Surveying Indexing Methods for the Internet of Things

Victor Kardeby(✉), Ulf Jennehag, and Mikael Gidlund

Department of Information and Communications Systems, Mid Sweden University,
851 70 Sundsvall, Sweden
{victor.kardeby,ulf.jennehag,mikael.gidlund}@miun.se

Abstract. The Internet of Things is expected to expand several magnitudes in the coming decade surpassing over 50 billion devices. In the Internet of Things there is a need to support complex queries on the massive amount of information that will be made available. The scalability of the indexes used to support the queries is therefore critical. This paper therefore investigate what type of index that could scale to the size required by the Internet of Things. We find that range query is an approach that support continuously changing information and fast updates with the lowest increase in signaling per participating device. We find that a Chord-based distributed hash table hosting a NUBL range query indexing scheme will scale to the required size while supporting multidimensional range queries.

Keywords: IoT · Internet of things · Survey · Indexing method · Distributed systems · Distributed hash table

1 Introduction

The Internet of Things (IoT) is defined in [1] as billions of devices having and sharing information with services in real-time. IoT promoters proclaim that there will be a torrent of new devices coming online to share information and provide new services within the next years. Some claim 50 or 100 billion devices by 2020 [2], while others are more modest, projecting 24 billion devices in the same period [3]. The need to quickly act upon acquired information requires devices to effectively find and acquire relevant sensor information. It becomes more difficult to acquire relevant information quickly as the amount of produced information increases; this is therefore one of several important research topics [4,5] in the IoT. However, other topics include for example device heterogeneity, privacy, and security. There are several viable approaches available for the current data volumes, but their fitness can be questioned [6] when there is both an information increase in the order of several magnitudes and an increased amount of dynamically changing information.

One approach to solve scalability issues is to increase the capacity of the current network infrastructure. However, increasing the capacity require either

B. Mandler et al. (Eds.): IoT 360° 2015, Part II, LNICST 170, pp. 282–291, 2016.
DOI: 10.1007/978-3-319-47075-7_32

heavy investments in additional hardware, or the exchange of current hardware with more improved versions. Both would have already been done if they were economically feasible. Alternative approaches to sharing, finding, and acquiring information from sensors therefore look toward distributed peer-to-peer networks [7]. Where alternatives to brute force searches or complete data replication can be employed.

1.1 Related Work

Traditional approaches to peer-to-peer networks [8] are constructed to support an index method for exact matches through the use of a Distributed Hash Table (DHT). While such functionality is adequate for rapid resolution of locations for named information, it is unsatisfactory for more rich queries that the IoT will require. A typical IoT query is likely to be multidimensional with a varying amount of parameters which require a different type of index method. In the simpler end of the scale there are queries for certain specified types of sensors within a region, asking for example "All temperature sensors in Sundsvall, Sweden" or "Any temperature sensor that have a humidity sensor attached and is located outdoor near my location". More complex queries could request sensors with some specified values, past or present, e.g. "All temperature sensors that are submerged, have had a measurement above 30 °C within the last week, and that is located within 5 km of my cabin".

In [7] the origin of current peer to peer based multidimensional indexing methods is visualized. It is shown that current research typically examines the query types: window query, range query, or k nearest neighbor query. Their conclusion is that there are two performance related problems that are always considered and thoroughly discussed in the literature. They are load balancing problems and update strategy. Load balance and search performance in a family of index schemes that support wildcard searches is examined in [9]. The related work is divided into different categories based on the indexing complexity. The categories are: (a) exact matches, (b) categorized topics, (c) ranged query, and (d) full text indexing. All DHT support category (a) as part of their basic operation. Then the other categories are added either by mapping an indexing structure, such as a search tree, on top of the DHT (*over-DHT*), or by modifying the DHT algorithm (*in-DHT*). Category (b) indexing is commonly supported by either creating multiple DHTs, one per category, or by arranging the structure differently, such as in [10]. However, with the expected size of the IoT reduces the value of finding information based on categories. For example, requesting every sensor that have a location among millions of devices where most have a meaningful location will only overload the requester with replies. Others, for example [11], use technologies such as prefix based routing and order preserving hash functions to enable category (c) searches [7,10,12,13]. Others [9] focus on category (d), detailing how to enable full text indexing in a DHT.

1.2 Problem Motivation

The problem is to find a type of indexing that is suitable for sensor data stored in a DHT that originates from 50 billion sensors and that is continuously being updated with high sampling frequencies. A suitable indexing scheme should support complex queries such as the examples in Sect. 1.1 while keeping resource consumption as low as possible. In all situations involving small computers, such as the Raspberry Pi, Sensor motes, smartphones, or similar, it is less expensive in terms of hardware resources to perform local processing than communication [14, 15], which mean that indexing schemes that favor local processing over network communication is preferable. The resource that is important to conserve is therefore communication. We therefore target a distributed topology where we measure communication by the discrete number of messages that is signaled as a part of routing requests between peers and performing maintenance of the indexing schemes. It is determined that the indexing category (a), exact matches, is supported by all options and cannot perform advanced queries. Category (a) is therefore of no interest to the problem. Category (b), categorical indexing, would include several categories that would contain every device and query results in these categories would then overload the requesting device and therefore be of little use to reduce the query. Full text indexing, category (d), would require extensive updates for continuously changing information and is therefore rejected as well. Therefore, this paper analytically investigates which algorithms for range queries (c) that could be supported by the IoT for dynamic information in a worst case scenario where large quantities of devices continuously insert new data and perform complex queries. Hence, an indexing scheme must support these following two requirements:

1. It must minimize the total number of signaled messages in a query.
2. It must reflect updated values quickly.

Our evaluation of requirement (1) focus on the total number of messages required for the query. Requirement (2) is evaluated by investigate the number of messages required to update the index in conjunction with a value update. The main scientific contribution of this paper is, (1) a survey of suitable indexing methods for continuously changing information for the volumes of information that is expected from the IoT, and (2), a recommendation of an indexing method that is indicated to scale well for the IoT.

1.3 Outline

The remainder of this paper is organized as follows. Section 2 details our approach to solve the problem. Section 3 examines several DHT to compare scalability of maintenance and lookup performance. Section 4 evaluates the listed DHTs, while Sect. 5 presents our conclusions.

2 Approach

In order to determine the scalability of range query indexing schemes we compare the number of messages that are required to route a request to the recipient and the number of messages that are transmitted during maintenance. The number of messages are compared to a typical well scaling DHT, Chord [16]. A number of different structures for both indexing and DHT are then selected to represent different topologies. Analytical expressions for the signaling are identified in the literature for every investigated indexing structure. *In-DHT* approaches are compared directly while *over-DHT* approaches are assumed to be used in conjunction with Chord if no alternative is given. All indexing schemes are then normalized to Chord and compared with each other. If an indexing scheme lacks analytical expressions for scalability then an extrapolation of presented measurements is used instead. An indexing scheme is determined to be scalable if its normalized values remain parallel to Chord. Any indexing scheme whose values remain less than Chord thus scale better and consequently any value above Chord scales worse. Intuitively we expect that well scaling complex indexing schemes remain parallel to Chord.

3 Indexing for the IoT

A number of DHT that support category (a), exact matches, are listed in Table 1 together with their individual routing and maintenance complexity. Chord [16] is selected as the base to compare other approaches to, since it is one of the most well-known DHT. It structures participating peers in a ring and maintain links, denoted fingers, to other peers and thus enable a distributed binary search to locate data logarithmically. CAN [17] is also a well-known DHT that form a Cartesian space that span the key space. Many researchers [18] have based their own approaches for supporting the other categories on these two DHT structures. Cycloid [19] and P-Grid [20] are examples of alternative topologies that contend with the performance of two previous DHTs. Cycloid distribute information over a hypercube with d dimensions where each vertex is a ring structure with 2^d identities. The total number of addressable peers is thus $n = d * 2^d$. Every peer in Cycloid maintains a list of 7 peers regardless of the number of participants. P-Grid is an example of a DHT that support more than index category (a). Every peer in P-Grid maintain a routing tree containing O(Log N) peers and can route messages in O(Log N) jumps.

3.1 In-DHT Approaches to Range Queries

In-DHT approaches modify the existing algorithms for routing and storage in a DHT to support new functionality. Table 2 displays a set of *in-DHT* approaches. The table presents the name, origin of the DHT, and the updated routing and maintenance cost. P-Grid support range queries in its basic operation and is therefore included here in addition to exemplifying a basic DHT. Chord# [12]

Table 1. Basic overlay networks

Name	Form	Routing	Maintenance
Chord [16]	Circle	$\frac{LogN}{2}$	$LogN$
CAN [17]	Multi-Torus	$\frac{d}{4} * N^{\frac{1}{d}}$	$d * 2$
Cycloid [19]	Hypercube	O(D)	7
P-Grid [20]	Mesh	$\frac{LogN}{2}$	$LogN$

remove the hashing function, sorting keys lexicographically, modify the data placement algorithm to restore load balancing properties, and exchange finger table algorithm to adapt to lexicographical searches. Finally the binary search algorithm for finding data is modified to enable b-dimensional searches. This changes the lookup signaling from O(Log$_2$ N) to O(Log$_b$ N) messages. SONAR [12] extends Chord# to enable multidimensional range queries. LORM exchange the algorithm that creates the identification for the data to use a data type category as one part of the two-dimensional identity. Where the other part is created using a locality-preserving hash algorithm [21].

Table 2. DHT with advanced indexes

Name	Origin	Routing	Maintenance
P-Grid [20]	P-Grid	$\frac{LogN}{2}$	$LogN$
Chord# [12]	Chord	$Log_b(N)$	$Log_b(N) * \frac{b-1}{2}$
SONAR [12]	Chord#	$-1.2 + 0.62 * log_2(N)$	$Log_2(N)$
LORM [10]	Cycloid	$Cycloid * d$	7

3.2 Over-DHT Approaches to Range Queries

An Over-DHT approach utilizes the existing structure established by the DHT to handle systemic issues, such as load balancing, churn handling, and routing by using the basic operations of the DHT. Table 3 show a representative set of *over-DHT* algorithms, their name, structure, maintenance overhead and signaling overhead. NUBL [22] is based on an earlier approach and construct a tree where leaf nodes are encoded and stored as object in the DHT. This enable peers to search through the tree to locate information within the desired range. The routing of a query is improved to O(Log$_2$(D/4)/2) using Chord with only O(D/2) in increased maintenance. DRing [13] constructs an additional overlay circle on top of the existing structure. This super-node structure consists of leaf nodes in a prefix based routing tree. The super-node structure reduce the tendency to overload the upper levels of the tree by traversing requests through caching. The DRing algorithm claim that optimal performance is achieved when queries greatly outnumber inserts.

Table 3. Over-DHT indexing algorithms

Name	Form	Routing	Maintenance
NUBL [22]	Tree	$\frac{Log_2(\frac{D}{4})}{2} * Chord$	$\frac{D}{2} + Chord$
DRing [13]	PHT	O(Log N)	O(Log N)

4 Results

The approaches are examined in a worst case scenario where up to 10^{10} devices continuously perform updates of their sensor data while performing complex queries in the DHT. Parameters, such as network topology, transmission delay, and query rate, are selected to be equal in each case and are therefore eliminated. All results are normalized to Chord, as it is determined to be well scaling for large number of devices. The figures will therefore show a message ratio where any solution that is parallel to 1 is determined to scale comparable to Chord. An approach that have a decreasing ratio that is <1 scales better than Chord, while having an increasing ratio that is >1 indicates that it scales worse.

In Fig. 1 the signaling is presented for the DHTs in Table 1, with the exception of CAN. Normalizing the routing and maintenance of CAN to Chord resulted in

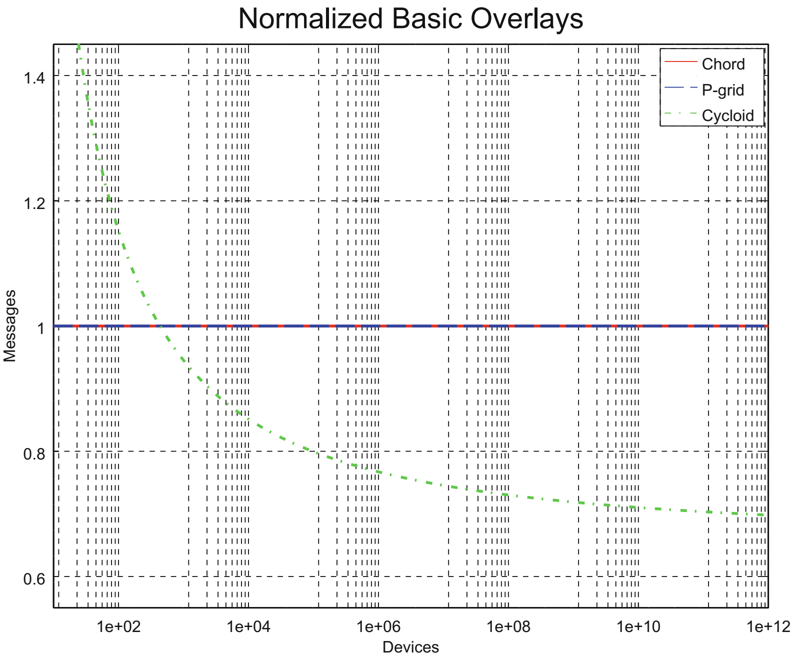

Fig. 1. Lin-Log plot of messages per node for basic overlay networks normalized to Chord for up to 10^{10} devices.

an exponential curve that ended at one thousand times the signaling in Chord for 10^{10} devices. Cycloid is shown to outperform Chord when the number of participating devices exceed 10^3. Since the maximum number of participants in Cycloid is determined as $n = d * 2^d$ where d is the dimension of the hypercube we opted for solving d for the current n along the x-axis. This simplification thus show the ideal solution at any given size and a practical implementation would require additional messages to be transmitted for each operation.

4.1 In-DHT Indexing Schemes

In Fig. 2 we examine the *in-DHT* approaches P-Grid, SONAR, and LORM. P-Grid follow Chord at 1. However, an analytic expression for SONAR routing were not found, and logarithmic coefficients were therefore found in [12] through logarithmic fitting. The discovered routing coefficients for SONAR are shown in Table 2. LORM is based on Cycloid and use the same equation as in Fig. 1. LORM is capable of maintaining multiple dimensions simultaneously but LORM is limited to a single dimension in this comparison since all other approaches would require multiple concurrent DHT to support the same. It is observed that there is an intersection between all three approaches in the vicinity of $5 * 10^2$ devices. After the intersection LORM outperform the other solutions. SONAR is

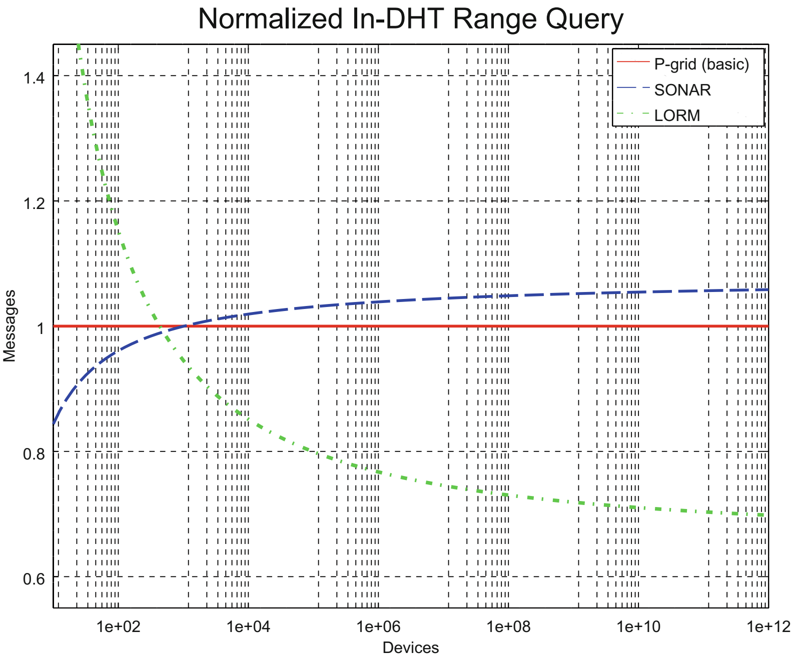

Fig. 2. Lin-Log plot of messages per node for *in-DHT* range queries normalized to Chord in peer-to-peer networks for up to 10^{10} devices.

performing better than the other approaches before the intersection and is then slightly more expensive than Chord. Since maintenance in LORM is constant, there is only one factor that scales logarithmic as opposed to two logarithmic scaling factors in Chord which implies that LORM will converge to 0.5 for very large numbers.

4.2 Over-DHT Indexing Schemes

Figure 3 display Chord and LORM which is compared with NUBL. Since NUBL is an *over-DHT* approach its routing is multiplied with that of Chord and its maintenance is added to the Chord maintenance. NUBL would perform better in conjunction with Chord than Chord by itself according to these calculations. Using only one dimension in NUBL yields a very low signaling when compared to Chord. Setting $d = 4$ make NUBL seems to converge with LORM for large networks with $>10^{10}$ devices. This behavior is caused by calculating logarithms for values less than one, making scalability estimations for $d < 4$ impossible using this method. The total routing term from Table 3 is less than one for $d < 8$ due to another division by 2.

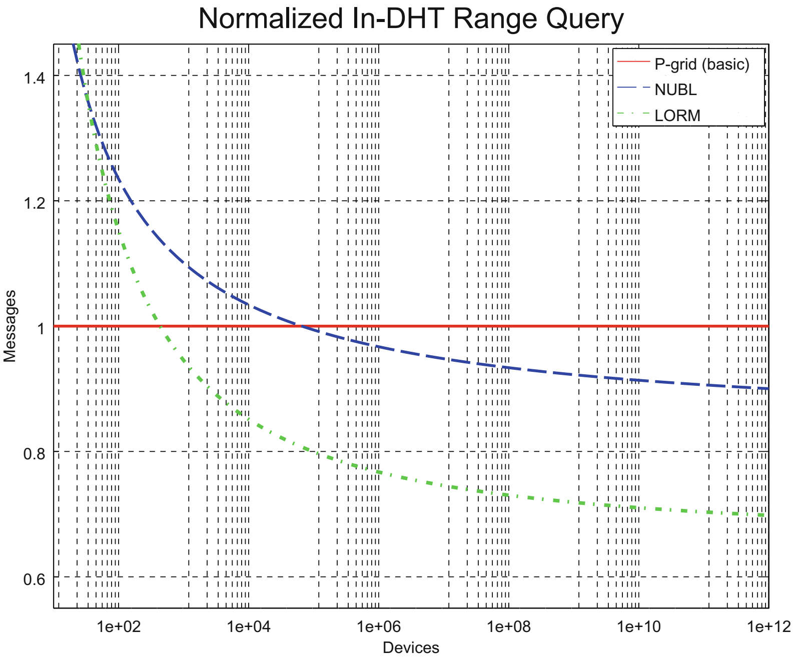

Fig. 3. Lin-Log plot of messages per node for *over-DHT* range queries normalized to Chord in peer-to-peer networks for up to 10^{10} devices.

5 Conclusions

The goal of this paper was to find what type of indexing the IoT would be capable of supporting for sensor information which is continuously changing. The focus was on distributed approaches that scale to billions of devices. The initial analysis of the literature concludes that range query indexing is the only indexing scheme in this survey that is viable for the scale that the IoT require. Several range query approaches are then investigated and evaluated based on two requirements. The indexing scheme must (1) transmit as few messages as possible in a query, and it must (2) reflect updates in information quickly. The approach was to compare the indexing schemes to a known scalable solution. Section 4 evaluated several indexing schemes that support range queries in terms of query length and maintenance. It is shown that LORM outperform the other alternatives, however the ideal parameters that were used in the evaluation of LORM are difficult to implement in practice. The second most effective algorithm is NUBL, which provide functionality for multidimensional range queries and still require less signaling than Chord. Both LORM and NUBL based approaches therefore require further investigation through both simulation and trial deployment in a testbed, to evaluate the possibility of selecting the parameters of LORM such that it is practical to implement and then compare that implementation to one of NUBL.

References

1. Xu, L.D., He, W., Li, S.: Internet of things in industries: a survey. IEEE Trans. Ind. Inform. **10**(4), 2233–2243 (2014)
2. Sundmaeker, H., Guillemin, P., Friess, P.: Vision and challenges for realising the Internet of Things. In: CERP-IoT, European, March 2010
3. Gubbi, J., Buyya, R., Marusic, S., Palaniswami, M.: Internet of Things (IoT): a vision, architectural elements, and future directions. Future Gener. Comput. Syst. **29**(7), 1645–1660 (2013)
4. Paganelli, F., Parlanti, D.: A DHT-based discovery service for the internet of things. J. Comput. Netw. Commun. **2012** (2012)
5. Perera, C., Zaslavsky, A., Christen, P., Compton, M., Georgakopoulos, D.: Context-aware sensor search, selection and ranking model for internet of things middleware. In: Proceedings - IEEE International Conference on Mobile Data Management, vol. 1, pp. 314–322, March 2013
6. Zhang, B., Mor, N., Kolb, J., Chan, D.S., Goyal, N., Lutz, K., Allman, E., Wawrzynek, J., Lee, E., Kubiatowicz, J.: The cloud is not enough: saving IoT from the cloud. In: 7th USENIX Workshop on Hot Topics in Cloud Computing (HotCloud 15). USENIX Association (2015)
7. Zhang, C., Xiao, W., Tang, D., Tang, J.: P2P-based multidimensional indexing methods: a survey. J. Syst. Softw. **84**(12), 2348–2362 (2011)
8. Lua, E.K., Crowcroft, J., Pias, M., Sharma, R., Lim, S.: A survey and comparison of peer-to-peer overlay network schemes. IEEE Commun. Surv. Tutorials **7**(2), 72–93 (2005)
9. Joung, Y.J., Yang, L.W.: On character-based index schemes for complex wildcard search in peer-to-peer networks. Inf. Sci. **272**, 209–222 (2014)

10. Shen, H., Xu, C.-Z.: Leveraging a compound graph-based DHT for multi-attribute range queries with performance analysis. IEEE Trans. Comput. **61**(4), 433–447 (2012)
11. Aberer, K., Cudré-Mauroux, P., Datta, A., Despotovic, Z., Hauswirth, M., Punceva, M., Schmidt, R.: P-Grid: a self-organizing structured P2P system. ACM SIGMOD Rec. **32**(3), 29–33 (2003)
12. Schütt, T., Schintke, F., Reinefeld, A.: Range queries on structured overlay networks. Comput. Commun. **31**(2), 280–291 (2008)
13. Hidalgo, N., Rosas, E., Arantes, L., Marin, O., Sens, P., Bonnairey, X.: DRing: a layered scheme for range queries over DHTs. In: Proceedings - 11th IEEE International Conference on Computer and Information Technology, CIT 2011, pp. 29–34 (2011)
14. Landsiedel, O., Wehrle, K., Gotz, S.: Accurate prediction of power consumption in sensor networks. In: The Second IEEE Workshop on Embedded Networked Sensors, EmNetS-II (2005)
15. Carroll, A., Heiser, G.: An analysis of power consumption in a smartphone. In: Proceedings of the USENIX Annual Technical Conference (2010)
16. Stoica, I., Morris, R., Liben-Nowell, D., Karger, D., Kaashoek, M., Dabek, F., Balakrishnan, H.: Chord: a scalable peer-to-peer lookup protocol for internet applications. IEEE/ACM Trans. Netw. **11**(1), 17–32 (2003)
17. Ratnasamy, S., Francis, P., Handley, M., Karp, R., Shenker, S.: A scalable content-addressable network. In: Proceedings of the 2001 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, pp. 161–172 (2001)
18. Zhang, L., Mitton, N.: Advanced internet of things, pp. 1–8 (2011)
19. Shen, H., Xu, C.Z., Chen, G.: Cycloid: a constant-degree and lookup-efficient P2P overlay network. Perform. Eval. **63**(3), 195–216 (2006)
20. Schmidt, R.: The P-Grid System - Overview, pp. 1–11 (2007)
21. Bongers, E., Pouwelse, J.: A survey of P2P multidimensional indexing structures, arXiv preprint arXiv:1507.05501 (2015)
22. Fu, Y., Hu, R., Chen, J., Wang, Z., Tian, G.: An improved lookup algorithm on over-DHT paradigm based P2P network. In: Gong, Z., Luo, X., Chen, J., Lei, J., Wang, F.L. (eds.) WISM 2011, Part I. LNCS, vol. 6987, pp. 200–207. Springer, Heidelberg (2011)