Towards a Resilient Message Oriented Middleware for Mission Critical Applications

Jinfu Wang[1], Beatriz Viñal Murciano[2], John Bigham[3]

MPI-QMUL ISRC[1]
Macao Polytechnic Institute & Queen Mary University of London
Macao SAR, China
jinfu.wang@mpi-qmul. org[1]

School of Electronic Engineering and Computer Science[1, 2, 3]
Queen Mary University of London
London, UK
john.bigham@eecs.qmul.ac.uk[3,]
beavimu@gmail.com[2]

***Abstract* — Message oriented middleware (MOM) provides a messaging service layer between the transport and application layer of the networking protocol stack. A resilient MOM system strives to provide a required level of message brokerage service in the face of bursty surges in workload demand, and failures in the underlay network or brokers. Resilience in our MOM system is achieved by a novel workload allocation mechanism which minimizes the quantified risk of workload exceeding capacity of a broker, while introducing redundant mirroring of workload; and also using resilient overlay routing and multi-homing to mitigate and recover from underlay network failure(s). This paper discusses the overall system architecture and the workload allocation and mirroring mechanism we employed. Comparing with round robin maximizing resource reserve ratio, our allocation algorithm provides superior resilience in minimizing the risk of correlated workload exceeding the capacity of system.**

Keywords- message oriented middleware; resilience; self-adaptive; publish-subscribe; quantifying risk; self-healing; overlay routing.

I. INTRODUCTION

Mission critical applications, e.g., financial data delivery banking transactions, and remote control commands all place different requirements on publish subscribe message oriented middleware (PSMOM) for reliability and performance. They are very sensitive to performance degradation and network and broker failures. This paper looks particularly at the provision of resilience to broker and link failures and performance degradation, and also at the related topic of assured delivery.

PSMOM is software used to support communication between components both within a system and between cooperating systems [1][2]. It is especially suitable when a message needs to be distributed to multiple subscriber applications as it reduces the number of point-to-point connections required for applications to communicate. MOMs are compatible with, and indeed can form a central component of the Enterprise Service Bus (ESB) architecture [3], where the MOM message broker acts as the “bus” between applications. In practice because of issues relating to differing enterprise responsibilities, geographical locations, bottlenecks in performance [4], mitigation of single points of failure, multiple brokers will be deployed using both high-availability clustering and broker federation.

Resilience is the ability of the MOM system to provide and maintain an acceptable level of service in the face of faults and challenges to normal operation, such as internal broker faults and external link failure and denial of service (DoS) attacks. To provide a resilient brokerage service, our PSMOM system is an overlay of federated brokers over a network, composed of several local domains, i.e., a few neighbour brokers instead of one centralized broker. The goal of the design is to provide adaptive configuration solutions and reactive mechanisms to provide a resilient service despite the risks from link or broker failures or degradation and surges in workload.

The resilience is achieved by: firstly, allocating the workloads to both primary brokers, and mirror brokers in the local domain - the mirror brokers provide redundancy at application end points in case of faults in primary broker; and secondly, employ overlay level multi-path routing in the overlay domain to provide reliable networking over wide area network between local domains in face of link faults and performance degradation. This paper discusses the first part, workload allocation and mirroring algorithm which quantifies and minimizes the risk of overloading brokers while exploring the correlations between messaging workloads; and also explains architecture and operation of the system.

This work has been driven by the requirements to produce resilient and secure MOM solutions appropriate for future real-time and business critical systems, which form the focus of the EU FP7 GEMOM project (Genetic Message-Oriented Secure Middleware) [5][13]. Besides the work that this paper is centered on, adaptive security and trust mechanism [7][8][11] are another strand of the GEMOM system. However, these will not be discussed in this paper.

The rest of the paper is structured as following. In Section II we present related researches and state of art MOM approaches. Then, in Section III the architecture and resilience mechanisms of our MOM system is explained. Section IV presents the risk-aware workload allocation and mirroring mechanisms, Section V explains the operation of system and Section VII describes an industrial application scenario. Finally, conclusions and ongoing works are briefly given in Section VI.

II. Background and Related Works

Most state-of-the-art industrial MOM systems, e.g., Apache's AMQP Qpid, offer the high-availability clustering (HA clustering) and federation functionalities, to enhance reliability, interoperability and scalability of messaging service. The two mechanisms can be combined and adapted to build MOM deployments with different topologies and to suit scenarios where applications have different performance and reliability requirements. The current HA clustering techniques and federation techniques are orthogonal concepts. Broker clustering techniques consist in creating groups of brokers that work closely together. HA clusters improve the reliability by replicating entire states and messages of the working broker to another broker. They support the clients to failover to another broker in the cluster if the working fails. Federation enables the communication between brokers. Federation thus both supports connecting brokers in different domains as a MOM overlay, and improves system scalability by distributing computation and bandwidth contention of the message brokerage to multiple brokers such as [12].

Some problems associated with the current HA clustering and federation motivates our research. Firstly, the bursty surge in demand of workloads will cause significant performance degradation [15] and the surge of correlated workloads will have a supper added influence on such degradation. A workload allocation mechanism that minimizes such problem is missing. In our approach instead of only considering mean value of workloads, our system employs a novel workload allocation and mirroring algorithm which quantifies the probability of workloads exceeding capacity of brokers, while exploring the correlation between messaging workloads with variance and covariance matrix based method; Secondly, introducing redundancy with HA clustering requires replicating the entire workload of a primary broker to another broker, which requires much computation resource and the replication intra LAN does not fit well with MOMs deployed for internet messaging. With our algorithm, we mirror the partitions of workloads (called items) on a primary broker to different neighbor brokers in the local domain, instead of replicating entire its workloads to another broker. Thirdly, the federation of MOMs across internet requires low converge time from random underlay IP network failures and faults. Our system employs overlay level multi-path routing for resiliency in networking between local domains.

Resiliency in P/S based MOM is a popular research field, with some related works focusing on specific contexts. Yoon [17] designed a set of protocols to replace a faulty broker with an extra spare broker. The replacement is through recovering the connections to neighbor brokers by contacting an external directory service and recovering subscription tables from the reconnected neighbor brokers. Their protocol is for on-demand replication which means there is no live redundancy for continued messaging. The entire workload of faulty broker is replicated together similar as the failover in HA clustering. Our approach provides live mirroring for continued messaging in the face of single broker failure with workload mirroring and reconfiguration solution for multiple broker failures.

Kazemzadeh [6] describes an approach to handle broker failures, while maintains order and provides exactly once delivery of publications. Each broker maintains a partial topology mapping of brokers in its publication tree, and a faulty broker in the mapping is replaced by reconnecting to the next working broker in the mapping. The cost of their approach is to trade for stability with latency comparing with multi-path routing over internet. Their approach is designed to fit content based, not topic based P/S.

Finally some works, such as Snoeren et al. [16], employs overlay level multi-path routing or multi-homing, which has evolved from similar network level approaches [9][14], to build a fault tolerant P/S system, by constructing redundant disjoint forwarding paths between subscribers and publishers. While we employ similar overlay level multi-path routing for connecting disjoint domains with QoS awareness, we also improve the application end point resilience by workload allocation and mirroring in local domains as described in the following sections.

III. Overview of The GEMOM System Architecture

The GEMOM resilience functions are designed to provide an acceptable level of continued messaging functionality in the face of faults or failures detected by the monitoring components. This is achieved by allocating primary workload and mirroring workload according to computed policies which minimize the risk of overloading brokers. As explained later the risk of system failure is the criterion that is used to assign load and mirrors. The overall traffic is allocated among a group of neighbor brokers intra cluster (e.g., in the case of Enterprise Messaging) or across internet (e.g., internet messaging). This group of neighbor brokers and the clients they serve is a local domain. Between disjoint local domains, overlay level multi-path routing can be employed to improve network link resilience, though this will not be discussed in the paper. This network model is shown in Fig. 1. The resilience architecture in GEMOM includes a Management Agent (MA) collocated with each broker. The Management Agent consists in components including Resilience Manager (RM), Overlay Manager (OM) and anomaly monitors.

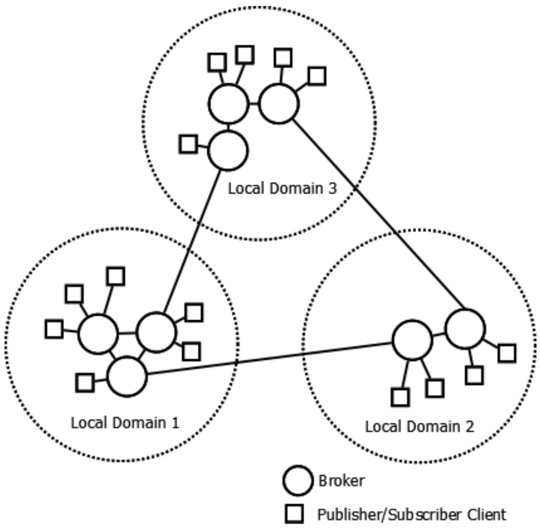

Figure 1. Network model of MOM system.

One of the MAs in the local domain is elected as Master MA, which is used to compute policies and distribute the policies and updates to other slave MAs. When a Master MA fails, one of the other MAs is elected as the new Master MA. The up-to-date policies in all MAs and the master election mechanisms provide resilience to master MA faults. Currently the host with most computation power is chosen as the Master MA, although different election algorithms can be applied. Inside the Master MA its RM and OM operate as following. RM computes the workload allocation policies and redundant workload replication policies which is called ***mirroring*** policies. These policies are computed for different possible state of the MOM system, e.g., initial allocation, single broker failure, double failure and RM builds a case base with these policies. Note that in low load situations the policies associated with switching off certain brokers are essentially the same as the failure case. The RM communicates with the OM and sends updates of case base of the new policy and its context. The OM is the management interface of to the broker and clients. One important functionality of OM is Broker Directory Service (BDS). Like DNS, BDS maps the topics to the physical brokers which act as their primary and mirror brokers. OM keeps BDS up to date according the policy case base and the current state of system. OM listens to the system state from monitoring tools. When a critical event, e.g., a broker failure, is under the radar, an optimal policy under current state of system retrieved from the case base is used by OM to update BDS. The Master MA distributes the changes in policy case base and BDS to all slave MAs to keep them up to date. These updates in BDS push brokers and clients to adapt to current policies.

The monitoring components consist of a set of anomaly detectors, namely broker failure, link failure and bottleneck detectors. The detected events from detectors which reside with each broker and each client trigger the adaptive actions in brokers and clients. The events are also passed to the Master MA to update current system policies.

A risk-aware workload allocation algorithm and hard constraints are used to compute the policies. Comparing with high-availability clustering, the policies are computed and applied on “item” level to reduce computation and the size of case base. Item is a system administrator defined partition of total workload carried by the MOM. An item can be considered as set of topics, treated by the RM as a single entity with respect to re-allocation of resources. A topic in a PS MOM is simply a label that a publisher (or publishers) can use to identify a message stream of particular type and subscribers can subscribe to the topics they wanted. Another difference is that a redundant replication, i.e., a mirror item, can be mirrored to live brokers instead of just to an idle slave broker to replicate entire workload and states of a live broker. Ideally this will allow a smaller set of brokers to carry the traffic at the prescribed risk bound. This algorithm is described in the following section. The overlay level-multipath routing between disjoint local domains is out the scope of this paper.

IV. MESSAGING WORKLOAD ALLOCATION AND MIRRORING

Mission critical applications such as remote control commands and finical data are very sensitive to the service degradation and message loss due to link or broker failures. Service degradation is often caused by the workloads on a broker rise and exceed the processing capacity of the broker [15]. Correlated workloads often exist in real world systems. Such correlated workloads will likely rise together, and thus have a super additive effect on the total workload. Uncorrelated or negatively correlated workloads, on the other hand, will make the total workload more stable. In our MOM system the workloads in a local domain are allocated and mirrored among a few neighbor brokers. RM employs an workload allocation and mirroring algorithm to quantify and minimize the risk of workload exceeding a broker while explore and mitigate such correlation. RM computes workload allocation policies and redundant workload replication policies, i.e., *mirroring* policies, under different possible states of the MOM system, e.g., initial allocation, single broker failure, double failures, using this algorithm. The policies are stored in a case base and retrieved according to current state of system as described in the previous section.

In comparison to HA clustering, where the entire workloads of a primary broker is replicated to another live broker, our mirroring strategy divides and assigns the mirror workloads of a primary broker to several live neighbor brokers in local domains, and does not necessarily introduce

ISBN: 978-1-61208-109-0

many extra idle brokers. The atomic unit for the workload allocation and mirroring is an *item*. An item, according the specific scenario where the MOM system is deployed, is a pre-defined partition from the total traffic carried in MOM system. For example, in the case of topic-based PS MOM, an item is defined as one or more topics that are bundled together in workload allocation. Every item is mapped to a primary and a mirror broker. A physical broker logically can be the primary broker for some and the mirror broker for some other items at the same time.

The RM computes workload allocation and mirroring policies to minimize the quantified risk of overload to the system, and also subject to satisfying important prescribed hard constraints such as requiring that the probability the workload exceeds each broker's capacity is lower than a prescribed upper bound or that the maximum number of topics that a broker is allowed to carry is not exceeded. This problem is formulated as allocating a set of items (an item is also a set, but is a set of topics) *j* to each broker *i*. Consider the optimal primary allocation case, i.e., the case where we simply want to maximize the utility function while allocating items to the brokers which are candidates that satisfy the pre-defined constraints. Then the goal is to find the optimal solution $S = \{j_1, \dots j_k\}$, which is k sets of items allocated to the corresponding k brokers, maximizing overall system utility, i.e., minimizing the overall risk, in the system. The peak period message rate of item set j is a random variable denoted by $\boldsymbol{V^j}$ and the capacity of broker *i* is a constant C_i (both measured in messages per unit time). Given the utility of current local domain U_{domain} , the problem is:

$$\arg\max_{S=\{j_i\}} (U_{domain} = \sum_{i=1}^{k} U_i) \quad (1)$$

$$where\ U_i = R_i^j - \int_{C_i}^{\infty} P_i^j(x - C_i) P_r\big(V^j = x\big) dx$$

The utility U_i of broker i carrying item set j is calculated from R_i^j which is the reward for carrying item set j and the penalty $P_i^j(x - C_i)$ for the arriving messages x exceeding the capacity of broker i. R_i^j is a value associated with item set j and P_i^j ideally is proportional to $(x - C_i)$.

By assuming V^j follows a normal distribution and relaxing P_i^j to a value associated with item set j, we simplify (1) as:

$$\arg\max_{S=\{j_i\}} (U_{domain} = \sum_{i=1}^{k} U_i) \quad (2)$$

$$where\ U_i = R_i^j - P_i^j P_r\big(V^j > C_i\big)$$

We can approximate by $P_r(V^j > C_i) = P_r(\frac{V^j-\mu_j}{\sigma_j} > \frac{C_i-\mu_j}{\sigma_j})$. $\frac{V^j-\mu_j}{\sigma_j}$ follows standard normal distribution, and the parameters are estimated by analyzing peak period message rates samples of all topics. The message rate sample is a series of messages arriving at unit time for each topic in the system over peak periods. We can have mean message rate $\mu_t$of each topics t. Given $V^j = \sum_{t\in j} V_t$, $\mu_j = \sum_{t\in j} \mu_t$. From the sample series of all topics T, we calculate a variance covariance matrix of $T \times T$. $\sigma^2 = \sum_{a\in j} \sum_{a'\in j} cov(V^a, V^{a'})$ where a and a' are any item in j; and $cov\big(V^a, V^{a'}\big) = cov(\sum_{t\in a} V^t \sum_{t'\in a'} V^{t'})$, where t and t' are any topic in a and a'. Intuitively covariance measures the degree to which two variables change or vary together, i.e., co-vary. The positively correlated items' message rates vary together in the same direction relative to their expected values, hence they result in a relative larger σ , which leads to a larger $P_r(V^j > C_i)\ and\ smaller\ U_i$. Hence by maximizing $\sum_{i=1}^{k} U_i)$ the allocation solution will result in a small system risk of being overloaded and positively correlated items are less likely to be allocated to the same broker.

To illustrate the allocation algorithm, there are overall 32 topics in a local domain, which are divided into 6 items, to be allocated to 3 brokers with different capacities in table 3. From sample series data, the mean and the variance covariance matrix among 8 items are calculated as in table 1 and 2. We find the solution for (2) with a Depth First Search using $R_i^j = P^j = \mu_j$ as the reward and penalty associated with j. We apply a prescribed risk threshold=0.007 as a hard constraint (3) to the search:

$$s.t.\ P_r(V^j > C_i) < threshold \quad (3)$$

We compare the allocation result with a round robin allocation which allocates each item in turn to a broker with maximum resource reserve ratio (4). The two solutions are shown in table 3.

$$reserve\ ratio = \frac{C_i - V^{j\,'}}{C_i} \quad (4)$$

The risks of two solutions are evaluated by the normalized gain (5). Risk in our solution and round robin solution is $\mathbf{\underset{1}{Pr}}\ \boldsymbol{and}\ \mathbf{\underset{2}{Pr}}$. A postive gain indicates $\mathbf{\underset{1}{Pr}} < \mathbf{\underset{2}{Pr}}$. The gain is shown in Fig.2.

$$Gain = \frac{(\underset{2}{\Pr}\left(V_i^{j\,'} > C_i\right) - \underset{1}{\Pr}(V_i^j > C_i))}{\underset{2}{\Pr}\left(V_i^{j\,'} > C_i\right)} \quad (5)$$

In this illustration, in term of the estimated risk our solution shows obvious advantage over round robin solution. It is because our allocation algorithm avoids allocating some highly correlated workloads to the same broker by exploring the correlation via the variance computed from the variance covariance matrix.

TABLE I. MEAN MESSAGE RATES OVER SAMPLED PERIODS

Item1	Item2	Item3	Item4	Item5	Item6
150	300	350	200	100	130

TABLE II. VARIANCE COVARIANCE MATRIX OF 8 ITEMS

	Item1	Item2	Item3	Item4	Item5	Item6
Item1	5310	2456	962	1071	840	409
Item2	2456	7544	611	418	1237	542
Item3	962	611	6972	-1622	3401	2387
Item4	1071	418	-1622	7538	-102	4266
Item5	840	1237	3401	-102	4992	3626
Item6	409	542	2387	4266	3626	4202

TABLE III. THE CAPACITY AND SOLUTION OF EACH BROKER

	Broker 1	Broker 2	Broker 3
C_i	700	750	600
Solution 1	Item 2,6	Item 1,3	Item 4,5
Solution 2	Item 1,2	Item 3,5	Item 4,6
$P_{r_1}(V^j > C_i)$	0.0018	0.0016	0.0010
$P_{r_2}(V^{j'} > C_i)$	0.0068	0.0054	0.0170

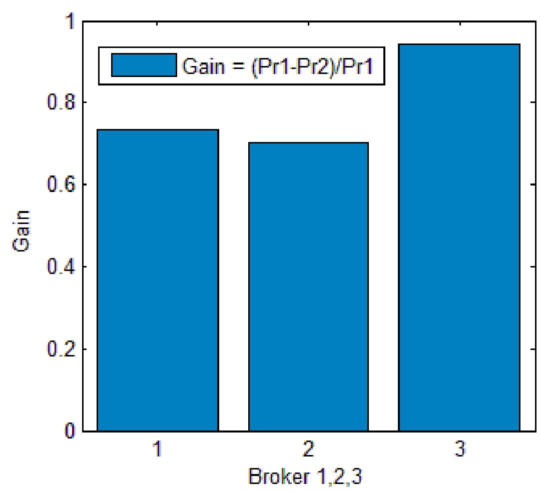

Figure 2. Normalized gain of our solution over round robin at each broker.

V. SYSTEM OPERATION

In each local domain, the Master MA computes the policy case base, and updates BDS in according to the current optimal policy as described in section III. The BDS entries is a 3-tuples in the form of [topic-name@domain, primary-broker, mirror-broker]. The BDS entries and are replicated to slave MAs and pushed to publishing/subscribing clients. Each client connects to both primary and mirror brokers according to the up-to-date BDS entries. The mirroring approach intrinsically trades off redundant resource to reliable messaging service. According to the QoS requirement of different topics, the critical topics with higher requirements for reliability and speed are published simultaneously to both primary and mirror broker; while the topics with lower requirements are published only to primary broker, however a session is pre-established to the mirror broker – the topic starts to publish to mirror broker at the moment when faults are detected in primary broker. Each message published by the same client (including the same message published through different paths) has an unique ID, in one overlay node, only at most one copy of message can exist. The local domain and inter domain messaging is illustrate in Fig. 3.

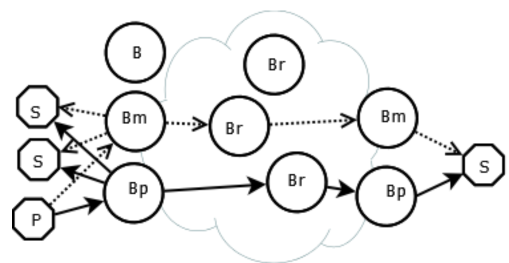

Figure 3. The operation of overall system.

In Fig. 3, primary and mirror brokers (Bm and Bp) provides the clients with both local domain redundancy and overlay level disjoint paths to remote domains.

VI. DEPLOYMENT SCENARIO

GEMOM is designed to emphasis on enterprise messaging where a few independent local domains which consist of a few neighbor brokers, are interconnected as the overlay domain by federation over WAN. For example in one of the case studies, a MOM overlay is deployed to provide messaging service for a national high way monitor and control system. Traffic sensors, mobile clients, toll booth, control centers and other components exchange messages via the MOM overlay. The local domain brokers are distributed across the different parts of the nation. The mirroring approach allows continued message delivery at the application end points in the case where primary broker(s) fail, whether from malfunction or as a consequence of degradation for a DoS attack, or indeed from simply switched off for maintenance. Although disjointedness of underlay network paths, e.g., not sharing a common border gateway, cannot be assured since we do not assume that we have no control over the substrate ISP network, the disjoint placement of brokers carrying primary and mirror workload, and overlay level multi-path routing between federated brokers improve resilience of messaging service in the face of networking faults and failures [9][14].

 ISBN: 978-1-61208-109-0

VII. CONCLUSION AND FUTURE WORK

In this paper, we presented the architecture and operation of the GEMOM system and the risk-aware workload allocation and mirroring mechanisms employed. The allocation algorithm minimizes the risk in a quantified way and mitigates correlated bursty workload to different brokers. The mirroring solution adapts the system to tolerate broker and faults and failures in the application end points in local domains. These mechanisms together with overlay level multi-homing and routing constitute the GEMOM approach towards a resilience MOM system.

In future work, we will develop a more accurate model that profiles broker capacity, and novel scheduling mechanisms are being researched to support QoS awareness and improve performance of overlay level multi-path routing connecting different domains.

ACKNOWLEDGMENT

This research is in the context of the EU project GEMOM grant agreement no 215327, approved by the EU Commission and TSB REMOM project. The authors acknowledge the efforts of all partners (Q-Sphere, NR, VTT, TXT, CNIT, JRC, Datel) who participated in the development of the GEMOM framework. The first author had been partly supported by Macao STDF funding.

REFERENCES

[1] P. T. Eugster, P. A. Felber, R. Guerraoui, and A.-M. Kermarrec, “The Many Faces of Publish/Subscribe,” Computing Surveys, vol 35, PART 2. USA, ACM, 2003, pp. 114-131.

[2] A. Demers, J. Gehrke, M. Hong, M. Riedewald, and W. White, “Towards Expressive Publish/Subscribe Systems,” Lecture Notes in Computer Science, NUMB 3896, Germany, Springer-Verlag, 2006, pp. 627-644.

[3] ESB. Enterprise Service Bus (ESB). Accessed on Sept. 8th, 2010, from http://www.progress.com/psm/sonic/enterprise-service-bus/index.ssp.

[4] P. Tran, P. Greenfield, and I. Gorton, “Behavior and Performance of Message-Oriented Middleware Systems,” Proc 22nd International Conf. on Distributed Computing Systems Workshops, IEEE press, 2002, pp. 645 – 650.

[5] J. Wang, J. Peng, J. Bigham, B. Chew, B. V. Murciano, M. Novkovic, and I. Dattani, "Adding Resilience to Message Oriented Middleware," Proc. 2nd International Workshop on Software Engenering for Resilient Systems (SERENE 2010), UK ACM Press, 2010 pp. 89-94.

[6] Y. Yoon, V. Muthusamy, and H.-A. Jacobsen “On-demand Replication for Failover in Content-based Publish/Subscribe Overlays,” Middleware Systems Reserch Group Technical Report, University of Toronto, November, 2009.

[7] H. Abie,, R. Savola, and I. Dattani, "Robust, Secure, Self-Adaptive and Resilient Messaging Middleware for Business Critical Systems," 1st International Conference on Adaptive and Self-adaptive Systems and Applications, Athens. IEEE Press, Nov. 2009, pp. 153 – 160.

[8] H. Abie, I. Dattani, M. Novkovic, J. Bigham, S. Topham, and R. Savola, “GEMOM - Significant and measurable progress beyond the state of the art,” Proc 3rd International Conference on Systems and Networks Communications (ICSNC 2008), IEEE Press October, 2008. pp. 191-196.

[9] V. Vasudevan, D. G.Andersen, and H. Zhang; “Understanding the AS-level Path Disjointness Provided by Multi-homing.”; Technical Report RCMU-CS-07-141, School of Computer Science, Carnegie Mellon University

[10] J. Mirkovic and P. Reiher, "A Taxonomy of DDOS Attacks and DDOS Defense Mechanisms," Computer Communication Review 2004, Vol 34, Numb 2. ACM USA ISSN 0146-4833, pp. 39-54.

[11] J. Wang, and Bigham, J. 2008. “Anomaly detection in the case of message oriented middleware,” Proc 1st International Workshop on Middleware Security (MidSec '08), ACM, New York, 2008, pp. 40-42.

[12] G. Marsh, A.P. Sampat, S. Potluri, and D.K. Panda, “Scaling Advanced Message Queuing Protocol (AMQP) Architecture with Broker Federation and InfiniBand,” technical report, available at ftp://ftp.cse.ohio-state.edu/pub/tech-report/2009/TR17.pdf, accessed on Sept. 8th, 2010.

[13] H. Abie, R. Savola, J. Wang, and D. Rotondi, "Advances in Adaptive Secure Message Oriented Middleware for Distributed Business Critical Systems,” To Appear on Proc 8th International Conference of Numerical Analysis and Applied Mathematics (ICNAAM 2010), Rhodes, Greece, September 2010

[14] A. Akella, J. Pang, B. Maggs, and S. Seshan, “A comparison of overlay routing and multihoming route control,” Computer Communication Review, 2004, ACM Press, VOL 34 PART 4, Pp. 93-106.

[15] N. Mi, G. Casale, L. Cherkasova, and E. Smirni, “Burstiness in multi-tier applications: Symptoms, causes, and new models,” Proc 9th ACM/IFIP/USENIX International Conference on Middleware (Middleware 2008), Springer-Verlag, New York, 2008, pp. 265-286.

[16] A. C. Snoeren K. Conley, D. K. Gifford, “Mesh-based content routing using XML,” in ACM SIGOPS Operating Systems Review Volume 35 , Issue 5, ACM Press, 2001. pp: 160 - 173

[17] R. Kazemzadeh and H.-A. Jacobsen “Reliable and Highly Available Distributed Publish/Subscribe Service,” Proc 28th IEEE International Symposium on Reliable Distributed Systems, IEEE Press, 2009, pp. 41-50.