

Network and content aware information management

Sergey Boldyrev, Ian Oliver, Ronald Brown, Juha-Matti Tuupola, Arto Palin and Antti Lappetelainen
Nokia Research Centre
Ubiquitous Architectures
{sergey.boldyrev, ian.oliver, ronald.brown, matti.tuupola, arto.palin, antti.lappetelainen}@nokia.com

Abstract

The presented approach addresses the problem of query and persistent query (subscription) resolution, taking into consideration distribution across multi-domains, network infrastructure and content management. This approach is particularly suitable for information-centric and cloud computing applications based around a mobile-device infrastructure.

1 Introduction

Traditionally, distributed information management has been based on sets of rules or policies which are predetermined by benchmarking and testing. These techniques may provide useful and accurate results when the architecture of a distributed information system is static. However, such solutions to distributed information management problems can generate non-optimal results when applied to dynamically changing environments, such as Smart Spaces.

Modern mobile device ¹ usage is moving towards implementation in Smart Spaces. A Smart Space may be an environment where a number of devices may use a shared view of resources and services through a non-centralised, information-centric and cloud computing infrastructure [14, 12]. In this regard, Smart Spaces can provide improved user experiences by allowing users to flexibly introduce new devices and access most or all of the information available in the multiple device system from any of the devices. However, a problem with information management in smart spaces can be that the location of desired information is not static. As a result, difficulties can arise in predicting the most efficient paths to access data included in the devices of the Smart Space.

Smart Spaces include the concept of a shared repository of information, also known as a whiteboard, blackboard or

¹particularly in our case though we obviously include 'fixed' devices such as personal computers

tuple-space. The whiteboard can be considered as a storage element to which every device or node in a Smart Space has access [15]. Logically, only one whiteboard exists in a Smart Space and hence every device has access to the very same information. However, the whiteboard is merely a logical concept, and in reality the whiteboard may be comprised of a multitude of physical data repositories.

2 Background

A challenge in providing consistent information management in a smart space is that the resources (e.g., the storage components comprising the whiteboard) are distributed over several physical devices, and as a result, the information consumers are not always in, or part of the same device as the desired information. For example, a number of content capturing devices (e.g., cameras, microphones, etc.) may provide information, while other content rendering devices (e.g., displays, speakers, etc.) may retrieve the information for presentation.

A trivial means for managing the shared information in these dynamic environments so that any device may readily retrieve queried information could be to copy all information to all devices. However, such a solution is typically unfeasible because some devices - especially mobile devices - may have limited capabilities (i.e. connectivity and bandwidth, storage space, computation abilities, etc.) and the cost (e.g. power consumption, impact to performance, etc.) of transferring all information over a communication channel to all devices may be excessive. Further, the dynamic architecture of the Smart Space also makes it difficult to guarantee that all the information will reach every device.

Alternatively, a combination of replication, stripping, dispersing and caching can be used to improve performance and information search efficiency, but it is not sufficient in a dynamically changing environment of a Smart Space. In such an environment a certain amount of statistical learning and clustering analysis is usually needed to forecast information request patterns with respect to the requested information location. However, even that still leaves the follow-

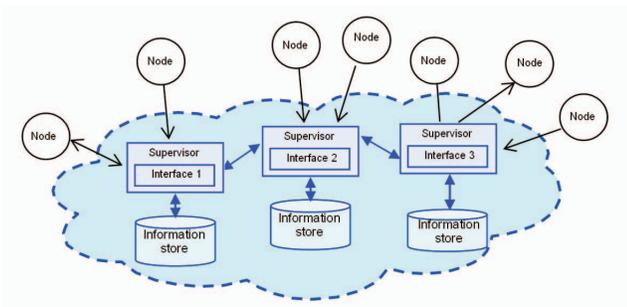


Figure 1. Smart Space architecture

ing challenges unsolved:

- Achieving high degree of predictability and responsiveness
- Making the management of the system adaptive to dynamic changes in the whole infrastructure and its content.

Futhermore we also see the need - based upon the above constraints - for any distribution system not to be simply a synchronisation mechanism but rather one that supports more intelligent asymmetric distribution that can take into consideration aspects such as context and even legal restrictions such as copyright and intellectual property.

Accordingly, the development of a mechanism for solving these problems while still providing the best possible user experience in terms of latency, reliability, or any other quality of service criterion is desirable. In particular, it would be desirable to provide a mechanism that can consider a dynamic network architecture, but also predict the location of desired information within the network.

3 Description

Based on the architecture described in [13] and visualised in Figure 1 the Information store of a Smart Space (Figure 1) stores the information of the Nodes, and any other information seen necessary by the Smart Space. This can include, for example, information of current state or activity, observations of the outside world or just maintenance information.

There are a fixed number of Information stores $\{Nstor\}$. Each store may have a connection to another which can be expressed in terms of a cost function. Nodes are the basis for all functionality within a Smart Space. Node has three important parts, namely External Interface, Information store Interface and Task. External interface considers Nodes interaction with the external world (e.g. with respect to end-user). Information store interface is used to put information to and react on the information in the information store. Task defines the relationship between these two

dimensions. Nodes can insert information, remove it, subscribe to the information by means of the persistent query (subscription) and cancel such subscriptions. There is Supervisor per each interface which monitors the frequencies of inserts and removals as well as subscriptions and appropriate cancellations are taken to determine the best place of subscription and inserts allocation. Thus the most demanded information is tracked, and the information gain can be constructed. The connectivity information gain is provided by the connectivity Controller, and the following key elements constitute the core of the approach:

- planning a path to target information - constructing an optimal path to the information, in the scope of the network and data properties, or, using “hints” to form the path.
- analysing information concentration (information divergence, in the sense of the query satisfaction and information location)
- using queries and subscriptions to derive additional knowledge for shaping the path to target information

While in the process of query/subscription distribution a certain number of Supervisor Interfaces and appropriate Information Stores can be determined as Information Gateways, which are the entry points and disseminators of all information related to query/subscription resolution. Thus, any further queries access only Information Gateways to fetch queried information, which is done by means of acyclic graph management. There, the Information Gateway constitutes the root structure of such graph.

A method is presented for optimal query/subscription distribution in a Smart Space using a combination of content and connectivity information.

3.1 The system

Considering the information access as a following 4 steps in resolving a Node query or subscription,

1. query information locally
2. check results as a match for the original query
 - (a) if not satisfied, then the query is forwarded to the next best information store candidate
 - (b) certain accompanying information redistribution actions should be taken then
3. actual access to information and referred content
4. respond to a query

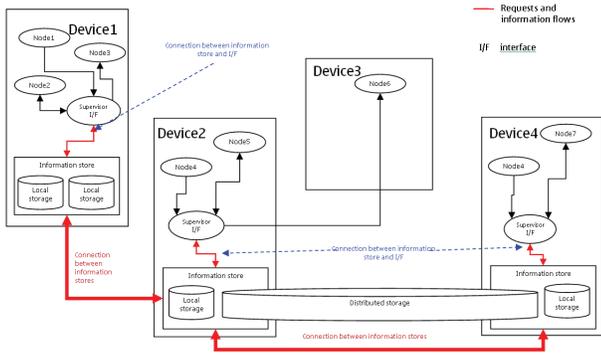


Figure 2. Distributed information framework

In fact the information distribution (Figure 1) can be represented as a model of the actual behavior of a certain system part (Figure 2).

Despite the fact of multi-domain convergence still distributed information management stands on complex and computationally expensive cross-domain relationships findings and extrapolations. Thus, most of the systems encounter issues in scalability, robustness and durability.

The main reason of such issues comes from the intelligent engine, which is based on complicated cross-correlation findings and analysis. Even though such complexity, and corresponding overhead, is compensated by resulting performance, it is obvious to see limited applicability in case of energy and computationally constrained mobile device. Therefore this raises the question of balanced, in energy and computational aspects, distributed information management construction - managing information allocation with minimum cost; allocation mechanism which is aware of information features and network properties at any moment of time.

Zooming in existed distributed information management infrastructure the following typical considerations are valid. A three tier architecture incremental low-level routing mechanism, a distributed query planner and a distributed directory management mechanism. Any layer above distributed directory management is defined to serve as a distributed information location and retrieval mechanism. The main purpose is to guarantee sustained information evolution management, serialization and access control. To undertake such actions in most efficient (intelligent) way a decision update mechanism considered. Decision update mechanism utilizes information gathered from meta-data (actual content and query related) and network, which is fused and delivered as a conditional rule.

Incremental low-level routing mechanism provides routing and message passing facilities between network mapping facilities and connection selection/transfer.

Below the incremental low-level routing layer a corresponding connectivity should be provided. By means of that any network specific information is delivered by connectivity layer and by means of intelligent workload control with service info which includes connection specific details. The granularity level of information can be adjustable.

As it is defined by [8], intelligent workload control can be based on external services specifications which are performance requirements and utilisation, e.g. access pattern. The central role of intelligent workload control plays is infrastructure resource provisioning subsystem which is based on workload-resource mapping and distribution-admission control. These two parts absorb information from actual network topology and service availability, and network conditions and traffic patterns. All elements above are converged by means of resource management and actual performance measurements.

Under current approach there are two main domains to gather information from. They are data specific and network specific domains. Data specific information can be delivered by distributed object file system infrastructure and can consist of commonly used meta-data, object distribution and hierarchy. Network specific information is delivered by connectivity layer and can consist of actual network topology, network conditions and traffic pattern information.

From the Network side the following information is relevant to network specific domain analysis (for example).

- Interface properties
- Adjacent nodes properties
- Last action type
- Timestamp of the last action
- Node access info

This information is provided by the intelligent connectivity controller. Intelligent connectivity controller is aware of its surrounding environment in addition to its local connectivity capabilities. It may also be aware of the physical limitations and of individual devices within the wireless range or within the wired network is associated. The aim of the interface provided by the controller and visualised in Figure 2 is the information distribution framework component is to provide enough yet abstract connectivity properties of the devices participating in to the distributing network. In addition the controller hides the complexity of multi-transport control mechanisms as a connectivity cost function interface.

In time domain the connectivity controller may have a multi-radio controller function that is responsible for allocating connectivity resource based on the communication medium activity, resource availability, and wireless spectrum availability to mention few. It has no direct effect to

the interface provided to the distribution framework. Connectivity controller implements protocols and data types for creating network topology map and connectivity technology map of the environment. These two property maps enable more sophisticated and power efficient transport selection upon data delivery. Also, the connectivity map protocols may be used to share information about the physical properties of each device in the operating in the network such as remaining battery life, available memory resources, computational capabilities etc. This information may further assist connectivity controller to draw conclusions about the most optimal next node to which the data should be delivered. One of the main advantages of such intelligent connection controller is that in the presence of multi-transport devices, heterogeneous networking technologies may be used to perform one data delivery task. For example, due to the dynamic nature of ad hoc networks the initial data delivery from one node to another may be done with Bluetooth but when the receiving node moves out of the range of Bluetooth radio the connectivity controller may open another connection between the nodes using WLAN and continue data delivery. The decision of such intersystem hand-over may be done based on the connectivity map information and physical characteristics of the participating nodes. Finally all the information may be provided as an input to the cost function which the distributing engine may use to determine data locality and availability.

3.2 Usage

Since there is similarity with storage infrastructure, two essential data related analyses are vital here – data locality and concentration. Data locality is analysed in terms of temporal and spatial locality.

Data locality shows the actual proximity of data to the potential consumer in terms of costs. Since cost function is rather compound dependency from several parameters, obviously, proximity is determined in terms of those parameters and is non-linear by nature. The information provided by the connectivity layer as it is described above, i.e. available connectivity technologies, neighbor devices and their connectivity capabilities etc. is the input to the cost function. Also the device characteristics and capabilities like remaining battery life, free memory resources etc. can be considered as input of the cost function when determining the data locality and the optimal path to that information.

Data concentration shows the number of available data pieces per certain locality (e.g within certain proximity). Data concentration serves as input parameter of local workload model and, is derived from content dispersing estimation and data tracker.

Therefore the path construction converges to the efficient query requests update rule mechanism which is based on

two domains information analysis and fusion. Since information which is used to update the management rule is rather independent (orthogonal) there are no ways to use correlation analysis and any derivative as well. Considerably different approach consists of domains (data and network) decomposition and fusion based on parameters covariance analysis.

The optimal operational mode above is determined by reliability (e.g. data loss, consistency), performance (e.g. latency, throughput) and power consumption strategies of the overall system.

The operational mode controls, for example, whether it is better to replicate information to several places for the easy query intersection or to pass the query to some close, however, remote location relative to the information original location while replicating information there.

3.3 Strategies

For different aspects of the system we explain the strategies used by the distribution system.

Power consumption: Different devices have different power budgets and constraints (e.g. battery vs. mains powered). The power consumption associated with the connection maintenance and data transfer between the devices varies and should be considered in calculating the overall cost.

Performance: The future latency requirement for the particular information access is estimated by monitoring the local and remote workloads [7]. The information subscribers may embed some meta-data for the additional latency requirements to the query.

Reliability: Due to the dynamic nature of the system some particular information access may not be granted at all times. This compromises the reliability and availability of the data. To counter this, the redundancy of the information and/or query can be increased. For example, the information can be replicated to several places or, by any other means, information can be dispersed to minimize the loss if one or more information sources become unavailable.

Operational strategies are to be selected in terms of the above mentioned three dimensions. For example a strategy can aim to minimize the power consumption while meeting the latency and reliability constraints. Latency minimization and reliability maximization strategies are constructed in a similar way. Further, by introducing weighting factors, multiple (or all) dimensions can be optimized simultaneously.

Operational strategies are taking in account storage space constraints as a prime limitation of any physical device. Due to the approach presented in the current mechanism, which is based on redundancy after stored information is dispersed, there is natural relationship between stor-

age space and above mentioned parameters. The process of content dispersion consumes a particular additional storage space which means additional spending of power, changes of reliability and performance (in particular latency). If dispersed information is replicated, obviously, these dispersed blocks incur additional storage space needs and consumption is growing further. In that scope any increase of a redundancy (in terms of pure disperse or replication, or both of them) leads to the changes in terms of power, data reliability and performance, thus providing the relationship between storage space consumption and all three parameters which are explained above.

Proposed path to the target information, the allocation of the subscriptions and the information fusion scheme can be characterized by the convex combination of two domains information covariances. As a baseline it takes a convex combination of the means and covariances in the information space (two domains space). Thus, cross-domain and communicative uncertainties are filtered out.

Therefore, from the context specific actions the following allocation and retrieval mechanism can be efficiently constructed utilising data-network specific information analysis and fusion (in terms of Smart Spaces):

- Group-based (physical device action, data reliability-energy-performance-timing)
- Subject-based (Smart Space node action)
- Content-based (codewords, say data blocks, action)

Algorithm sketch can be presented as follows (Figure 3):

1. Decision rule bootstrapping and information gain construction
2. Domain X (network related) feature selection
3. Domain Y (data related) feature selection
4. Features weighting and filtering
5. Covariance calculation
6. Convex covariance calculation [1, 2]
7. Decision rule updates (p. 2)
8. Append updates to workload model (encode updates by means of action model)
9. Back propagate using intelligent workload control infrastructure [8]

To satisfy the needs of proposed approach each Smart Space physical device should to provide feature selection from the data specific and the connectivity specific domains. The fusion and update scheme should be constructed afterwards.

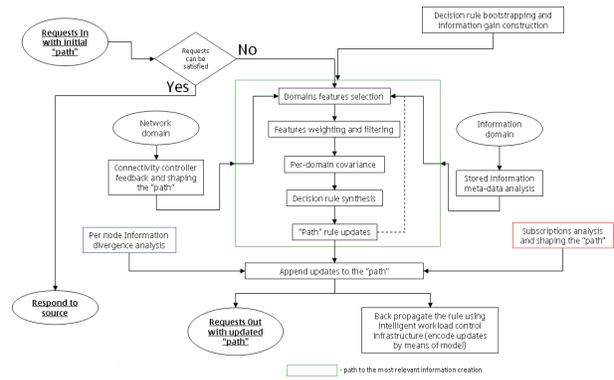


Figure 3. Algorithm flowchart, as it is seen by the particular mobile device within Smart Space

The connectivity related information gain (characteristics gain) is obtained from the Connectivity controller, as it is presented above.

Data specific feature selection consists of meta- and raw data filtering and decomposition for the further fusion. The connectivity specific feature selection consists of network/connectivity information filtering and decomposition for the next step fusion. The process of feature selection can utilize any existing filtering technique like particle filtering [6]. In particular, information gain is obtained by aggregation of meta-data from low-level data allocation analysis, meta-data like “inode” straight through the meta-data provided by storage filesystems, bindings between high abstract information and Information stores content.

Considering that the information gain of the data related domain can be constructed by the assumptions presented in [10] by taking a set of aggregate queries $Q = \{q_1, \dots, q_m\}$ over a set of k distinct data sources, a set of data source readings may be a vector $x = (x_1, \dots, x_k) \in \mathbb{R}^k$. There is an assumption that any query requests are aggregate value of some subset of the data sources at some desired frequency. This then allows each query to be expressed as a k -bit vector: element j of the vector is 1 if contributes to the value of x_j , and 0 otherwise. And, as a result, the value of query q_j on data source readings x is expressed as the dot product $q_j \cdot x$.

Once the features are extracted and the information gains are constructed they can be passed to the actual fusion stage. Since fusion of the convex combination of two domain covariance is used, the process include the following [1, 2]. There, starting from two given estimates \hat{x}_1 and \hat{x}_2 of the true state x with corresponding positive definite error variance matrices P_1 and P_2 , a combined estimate \hat{x} with error variance matrix P can be sought. The overall estimate can

be given by basic convex combination of the two estimates, and, it is as follows [1, 2]

$$P_{xx}^{-1} = \omega P_{x_1x_1}^{-1} + (1 - \omega) P_{x_2x_2}^{-1}$$
$$P_{xx}^{-1} \hat{x} = \omega P_{x_1x_1}^{-1} \hat{x}_1 + (1 - \omega) P_{x_2x_2}^{-1} \hat{x}_2$$

where $\omega \in [0, 1]$. The free parameter ω manipulates the weights which are assigned to x_1 and x_2 . Different choices can be used to optimize the update with respect to different performance criteria.

The resulted estimation is used to update initial set of aggregate queries and to embed the necessary changes to the initially received queries. It is important to note that current approach can be easily extended to the information concentration management, as it was mentioned above. There, the updated estimate will be used to track the allocation for the certain information distributed in the network and stick it to the set of aggregate queries that are targeted for that information. Thus, a dual side optimization is possible, from the query and from the information location, concurrently.

Current approach is defined as glue for the intelligent workload management and query distribution. Since access analysis behaviour is different from query distribution and triggering, some assistant four step analysis can be defined. Query and inserts analysis consists of inserts or query issue, respond, match and actual access decomposition. Currently, workload management [8] is able to create workload models generated by actual accesses. Three other steps can be considered separately. As was defined above current approach treats any nodes activities as an abstract information gain vector and feeds with that fused entity. As a result query distribution and subscription allocation can be optimised and updated due to the actual network and data specific domain status.

4 Discussion and Conclusions

The presented approach provides robust scalable mechanism which is device energy efficient and efficiently utilises different computing platform.

An implementation of this has been made and is currently implemented in our Smart Space infrastructure which utilises Semantic Web principles [13, 14]. Use of this distribution allows us to automatically expand the search space and assists in controlling the privacy and sharing of the user's information implicitly.

The effect upon the performance of the system as a whole is negligible with network latency over mobile networks and even fixed networks being more of a problem. Such problems can be solved by selective caching and more advanced information representation mechanisms - these are the focus of our current research.

5 References

- [1] Simon J. Julier, Jeffrey K. Uhlmann, A non-divergent estimation algorithm in the presence of unknown correlations, 1997
- [2] D. Franken, A. Hupper: Improved fast covariance intersection for distributed data fusion. IEEE&ISIF Intern. Conf. on Information Fusion (FUSION), 7 pages, Philadelphia, USA, July 2005.
- [3] L. Huston, R. Sukthakar, R. Wickremesinghe, M. Satyanarayanan, G. Ganger, E. Riedel, A. Ailamaki. Diamond: A Storage Architecture for Early Discard in Interactive Search. Proceedings of USENIX Conference on File and Storage Technologies, 2004.
- [4] G. Bradski, A. Kaehler, Learning OpenCV: Computer Vision with the OpenCV Library, O'Reilly Media, 2008.
- [5] Sergey Boldyrev, Eugeny Linky, Distributed Objects Allocation/Retrieval System for Heterogeneous P2P Network, Nokia Research Centre, SUAI, Finnish-Russian university co-operation program in telecommunication, 2008
- [6] J.T. Tou and R.C. Gonzalez, Pattern Recognition Principles, Addison-Wesley, Reading, MA, 1974
- [7] C. Y. Chong, S. Mori, Graphical models for nonlinear distributed estimation, Proc. Of ISIF, pp. 614-621, Jul. 2004.
- [8] S. Boldyrev, S. Balandin, Illustration of the Intelligent Workload Balancing Principle in Distributed Data Storage Systems. Proceedings of workshop program of the 10th International Conference on Ubiquitous Computing, September 2008
- [9] A. Lappetelainen, Juha-Matti Tuopola, Arto Palin, Timo Eriksson, Networked systems, services and information, The ultimate digital convergence, 1st International NoTA conference, Helsinki, Finland, 2008
- [10] Niki Trigoni, Yong Yao, Alan Demers, Johannes Gehrke, and Rajmohan Rajaraman. Multi-Query Optimization for Sensor Networks. International Conference on Distributed Computing in Sensor Systems (DCOSS 2005).
- [11] Nicholson D., Lloyd C.M., Julier S.J., Uhlmann, J.K. Scalable distributed data fusion, Information Fusion, 2002. Proceedings of the Fifth International Conference, 2002.

- [12] Reto Krummenacher, The Smartest Space of All: A Global Space of (Machine-Understandable) Knowledge, 1st Russian Conference on SmartSpaces, ruSmart 2008, St.Petersburg, Russia, September 2008
- [13] Ian Oliver and Jukka Honkola, Personal Semantic Web Through A Space Based Computing Environment, Middleware for Semantic Web 08 at ICSC'08, Santa Clara, CA, USA, August 2008
- [14] Ian Oliver and Jukka Honkola and Jurgen Ziegler, Dynamic, Localised Space Based Semantic Webs, WWW/Internet Conference, Freiburg, Germany, October 2008
- [15] Lyndon J. B. Nixon and Elena Simperl and Reto Krummenacher and Francisco Martin-Recuerda, Tuplespace-based computing for the Semantic Web: a survey of the state-of-the-art, The Knowledge Engineering Review, June 2008, 23:2 pp181-212
- [16] Tim Berners-Lee and J Hendler and O Lassila, The Semantic Web, Scientific American, May 2001, vol 284:5 pp34-43