch is chosen and the lookup occurs again from this node. Hence the lookup should end at the next adjacent or nearby range to the left from where the previous backtracking occurred.

Because IP ranges are not necessary adjacent, the lookup of an IP address may not land at the second range even after tree backtracking. In this case the backtracking is repeated moving generally in left direction (since ranges are «extending» to the right) until a suitable range is found.

V. Conclusions

While it is generally accepted that the computational complexity for the binary tree lookup varies from $O(\log n)$ to $O(n)$ our test runs have confirmed the observed speedup in comparison to the SQL approach corresponds to $O(\log n)$, since only 3.1% of the IP addresses from the testing BitTorrent swarm required tree backtracking at all, and 1.7% required tree backtracking performed more than once to complete range lookup. On the same hardware as mentioned earlier, the average lookup timeout was 120 msec.

Thus we have conclusively demonstrated that using the specially designed binary tree may not only decrease the memory footprint requirements for IP address range lookup routines but also significantly increase the lookup speed.

VI. Bibliography

1. Poryev G. V., Schloss H., Oechsle R. CARMA: A distance estimation method for internet nodes and its usage in P2P networks // International Journal on Advances in Telecommunications.— 2010.— Vol. 3, no. 3,4.— Pp. 114–128.
2. iPlane: An information plane for distributed services / H. Madhyastha, T. Isdal, M. Piatek, C. Dixon // In Proceedings of the 7th USENIX Symposium on Operating Systems Design and Implementation.— USENIX, 2006.— Pp. 367–380.
3. Poryev G., Poryev V. Improved CARMA locality estimation model for peer list reordering and its experimental validation // Реєстрація, зберігання і обробка даних.— 2011.— no. 4(13).— Pp. 3–11.
4. G.V.Poryev. Direct observation of performance improvement in filesharing p2p networks using CARMA techniques // Journal of Qafqaz University.— 2011.— no. 31.— Pp. 16–22.