CIGMA : aCtive Inventory service in Global e-Market

Korea Advanced Institute of Science and Technology Su Myeon Kim
Seungwoo Kang · Heung-Kyu Lee · Junehwa Song

1. Introduction

A fully connected Internet business environment will introduce a high level of dynamics to business process. Business processes can be promptly summarized and reported, and decisions can instantaneously be reflected. By effectively adapting to such a dynamic environment, merchants can increase their profits. However, such situation may result in very frequent changes in business decision. Thus, information of sales items may undergo constant changes. Unfortunately, to ordinary customers, this introduces a high degree of inconvenience in online shopping. An example can be seen from the CISCO's case. By effectively adapting their business processes to the Internet environment, CISCO has shorten the term of their business closing to every 6 hours. Currently, they are trying to further shorten it to three hours. This means that they can estimate the accurate cost of their products, and hence change the sales conditions, every three hours. However, supposing that the sales conditions of popular goods are changed several times a day, customers may not be sure whether the conditions they saw at the merchant site ten minutes ago are correct. Thus, they can fall in a situation in which making shopping decision becomes very difficult.

Such inconvenience in online shopping is compounded by a flood of similar shopping sites on the Internet. For example, there are 7513 registered sites in the Yahoo booksellers' directory[1]. After narrowing the scope to the children's booksellers, there are still 298 sites. Online customers usually think that there are certainly other shopping sites with better sales conditions than the ones they have chosen and are not sure which site is the best one. This discomfort is much more serious than that of offline shopping due to the increased expectation to online shopping. In the conventional commerce environment, customers readily exclude visiting far-off shops considering the limits of time and distance. However, most online shopping sites have nearly the same accessibility. To make matters worse, merchants contend with each other. Online merchants compete to provide better sales conditions than the ones provided by the others. In the end, such competition may cause a chain reaction among similar shopping sites within a short time[2]. After all, online shopping under highly dynamic business environment can be described in a sentence: "NO ONE KNOWS HOW TO PURCHASE GOODS UNDER THE BEST CONDITION."

To overcome such inconvenience in online shopping, customers need proper support. However, existing service models or systems cannot effectively reflect such a highly dynamic environment and support ordinary online cus-

1) From the Yahoo Web site(www.yahoo.com) in Mar. 14th, 2002

2) Price War is the term used in economics to indicate such a chain reaction in online prices[7].

tomers. This paper describes a new EC service called "the aCtive Inventory service for Global e-Market"(CIGMA). The CIGMA provides catalog comparison services over multiple shopping sites for ordinary online customers. In addition, the CIGMA provides a one-stop shopping opportunity for online customers by mediating purchase transactions between the customers and merchants. A key characteristic of the CIGMA is that the services is provided, based on up-to-date information by reflecting frequent changes in catalog information in real-time. In addition, the service matches the desire of the online customers for the fast response since the service is provided based on the data cached in a high performance caching system. This paper overviews the system architecture and implementation of a working prototype system. Discussion on technical details can be found in [9].

Price comparison services [1, 2, 3, 4] can partly address the problems arising from the situation where too many similar shopping sites co-exist. It gathers price information for an item from many shopping sites. Then, it provides customers with the comparison information. However, it can not guarantee the correctness of the comparison information because it generally updates the gathered price information periodically. So, the comparison information may not be up-to-date at the moment of access. In the end, customers need to visit the original sites to check the validity of the given prices. There also exists an approach based on instant gathering of price data [5]. This approach may provide up-to-date price comparison information. However, it may incur a long delay which most online customers can not endure.

This paper is organized as follows. In Section 2, the challenges of online shopping are illustrated by a motivating scenario. In Section 3, the design of the CIGMA is described. Section 4 discusses related work in two aspects: existing e-commerce services and cache consistency mechanisms in dynamic data caching. We conclude this paper in Section 5.

2. Motivating Scenario

Consider two customers, Tom and Mary, who decide to buy a book, Harry Potter and the Condition of Sales at online bookstores.

There are numerous online bookstores on the Internet. Most of them sell the book since it is very popular. For instance, Amazon, Barnes&Nobles and AllDirect[3] sell the book at the price of $16.30, $15.55 and $15.20, respectively. Suppose that at this time, Amazon decides to lower its price to $14.50 to increase its market share. In a few minutes, Barnes&Nobles also lowers its price to $14.29 in accordance with its business policy: Lower price than Amazon, Better service than all others. These events also force AllDirect to similarly change its selling price. In the end, this price change and competition cause a chain reaction in the online book market.

Now, one of the customers, Tom does not know much about e-commerce; he has only heard about the advantages of online shopping several times from his friends. He visits only the two well-known online bookstores: Amazon and Barnes&Nobles. He finds out that Barnes&Nobles has proposed a lower price than Amazon and he decides to purchase the book at Barnes&Nobles. At that time, he realizes that there are a lot of other bookstores on the Internet and he begins to worry about whether or not his purchase at Barnes&Nobles is a good deal. He decides to visit more bookstores but realizing that there are so many, he randomly chooses several more bookstores and compares their prices for the book. In the end, he purchases the book from one of the

3) The bookstores Web sites are real. However, the characteristics of the sites mentioned in this paper are irrelevant to the actual characteristics of the sites. We use their names because they are popular. Other site names in this paper are used for the same purpose.

bookstores he has visited recently. However, he feels that it is very uncomfortable to do shopping on the Internet. He feels that there must be better a sales conditions elsewhere and that he can find it if he has more time.

Mary, the other customer, as an experienced online shopper, knows that there are too many bookstores to visit and thus, visits MySimon, a popular Web site with price comparison service for well-known bookstores. After several mouse clicks, she reaches the page showing a comparison table about the book and decides to buy it at AllDirect because of the lowest price. But at that moment, she notices a warning at the bottom of the Web page which states: Please note: PRICE shown here may NOT reflect the CURRENT price at the particular store. Be sure to CHECK both the price and the availability when viewing the product at the store. After reading this, she visits the AllDirect Web site and finds out that the book is already sold out, and she is compelled to visit each bookstore listed for comparison at MySimon to verify the price and the availability. In the end, she is able to purchase the book at the lowest price but only after a long series of manual Web search. Both Tom and Mary need to undergo an exhaustive Web search that may be impractical and at many times, useless.

By using the CIGMA service, both Tom and Mary can do a shopping easily and conveniently. Once they visit the CIGMA Web site and request the comparison information about the book they want to buy, they can immediately find the book with the best sales conditions through the comparison information which the CIGMA displays.

In addition, through the CIGMA's purchase mediation service they purchase the book at the CIGMA Web site without having to access the bookstore's Web site. The remarkable difference between the CIGMA and the conventional price comparison service can be described in an advertising sign at the top of the CIGMA Web page: Please note: We GUARANTEE the FRESHNESS of catalogs and comparison information which is shown here. NEVER waste time visiting each shopping site to VERIFY the provided information and to BUY goods.

3. System Design

The CIGMA system has a modular structure. It consists of five server components and one remote component. The remote one, called Merchant Wrapper (MW), runs on each merchant server. The server components are Client Manager (CM), Merchant Manager (MM), Transaction Manager (TM), System & DB Manager (SDBM) and Catalog Cache DB (CCDB). Figure 1 shows the CIGMA system along with a merchant and a customer.

3.1 Internal Service Flows

There are three kinds of external service in the CIGMA: catalog browsing, item purchase, and catalog update. The first two are for customers and the last, for merchants with each requests being handled differently. The internal service flows for these requests are shown in Figure 1.

The catalog browse request is the most frequent one. It is first received by the Web server, which is a part of the CM. The CM then parses the request and constructs an equivalent DB query string. It sends the query string to the CCDB to retrieve the requested information such as a catalog comparison table, an entire catalog, an item category list, etc. Lastly, it replies to the request with an HTML page which is dynamically generated with the query result.

To handle a purchase request, a safety mechanism is required because it generally includes important information such as a credit card number, address, phone number, etc. To secure the purchase transaction processing, the CM communicates with customers using the secure HTTP protocol. Then, it forwards the received

data to the TM. The TM also forwards the data to a proper merchant server at once. The result of the transaction is delivered from the merchant server to the customer in the opposite direction.

In case of an update request, immediate processing is important. Each modification of source data at the merchant's DBMS is detected by the MW module. Then, the module constructs and sends an update request message to the MM at once. The MM forwards the message to the SDBM after verifying the integrity of the message. Then, the SDBM actually updates the cached data of the CCDB by composing a DB query string based on the message. After the CCDB update, the return value of the update operation (i. e., OK or NOT_OK) is forwarded to the MW in the reverse order.

3.2 Major Components

Each component is designed as a separate process. For communication among the components, we use a message-passing-based mechanism. It is preferable in terms of scalability and flexibility because the system can be easily deployed in the multi-node environment without any modification and each component can be modified or upgraded independently as long as external interfaces are not changed. Providing scalability is very important in the CIGMA service because the CIGMA system has to handle a large number of online customers' requests and a number of updates from merchants. The following describe the function of each component.

3.2.1 Client Manager(CM)

The CM takes charge of every communication with customers. It includes a Web server and a Web application server. The Web application server is required to generate responses with dynamic contents[4]. Via these servers occurs each interaction with customers. The CM receives both catalog browse requests and purchase requests from customers. As mentioned before, the secure HTTP is used to handle purchase requests with security.

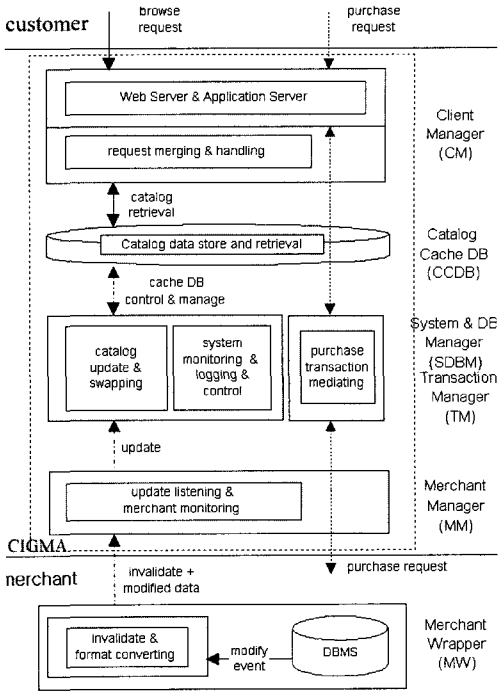

Figure 1 Service flows and major functions of the CIGMA

The CM also has the capability to merge identical requests from different customers. The CM inspects the request queue and merges the same requests from different customers into a single internal request. This improves the system performance when a lot of customers simultaneously send the same request for a popular item.

3.2.2 Merchant Manager(MM)

The MM manages most interactions with merchants except those required for purchase transactions. Upon a catalog modification, the MM receives an invalidation message from a merchant server. Then, the MM immediately notifies the SDBM of the event to request the update of the CCDB.

4) Popular server applications are constructed by using JAVA Servlet, Java Server Page (JSP), Active Server Page (ASP), etc.

The merchant server monitoring should be noted. There could be a certain situation such as network partitioning, merchant server failure and network congestion, in which the MM cannot receive invalidation messages. This may result in incorrect catalog information. To avoid this problem, the MM monitors heart beat messages from merchants to check the status of servers. If a merchant sends no messages during a predefined period, the MM notifies the SDBM of the situation. Then, the SDBM will initiate a suitable action. For example, the SDBM may invalidate all the catalogs from the unreachable merchant server until the server responds again. A merchant is forced to send an empty message at the end of the predefined interval if there is no update.

3.2.3 Catalog Cache DB(CCDB)

The CCDB manages storing and retrieval of cached catalog data. Generally, the performance of DBMS is very critical to the overall performance of an e-commerce site. The CIGMA is expected to process a much higher number of update requests as well as customer requests than general shopping sites. Therefore, the performance of DBMS is even more critical in the CIGMA.

We use main memory as the primary data storage to process a high rate of requests. Disks are used as a supporting medium to compensate the shortage of main memory space. We argue that persistency is less critical in the CIGMA since it deals with only the cached data which has replication in the original merchant servers. In addition, it has been reported that, in disk based caching systems, the overhead of disk accesses is a significant portion of overall load. Thus, in a busy system like the CIGMA with frequent requests and updates, the overhead of keeping data in disks will be even more significant. Our system trades off the lack of persistency wit h such an overhead.

3.2.4 System & DB Manager(SDBM)

The SDBM is the most important and complex component in the CIGMA system. The SDBM has three main functionalities: updating cache data, monitoring and controlling the system, and swapping catalog.

The SDBM takes charge of updating cache data. Upon update requests from MM, the SDBM composes a proper query string based on the incoming request. Then, it updates cache data by sending the query to CCDB. After the update, the return value of the update operation (i.e., OK or NOT_OK) is forwarded to the MW through the MM.

The SDBM monitors the status of the system and controls system components. For example, the SDBM monitors the frequency of the requests from customers and merchants and collects its statistics. Based on the collected data, it controls major facilities of the components to improve the performance.

3.2.5 Transaction Manager(TM)

The TM mediates purchase transactions. This is for providing online customers with a one-stop shopping opportunity. The CIGMA does not process the transaction on behalf of merchants, but simply relays all the purchase-related information, i.e., order form, payment information, buyer information, and transaction result, between a customer and a merchant. When a customer fills out and submits the form, the TM receives the purchase-related information through the CM and forwards it to the merchant server related to that purchase transaction. This approach enables the CIGMA to avoid the complication and overhead which may be incurred by related business issues.

Figure 2 shows the basic procedure for the purchase mediation. Upon a customer's purchase request for an item (1), the CIGMA responds by sending the proper order form for the item (2). For the purchase, the customer simply fills out the delivered form (3). Then, upon receiving the purchase-related data from the customer (4), the

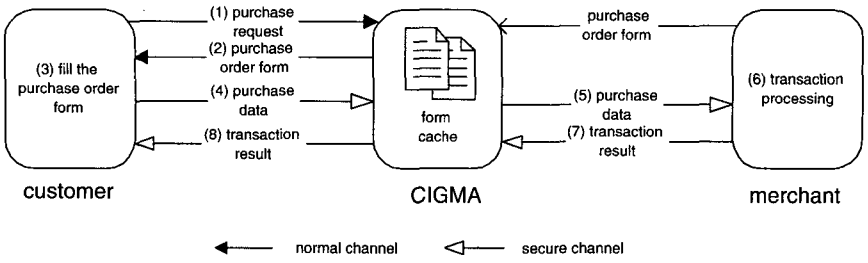

Figure 2 Message flow for Purchase Mediation

CIGMA forwards that data to the original merchant (5). After processing the transaction (6), the original server sends the transaction result to the CIGMA (7). Lastly, the CIGMA forwards the result to the customer (8).

3.2.6 Merchant Wrapper(MW)

The MW covers all the required works for a merchant to interact with the CIGMA. As shown in Figure 1, the MW is deployed on a merchant server for the merchant to join the CIGMA service. Two main functions of the MW are the update propagation and the heart beat message transmission.

The update propagation is performed according to the real-time update scheme. The MW also periodically sends out heart beat messages to the CIGMA, if there is no update event during a predefined interval. Although the heart beat message carries no information, the message itself is important because it notifies the CIGMA of the fact that the merchant server is alive. If there are neither update messages nor heart beat messages, the CIGMA regards the merchant as having some critical problems. Thus, the CIGMA can invalidate all catalog data from the merchant to avoid providing catalog data which may be stale.

3.3 Service Example

We have built two sample shopping sites: Total Shop and Fragrance Online[5]. Total Shop has a plan to become a general shopping mall and is trying to handle various items. Currently, Total Shop carries three kinds of items: MP3 players, perfumes, and computers. Contrary to Total Shop, Fragrance Online specializes only in fragrances. Both shopping sites have joined the CIGMA service to increase their business profit. Two shopping sites have already specified the category classification and catalog conversion information for each catalog and made a CMT[6] (Catalog Mapping Table), respectively. Then, they have installed the MW modules and DB triggers.

The CIGMA is configured to handle various kinds of items. The CIGMA has its own catalog format as it has already contracted with both shopping sites. The catalogs of the two shopping sites are automatically reclassified and converted according to the CMT. In addition, although it is difficult to show here how the real-time update occurs, the catalog data of the CIGMA is updated immediately whenever the original data in the merchant DB is modified.

4. Related Work

As described before, price comparison services [1, 2, 3, 4] is similar to our work in terms of providing comparison information. However, they

5) You can view the CIGMA and the two sample sites. The URL of the CIGMA is http://nc9.kaist.ac.kr/cigma/. The two sample merchant sites can be accessed via http://nc2.kaist.ac.kr/merchant/ and http://nc2.kaist.ac.kr/merchant2/

6) The CMT is a table which contains descriptions for catalog conversion. The details can be found in [9].

do not guarantee the accuracy of the provided information because they update the data manually or periodically. Thus, these services can not effectively support customers to shop online under a dynamic e-commerce environment.

AddAll [5] provides price comparison service based on instant catalog searching and gathering upon customer's requests. Thus, it can provide fresh information about an item. However, it has a serious problem: the response time may be very long because it has to visit many shopping sites in an on-demand fashion. In addition, it issues multiple Web requests per each customer request and thus may cause a heavy traffic on the Internet.

Recently, several researches have proposed techniques for dynamic data caching [12, 10, 16, 8, 11, 17, 6, 14]. Although consistency mechanism has not been described explicitly in [16, 15, 17, 6], cache consistency is an important issue in dynamic data caching. Generally, approaches for maintaining cache consistency can be classified as client-pull and server-push approaches.

approaches: In client-pull approaches, cache is responsible for maintaining the consistency. Cache checks a data modification in a server to revalidate or update cached data without the knowledge of update. Most current Web proxy caches use the If-Modified-Since request header to revalidate cached objects [13]. This method may incur much unnecessary traffic during the revalidation. The mechanism proposed in [12] maintains the consistency of a cached content by setting a TTL field for different CGIs. This method is acceptable because the paper focuses on the digital library applications, where most contents is not changed frequently.

Server-push approaches: In server-push approaches, a server notifies caches of a modification for the invalidation or update of the cached data. Server-push approaches can relatively well maintain the consistency. However, they incur heavy load on a server in the case of Web proxy caching because there could be a large number of proxy caches that a server should notify of a modification. Thus, server-push approaches are usually used in caching for a specific server: reverse proxy caching [10, 8, 18] or database caching [14]. In [10, 8, 18], dynamically generated HTML pages or fragments are cached. These researches focus on how to identify the dependency between the modified underlying data and cached objects. Invalidator module proposed in [8] uses DB log files or periodic polling to detect a modification of data in a database and invalidate related objects. DBCache proposed in [14] uses a tool for asynchronous data replication. In this approach, an administrator subscribes replication requests or specifies the frequency of update propagation. Contrary to these two approaches, a modification of data can be detected automatically and there is no unnecessary overhead for periodic polling in our real-time update scheme. An algorithm for update propagation using a dependency graph was proposed in [10]. It is similar to our system that it uses the trigger mechanism of a DBMS. However, as mentioned above, that approach is used in the HTML fragments caching for a specific server.

5. Conclusions

E-commerce (EC) has already become a major facilitator of business. It will become more important as the Internet grows rapidly. However, despite many advantages of the EC, online shopping is an overloaded activity for many customers due to a highly dynamic e-commerce environment.

We have proposed the CIGMA service which can effectively support customers to do online shopping under a highly dynamic e-commerce environment. The CIGMA suggests the best sales conditions over multiple shopping sites and provides a convenient shopping environment for customers. Customers can greatly save shopping overheads including time and money by using the CIGMA service. We have presented the architecture and implementation of a working

prototype system. We have also shown the validity of the system via a service example.

Acknowledgement

This work was partly supported by the Korea Science and Engineering Foundation(KOSEF) through the Advanced Information Technology Research Center(AITrc).

References

[1] http://www.nextag.com/, NexTag－The Search Engine for Shoppers.

[2] http://www.mysimon.com/, mySimon－Compare products and prices from around the Web.

[3] http://www.dealtime.com/, DealTime－compare products, prices & stores.

[4] http://www.pricescan.com/, PriceSCAN－Com parison Shop for the Lowest Prices.

[5] http://www.addall.com/, AddALL - Book Search and Price Comparison.

[6] A. Brown, G. Kar, and A. Keller, An active approach to characterizing dynamic dependencies for problem determination in a distributed environment, In Proc. of IFIP/IEEE International Symposium on Integrated Network Management, pages 377 - 390, May 2001.

[7] Luis M. B. Cabral, Introduction to Industrial Organization, 1st edition, MIT Press, 2001.

[8] K. S. Candan, Wen Syan Li, Qiong Luo, Wang Pin Hsiung, D. Agrawal, Enabling Dynamic Content Caching for Database-Driven Web Sites, In Proceedings of SIGMOD 2001, May 2001.

[9] Su Myeon Kim, Seungwoo Kang, Heung-Kyu Lee, and Junehwa Song. CIGMA: aCtive Inventory service in Global e-MArket for enabling one-stop shopping over internet shopping sites. Technical Report CS-TR-2002-178, Korea Advanced Institute of Science and Technology, 2002.

[10] Jim Challenger, Arun Iyengar, Paul Dantzig, A Scalable System for Consistently Caching Dynamic Web Data, In Proceedings of INFOCOM 99, pages 294-303, 1999.

[11] Jim Challenger, Arun Iyengar, Karen Witting, Cameron Ferstat, Paul Reed, A publishing system for efficiently creating dynamic web content, In Proc. of INFOCOM00, pages 844-853, 2000.

[12] Vegard Holmedahl, Ben Smith, and Tao Yang, Cooperative caching of dynamic content on a distributed web server, In Proc. of the 7th IEEE International Symposium on High Performance Distributed Computing, pages 243-250, 1998.

[13] Balachaner Krishnamurthy and Jennifer Rexford, Web Protocols and Practice - HTTP/1.1, Networking Protocols, Caching, and Traffic Measurement, Addison-Wesley, 2001.

[14] Qiong Luo, Sailesh Krishnamurthy, C. Mohan, Hamid Pirahesh, Honguk Woo, Brude G. Lindsay, and Jeffrey F. Naughton, Middle-tier database caching for e-business. In Proceedings of SIGMOD02, June 2002.

[15] Qiong Luo and Jeffrey F. Naughton, Form-based proxy caching for database-backed web sites, In Proc. of the Int. Conf. on VLDB, September 2001.

[16] Qiong Luo, Jeffrey F. Naughton, Rajasekar Krishnamurthy, Pei Cao, and Yunrui Li, Active Query Caching for Database Web Servers, In Proc. of WebDB'2000, pages 32-35, 2000.

[17] Khaled Yagoub, Daniela Florescu, Cezar Cristian Andrei, and Valerie Issarny, Building and Customizing Data-Intensive Web Sites Using Weave, In Proc. of the Int. Conf. on VLDB, September 2000.

[18] Huican Zhu and Tao Yang, Class-based cache management for dynamic web content, In Proceedings of INFOCOM01, April 2001.

Su Myeon Kim

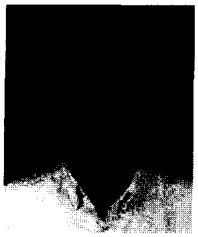

1995 Yonsei University, Mathematics, B.S.
1997 Korea Advanced Institute of Science and Technology, EECS, M.S.
1997 Korea Advanced Institute of Science and Technology, EECS, Ph.D. course
E-mail:shkim@casatum.kaist.ac.kr

Seungwoo Kang

2000 Yonsei University, Electronic Engineering, B.S.
2002 Korea Advanced Institute of Science and Technology, EECS, M.S.
2002 Korea Advanced Institute of Science and Technology, EECS, Ph.D. course
E-mail:swkang@nclab.kaist.ac.kr

Heung-Kyu Lee

1978 Seoul National University, Electronic Engineering, B.S.
1981 Korea Advanced Institute of Science and Technology, Computer Science, M.S.
1984 Korea Advanced Institute of Science and Technology, Computer Science, Ph.D.
1984~1985 University of Michigan, Ann Arbor. Post Doc.
1986~ Korea Advanced Institute of Science and Technology, EECS, Professor
E-mail:hklee@casatum.kaist.ac.kr

Junehwa Song

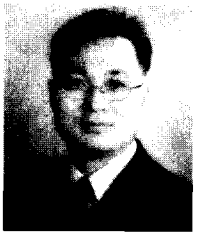

1988 Seoul National University, Computer Sc and Statistics, B.S.
1990 State University of NY, Computer Science, M.S.
1997 University of Maryland, Computer Science, Ph.D.
1994~1996 IBM T.J. Watson Research Center, S/W Engineer, Digital Library Group
1996~1997 IBM T.J. Watson Research Center, S/W Engineer, Digital Media Solutions Dept.
1997~2000 IBM T.J. Watson Research Center, Research Staff Member, Parallel Commercial Systems Dept.
2000~ Korea Advanced Institute of Science and Technology, EECS, Assistant Professor
E-mail:junesong@nclab.kaist.ac.kr