

# 네트워크 관리에서의 사건상관기술

정회원 박용석\*

## Event Correlation Technologies in Network Management

Yongseok Park\* *Regular Member*

### 요 약

네트워크관리는 오늘날 통신 및 기업네트워크의 복잡하고 엄격한 성능과 유용성을 극복하는 방법으로 최근에 많은 주목을 받고 있다. 네트워크에서 발생하는 사건들을 통합하는 사건상관은 이 문제점들을 해결하는데 가장 효과적이고 필요한 방법으로 여겨지고 있다. 이 분야에 대한 연구가 10년도 넘게 진행되었지만 현재까지 개발된 방법들에 대해 총괄적으로 다룬 논문을 찾기가 어렵다. 이 논문에서는 사건상관이 왜 필요한지를 보이며, 여러 논문들과 상업제품에서 발견되는 방법들을 소개하며, 그리고 마지막으로, 해결되어야 할 문제점들을 고려한다.

### ABSTRACT

Network Management has recently gained much attention as a vehicle to overcome the increasing complexity and stringent performance/availability requirement of today's telecommunications and enterprise networks. Event correlation, which refers to consolidation of collected events occurring in networks, is considered as one of the most effective and desired functions to meet the requirements. There have been a lot of research on this subject for more than a decade, but it is difficult to find a literature that provides comprehensive survey on the existing state-of-the-art technologies. In this paper, we show why event correlation is necessary, illustrate the existing approaches found in literature and commercial products, and finally, consider the challenges that need to be overcome.

### I. 서론

Today's computer and communication networks are characterized by complexity and fast pace of evolution. Network Management, which refers to overall activities of monitoring and controlling network resources, has become a vital task in providing reliable services over the networks. Outages in networks cause not only enormous financial damages but also potential loss of customers to service providers. However, latest outages of AT&T's frame relay service<sup>[1]</sup>, and MCI Worldcom's Internet service<sup>[2]</sup>, show that it takes a lot of time not only to recover but also

to find the root cause with current network management environment.

Network management systems collect and receive many events occurring in the networks. The events include traps out of the polls on the states of network elements, alarms, configuration or provisioning change notifications, state change notifications, and user/operator actions. It is alarms, among the events, that requires most urgent treatment, because it is the most service-affecting.

It has been common that network management

\* Lucent Technologies, Holmdel, New Jersey, USA, Optical Networking Division (y.park@ieee.org)

논문번호 : 99343-0823, 접수일자 : 1999년 8월 23일

※ This work has been done when the author was with AT&T Labs.

operators are overwhelmed by a lot of alarms in a very short time period (event storm/flood) from all over the networks, although the alarms were caused by a single failure of network equipment. For example, suppose that a T3 interface card has a bad clock (this scenario is excerpted from [3]). The intermittent loss of synchronization due to the bad clock results in bit errors in T3 signals. This problem in the physical layer results in the frame errors in the link layer, packet losses in the network layer, and retransmission at transport layer. Evidently, the applications using the network undergoes the network slowing down.

The most time-consuming part of network management is in identifying events and drawing conclusions on many events occurring in the networks. That's because events occur as a result of very complex interaction of many physical and logical components of networks, and the overall system is not well understood (see [4] for general problem of fault management). That's why many artificial intelligence techniques have been tried and applied so far (will be evident in Section II).

Event correlation is a process that consolidates events by correlating events based on causal relation between events or by referring to the knowledge base that keeps information about the system (see Figure 1 for general architecture of network management). Event correlation removes redundant events, detects event patterns, and isolates faults. Event correlation is different from event filtering in that the former does not use the information except the event itself but latter utilizes other knowledge on the system. One typical application of event correlation is fault diagnosis or root-cause analysis. Other applications of event correlation include personal news, personal location tracker, and stock-change services.

Thanks to the advance of computer technologies, many aspects of network management have been automated, but still much portion is being manually done. For example, typical fault




Fig. 1 Architecture of Network Management Systems.

management process involves many independent operations support systems and intervention by human operators. Figure 2. shows a typical scenario of fault isolation in international data services using both Time-Division Multiplexing (TDM) and Asynchronous Transfer Mode (ATM) over Synchronous Optical Networks (SONET). It is Customer Care (CC) that gets customer's complaints when network slows down. It is also possible that some proactive monitoring systems or operators notice problems and report. The first thing that CC does is to check whether the problem has been caused domestically, by running circuit loopback tests. If the problem turns out to belong to domestic area, then the trouble information is transferred to TDM facility maintenance center (FMC). The operators at FMC use several network management systems and try to find a root cause of the problem. If the problem turns out ATM or SONET problem, then the trouble information is transferred to ATM or SONET maintenance centers accordingly, and resolved there.

If event correlation system is deployed and operating efficiently, we can achieve the dramatic improvement in network availability as follows:

- 1) avoid overloading operators with events,
- 2) reduce handoffs among operations centers/operators/systems/windows,
- 3) automate human expert's tasks of fault detection and diagnosis, and
- 4) provides scalability when future expansion of networks is expected.

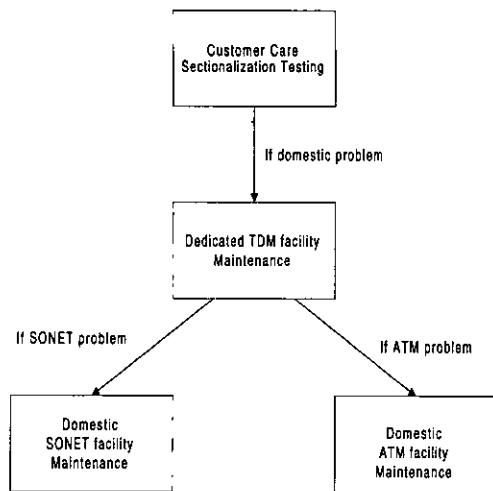


Fig. 2 Handoffs at Network Operations Centers.

The organization of the rest of the paper is as follows. In next section, we present and discuss the current state-of-the-art approaches to event correlation. In Section III, we consider the challenges to overcome for event correlation to be successful. In Conclusions, we consider future research directions.

## II. Event Correlation Approaches

The approaches to event correlation can be classified based on 1) how to acquire knowledge, 2) how to represent the knowledge, and 3) how to reason a set of events against the acquired knowledge base.

### 1. Rule Based Expert Systems

Rule-based Reasoning (RBR) or Expert Systems is assuming the existence of human experts in the problem domain. The rules are acquired by interviewing the experts, and then represented with pre-defined computer languages such as Prolog and Lisp. As an example, CardUp/CardDown rule for a network equipment can be such as: "Post a critical alarm when the card down event is received, clear the cardDown event when the cardUp event is received. When the cardDown is received, filter out all subsequent events from physical ports and logical channels

contained in the card."

There has been extensive research in this area with applications to planning and diagnosis. RBR is easy to implement as long as human experts are available, which becomes a problem in Today's ever changing networking technologies and short life time of equipments. Other well-known problems with RBR are knowledge acquisition bottleneck, brittleness, slow reasoning speed, scalability problem, and weakness to noise. With all these limitations, Expert Systems have still been the most popular approach in the event correlation implementations due to its simplicity (see [5] for detailed overview). Alarm Correlation Engine from GTE is a rule-based system for a telephony switch<sup>[6]</sup>.

### 2. Structural Model Based Approach

Structural models for network equipments are in some cases available as physical or logical objects at Management Information Base (MIB) definition. Connection information is also available at configuration files of associated network management systems. The mapping information of customers to network connections is typically available at Network Service Provider's Operations Support Systems (OSS). With all these information about the network topology and equipments, It is natural to exploit them for event correlation. Well-designed network model and configuration data can be very effective for event correlation. IMPACT from GTE uses network element class hierarchy with connect, within, and contain filters<sup>[7]</sup>. Bouloutas et al. uses a network structure modeled by Phase Structure Model, apriori probabilities of faults, and observed alarms to compute the most probable faults<sup>[8]</sup>. Katzela and Schwartz uses belief network, which is a dependency graph with weight of arcs representing belief, for alarm correlation<sup>[9]</sup>. Gruschke describes an event correlation prototype for IP connectivity service, using dependency graph, which is built on MIB-II information and a list of hosts<sup>[10]</sup>. Meira and Nogueira proposed a layered model

consisting of user network, management network, Intelligent Networks, signalling, cellular, transport, local switching, and access infrastructure, where each layer offers services to other layers<sup>[11]</sup>.

### 3. Coding Approach

Messages related to network management typically have lower priority than those related to data. For example, most SNMP messages are delivered over UDP (User Datagram Protocol), which is unreliable in nature. Therefore, it is possible that events are lost, delayed, and spuriously generated. The event correlation approaches that are rule-based or finite-state machines (will be discussed later in Section II.7) based are very sensitive to this noise in the events related to network management.

InCharge employs the codebook approach to solve this problem<sup>[3]</sup>. The codebook refers to a table that relates lists of alarms (codes) and their root-causes. The procedure for creating the codebook is as follows. They first create a matrix that represents the relationship between fault (problem) events and alarms (symptoms). Then, they create the causal graph that represents the relationship between alarms using the object-oriented structural model of the networks. Finally, they remove redundancy using the causal graph, and selects the codes that are separated far enough from each other.

Figure 3 shows an example of the relation between problems and symptoms. The code [1 1 0 0], which represents that observation of both




Fig. 3 Relationship Between Problem Events and Symptom Events.

symptoms a and b uniquely points out that the root-cause of the symptoms is problem A.

InCharge provides robustness to event correlation, but it also has shortcomings. First, its codebook should be recomputed whenever there is a change in the networks. Secondly, it is not good for correlating events with temporal relationship.

### 4. Case Based Approach

Case-based reasoning (CBR) stems from realization that human solves problem based on his experience rather than specific rules. Figure 4 shows an architecture for CBR.

CBR keeps the problem-solving cases into knowledge base using sophisticated indexing scheme. When a fault occurs, the CBR system retrieves a similar case from the knowledge base, and tries to apply to the current problem with adaptation. The solution is evaluated afterwards. If it was successful and worth as a new case, it is inserted to the knowledge base. Otherwise, the system learns that the retrieved case was not applicable to the current problem. Therefore, CBR overcomes the knowledge acquisition bottleneck and brittleness problems of the RBR. Lewis and Dreo suggests the use of fuzzy logic for correlation of low-level network events and high-level case description of the problem<sup>[12]</sup>. Their approach is implemented into SpectroRx from Cabletron<sup>[13]</sup>. The commercial CBR system complements the inductive reasoning module of the Spectrum network management system as an add-on.




Fig. 4 Architecture of Case-Based Reasoning.

### 5. Neural Network Approach

Neural Network (NN) is a pattern recognition method based on networks of nonlinear (mostly sigmoid) function nodes. NN provides robustness to pattern recognition due to smooth continuity property. NN has been successfully used for proactive performance management, for example, at Neugent from Computer Associates<sup>[14]</sup>.

Wirtgreffe et al. proposed using NN for event correlator<sup>[15]</sup>. Figure 5 shows an NN which relates faults and alarms. Trouble tickets and their resolutions are used to train the NN. A problem with NN approach is that it requires re-training whenever there is a network topology change.




Fig. 5 NN Architecture for event correlation.

### 6. Circuit Approach

Hewlett-Packard (HP) OpenView, which is the most widely used network management platform, did not have event correlation in its product line until it introduced Fault Management Platform (FMP)<sup>[16]</sup>. The FMP is based on HP OpenView Distributed Management Platform, and provided utilities for trouble ticketing, performance measurement, alarm handling, and network map. The front end of the FMP is Mediation Device Block, which has a module that correlates repeated, transient, and related alarms. Their correlation technique is based on the causal relationship between alarms. Figure 6 shows an example of causal relationship between events. Alarm a causes alarms b and c; alarm c in turn




Fig. 6 Example of Causal Relationships between Alarms.

causes alarm d.

For higher level of event correlation, HP provides Event Correlation Services (ECS) (see Figure 7 for architecture of ECS)<sup>[17]</sup>. The ECS consists of Correlation Circuit, Data/Fact Store, and Annotate External Server. The Correlation Circuit is an acyclic circuit with 14 function nodes such as source, sink, filter, delay, table, combine, create, and annotate. The annotate node sends request to external annotate server for a data, and receives response. Data Store contains information on parameters, and fact Store on relationships between events. They serve as internal databases for Correlation Circuit. The ECS also has capability of correlating events arriving at different time instances thanks to its delay node.

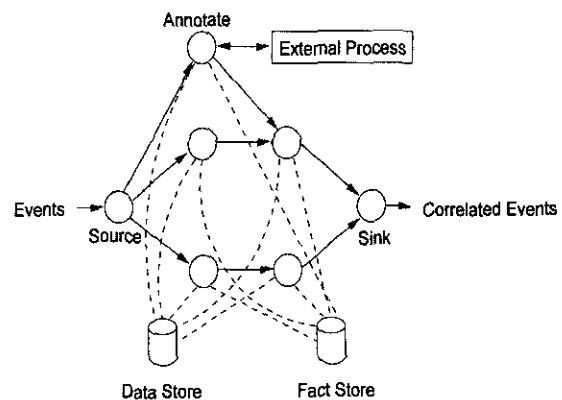


Fig. 7 Architecture of Event Correlation Service.

Although ECS is a very powerful event correlation tool, it has been reported that designing the ECS circuits is very tedious and time-consuming. Another problem with ECS is that it needs to review all the circuits whenever there is change in the network topology.

### 7. Behavioral Model Based

This approach provides the most comprehensive event correlation if a proper behavior model is available. NerveCenter Pro from Seagate Software is a probing and event correlation tool for network management<sup>[18]</sup>. It can be used as a stand-alone, but in most cases is used with other network management systems such as HP OpenView, Spectrum from Cabletron, and Netview from IBM. It is also integrated into Solstice Enterprise Manager from SUN.

NerveCenter polls for the attributes of managed objects, and receives event notifications from the agents of managed objects through requests. The requests are represented by finite-state machines (more rigorously Mealy machines). The requests can be designed using the pre-defined request templates. NerveCenter uses the requests for correlation of polls and event notifications, and issuance of alarms. Figure 8. shows an example of finite-state machine for the request that checks reachability of a machine (excerpted from [18]). Suppose there is no response from Simple Network Management Protocol (SNMP) agents on a machine. It is possible that either the agent or the machine is down. At this stage, by using ping, we can identify which is down.

NerveCenter provides a very sophisticated mechanism for detecting and identifying some faults, but the scenarios are very specific to a limited number of faults. Therefore, it is quite difficult to cover all the possible fault scenarios. It also has the same tedious design problem as ECS.

### 8. Constraint Based Approach

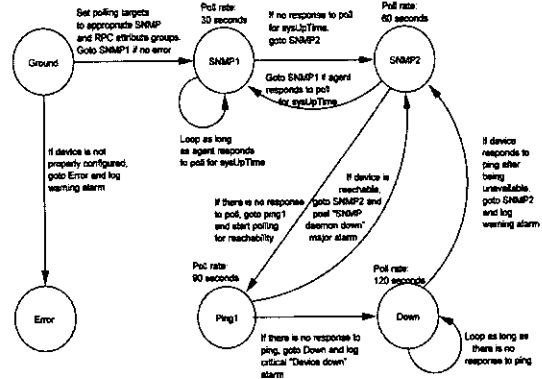


Fig. 8 State Diagram for Machine Reachability.

A significant portion of root-causes for network and service problems is due to configuration mistakes. Constraint based approach is a model based approach when we have 1) a comprehensive specification or constraint among model components, and 2) the components are observable. In this case, it is easy to detect faults as violation of constraints, which also becomes the root-cause. Sabin et al. presented ADNET, a constraint modeling environment for fault diagnosis, and demonstrated its application to diagnosing DNS (Domain Name Service) problems<sup>[19]</sup>. In [20], Sabin et al. extended the result of [19] by adding model scalability, proactive diagnosis, and better problem explanation capability.

### 9. Event Specification Approach

As many distributed systems emerge, there has been need to integrate and automate event notification among the distributed systems. The examples include business messaging, network and systems management, application management, on-line auction, calendar service, stock price change notification, personal news, and on-line sports broadcast systems. There has been some effort to develop general purpose event notification systems. Yeast by Krishnamurthy et al. was a client-server system in which clients register event-action specifications to a central server for event notification services<sup>[21]</sup>. READY by Gruber et al. is an event notification system being

developed at AT&T, which extends Yeast by adding event structures, quality of service directives, and session concepts<sup>[22]</sup>. The READY system is being integrated to Marvel, a Java based network management system, as an event correlation module<sup>[23]</sup>. Rules acquired from human experts or network model analysis can be represented as simple or compound event specifications of the READY system for alarm correlation. JECTOR by Liu et al. is an alarm correlation system based on an event specification language, which is used to write temporal relationships between events<sup>[24]</sup>.

### 10. Hybrid Approach

As explained so far, many approaches for alarm correlation have both advantage and disadvantage in most cases. The hybrid approach is the one that tries to complement the weakness of one approach by the strength of the other approach. G2 from Gensym uses both model and rule based approaches<sup>[25]</sup>. They use Specific Fault Propagation Model (SFPM), which is constructed by combining General Fault Propagation Model (GFPM) and Specific Domain Model (SDM). GFPM is built from human operator's expertise, and SDM is a local network model reflecting network entities and connectivity. The SFPM is recomputed when network topology changes, new rules are inserted, or old rules are removed. Grimes and Alley proposed using expert systems for advanced alarm filtering, neural networks or belief networks for alarm correlation, and case-based reasoning for alarm identification<sup>[26]</sup>.

## III. Challenges to Overcome

In this paper, we surveyed many state-of-the-art approaches to alarm correlation in telecommunication networks. It has been more than a decade's research on event correlation, but still the technology has not been widely deployed and used. The obstacles that should be overcome are as follows. First, there is an issue in data quality which alarm correlation systems use. The network

models are built using the configuration data in a database, which should be Database of Record (DBOR). However, the reality is that the DBOR is not 100% correct, because many data are still manually inserted, and proper validation is not in place. Therefore, it is common to find the databases for fault management systems and those of provisioning systems are out of synchronization. Alarm correlation using the unreliable information source may mislead so that lose its effectiveness. Secondly, the events may not represent real-time information themselves. Many events are generated as a result of polling, which is periodic typically with 5 to 15 minutes. That means the event may represent 5 to 15 minute old information at the worst case. Thirdly, it is difficult to find human experts in Today's fast changing network technologies era. Some networks such as SDH/SONET have been managed with a well defined object oriented model (see [27] for overview of TMN), but IP networks have not. It is a challenge how to complement the absence of experts and good models. Fourthly, as networks become complex, e.g., cable and DSL networks with millions of network elements and IP services, scalability becomes a problem. How to distribute events, event correlation functionalities, and information models is a big challenge. And lastly, there is time windowing issue for temporal correlation of alarms. If the window is too small, then we lose some of correlation capability such as pairwise cancellation of alarms. If timing window is too large, then a lot of computer resource is necessary, and the correlation result might be too late.

## IV. Conclusions

In this paper, we presented why event correlation is important in network management, and gave a comprehensive survey on existing state-of-the-art technologies for event correlation. We also considered several standing issues to be resolved for event correlation to play a major role in network management.

After reviewing the event correlation technologies, it seems that one technology is not enough for comprehensive event correlation. Therefore, it is necessary to employ the hybrid approach (see Section III.10 for detail.). Now the question is which technologies to select and how to distribute events to the systems that are based on the selected technologies. The answer should be based on the following factors. First, it should be studied whether information models are available from standards or vendors. Secondly, the human operators and tier supports who have been troubleshooting the networks should be interviewed for their expertise. Thirdly, event logs, which tend to be enormous amount of data, should be analyzed for event statistics and patterns (some works in this context are reported in [28]).

### References

- [1] Network World Inc., "AT&T's big fix," *Network World*, <http://www.nwfusion.com/news/1999/0322att.html>, March 22, 1999.
- [2] USA Today, "MCI WorldCom service restored, but many users still angry," *USA Today*, August 16, 1999.
- [3] S. A. Yemini, D. Kliger, E. Mozes, Y. Yemini, and D. Ohsie, "High speed and robust event correlation," *IEEE Communications Magazine*, pp.82-90, May, 1996.
- [4] A. Dupuy, J. Schwartz, Y. Yemini, G. barzilai, and A. Cahana, "Network Fault Management: A User's View," *Integrated Network Management I*, pp. 101-107, 1989.
- [5] R. N. Cronk, P. H. Callahan, L. Bernstein, "Rule-Based Expert Systems for Network Management and Operations: An Introduction," *IEEE Network*, pp. 7-21, Sep., 1988.
- [6] P. Wu, R. Bhatnagar, L. Epshtein, M. Bhnadaru, and Z. Shi, "Alarm Correlation Engine (ACE)," *Proceedings of the IEEE/IFIP 1998 Network Operations and Management Symposium (NOMS)*, New Orleans, LA, USA, pp. 733-742, Feb., 1998.
- [7] G. Jakobson and M. D. Weissman, "Alarm correlation," *IEEE Network*, Nov. 1993.
- [8] A. T. Bouloutas, S. Calo, and A. Finkel, "Alarm Correlation and Fault Identification in Communication Networks," *IEEE Trans. Communications*, vol. 42, no. 2/3/4, pp. 523-533, Feb./Mar./Apr., 1994.
- [9] I. Katzela and M. Schwartz, "Schemes for Fault Identification in Communication Networks," *IEEE/ACM Trans. Networking*, Dec. 1995.
- [10] B. Gruschke, "Integrated event management: event correlation using dependency graphs," *Proceedings of the 9th IFIP/IEEE International Workshop on Distributed Systems: Operations & Management (DSOM)*, pp. 130-141, Oct. 1998.
- [11] D. M. Meira and J. M. S. Nogueira, "Modelling a Telecommunication Network for Fault Management Applications," *Proceedings of the IEEE/IFIP 1998 Network Operations and Management Symposium (NOMS)*, New Orleans, LA, USA, pp. 723-732, Feb., 1998.
- [12] L. Lewis and G. Dreo, "Extending trouble ticket systems to fault diagnosis," *IEEE Network*, pp. 44-51, Nov. 1993.
- [13] Cabletron Systems, "SpectroRx Resolution Expert," <http://www.cabletron.com/products/items/SA-CSI1016>, 1998.
- [14] Computer Associates, "The Necessity of Neugents: Mastering Complex Systems," <http://www.cai.com/products/whitepapers/neugents>, 1999.
- [15] H. Wietgreffe, K. Tuchs, K. Jobmann, G. Carls, P. Frohlich, W. Nejd, and S. Steinfeld, "Using neural networks for alarm correlation in cellular phone networks," *Proceedings of the International Workshop on Applications of Neural Networks in Telecommunications*, 1997.
- [16] S. Hajela, "HP OEMF: alarm management in telecommunications networks," *Hewlett-Packard Journal*, Oct. 1996.
- [17] K. R. Sheers, "HP OpenView event correlation services," *Hewlett-Packard Journal*, Oct. 1996.
- [18] Sunsoft, "Solstice Enterprise Manager



- Administration Guide Release 2.0," 1997.
- [19] M. Sabin, R. D. Russel, and E. C. Freuder, "Generating Diagnostic Tools for Network Fault Management," *Integrated Network Management V*, pp. 700-711, May, 1997.
- [20] M. Sabin, A. Bakman, R. D. Russel, and E. C. Freuder, "A Constraint-Based Approach to Fault Management for Groupware Services," *Integrated Network Management VI*, pp. 731-744, May, 1999.
- [21] B. Krishnamurthy and D. S. Rosenblum, "Yeast: A General Purpose Event-Action System," *IEEE Trans. on Software Engineering*, vol. 21, no. 10, pp. 845-857, Oct., 1995.
- [22] R. E. Gruber, B. Krishnamurthy, and E. Panagos, "The Architecture of the READY Event Notification Service," *Proceedings of the 19th IEEE International Conference on Distributed Computing Systems Middleware Workshop*, Austin, Texas, USA, May, 1999.
- [23] S. Yucel and N. Anerousis, "Event Aggregation and Distribution in Web-based Network Management Systems," *Integrated Network Management VI*, pp. 35-48, May, 1999.
- [24] G. Liu, A. K. Mok, and E. J. Yang, "Composite Events for Network Event Correlation," *Integrated Network Management VI*, pp. 247-260, May, 1999.
- [25] G. M. Stanley and R. Vaidhyanathan, "A Generic Fault Propagation Modeling Approach to On-line Diagnosis and Event Correlation," *Proceedings of the 3rd IFAC Workshop on On-Line Fault Detection and Supervision in the Chemical Process Industries*, Solaize, France, June, 1998.
- [26] D. W. Gurer, I. Khan, and R. Ogier, "An Artificial Intelligence Approach to Network Fault Management," *Proceedings of the AI in Distributed Information Networks Workshop of the International Joint Conference on Artificial Intelligence (IJCAI)*, Montreal, Canada, pp. 25-33, 1994.
- [27] ITU-T, "Overview of TMN Recommendations," M.3000, Oct., 1994.
- [28] T. Oates, D. Jensen, and P. R. Cohen, "Discovering Rules for Clustering and Predicting Asynchronous Events," *Workshop Notes of the AAAI Workshop on Predicting the Future: AI Approach to the Time Series*, pp. 73-79, 1998.

박 응 석(Yongseok Park)

정회원



1986년 2월 : 서울대학교

전자공학과 졸업

1988년 2월 : 서울대학교

전자공학과 석사

1996년 5월: Purdue대학교

전기 및 컴퓨터공학과  
박사

1996년 5월~8월 : IBM Austin Research Laboratory

1996년 8월~2000년 11월 : AT&T Labs

2000년 11월~2001년 8월 : Core Networks

2001년 9월~현재 : Lucent Technologies

<주관심 분야> Communication Network Architecture,  
Protocols, and Management