MARKETNET: A MARKET-BASED ARCHITECTURE FOR SURVIVABLE LARGE-SCALE INFORMATION SYSTEMS[1]

Y. Yemini, A. Dailianas, D. Florissi

Department of Computer Science
Columbia University,
450 Computer Science Bldg.,
New York, NY 10027
Email: {yemini, apostolo, df}@cs.columbia.edu

Abstract

This paper describes the MarketNet architecture for survivability of large-scale information systems. In MarketNet, access to information system resources is governed by a market economy. Processes purchase services from resource managers using their budgets and seek to optimize the Quality of Services (QoS) they obtain, while resource managers seek to maximize their revenues. The resources traded include physical resources such as bandwidth, processor cycles, storage, I/O devices or sensors as well as higher-level services such as file storage, name service, database or web service. Prices reflect a balance between supply and demand for a resource while the budgets assigned to client processes determine their priority in gaining QoS access to resources.

MarketNet provides hedging mechanisms to support survivability under loss or attacks. Resource managers replicate resources to hedge against loss. In deciding on creation and location of such backup replicas, a resource manager seeks to balance the investment in current costs of maintaining replicas, against the future value of the backup in case of a loss. Currency provides a uniform resource instrumentation that reflects the net-present-value of resource loss and admits optimized load redistribution and graceful degradation upon loss. A loss of a resource reduces the supply, thus automatically causing clients to redirect their demand to backup replicas. Reduced availability results in rising prices of the replicated resources creating a natural selection process where high-priority clients can apply their budget to continue and obtain high QoS, while low priority clients are priced-out. Thus a loss results in graceful selective degradation of services that optimizes the balance between available resources and demands.

1. Introduction

Future large-scale information systems will provide a multitude of services to ever-growing users (consumers) of various applications. The technical challenges that arise include the design of a robust, decentralized, efficient, self-organizing, self-protecting, adaptive information system.

To face the above challenges, the MarketNet project is developing a set of novel mechanisms, based on economic paradigms and principles, that ensure the systematic, quantifiable and predictable reliability of large-scale information systems. These mechanisms provide the means for meaningful replication of resources, predictable system behavior upon loss of resources, ways for applications to protect against the risks associated with their operation in unreliable environments, and ways to quantify and tune the reliability of resources and systems.

Survivability of information systems is of great importance both due to concern for loss of services (i.e., availability) and loss of — possibly vital — data (i.e.,

[1] This research is sponsored in part by the USAF, Air Force Materiel Command, under contract F30602-97-1-0252, "MarketNet: A Survivable, Market-based Architecture for Large-scale Information Systems". The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied of the Defense Advanced Research Projects Agency (DARPA), the Air Force, or the U.S. Government.

reliability). The loss of services and data can occur due to either failure of components or due to malicious attacks. To improve reliability and availability of a large information system, one must store redundant information while incurring the least cost in system resources. This is particularly important in real-time systems such as multimedia information systems where recovery from failure must occur in real-time. One must also account for the importance of resources, services, and data to users, which changes over time when deciding what resources and services to replicate.

Networks are inherently unreliable environments. Applications operating in such environments are prone to the uncertainties of unavailability due to failures or reduced availability due to congestion for resources. In MarketNet degradation due to failure is graceful and selective. Automatic load balancing mechanisms redirect service requests to available replicas favoring high-priority customers. Further, applications can reduce susceptibility to network uncertainties and operate in a more predictable and reliable fashion by purchasing protection against the risk of resource unavailability.

Survivability at the level of individual components and services is vital, but cannot by itself guarantee the reliable and robust operation of large-scale networks. MarketNet provides system-architecture-level organization structures that can effectively limit the spread of faults and attacks imposing quantifiable bounds on the attack power of network entities.

The rest of the paper is structured as follows. Section 2 presents the market-based resource management approach used by MarketNet to address the issues of survivability in large-scale information systems. Section 3 describes optimal survivability of services and the role of resource managers in the network market economy. Section 4 describes the operation of customers in MarketNet and their role in achieving optimal survivability of applications. Section 5 describes the interaction of players (resource managers and customers) in a market based information economy, mechanisms to automatically limit access to high-priority customers in case of loss, and to prevent attacks to sensitive resources and services. Section 6 describes the use of hedging mechanisms used against uncertainties in networks by purchasing protection in the form of instruments. Section 7 describes the organization of resources in currency domains and their use in addressing security and survivability concerns. Section 8 concludes the paper.

2. The MarketNet Approach: Market-based Resource Management

To ensure the systematic, quantifiable and predictable reliability of large-scale distributed systems, MarketNet develops and demonstrates novel technologies, based on economic mechanisms. These technologies include mechanisms, protocols, and algorithms to support the reliability of networks; and the adaptation of network clients and services to changing resource availability.

Access to resources in MarketNet is through a market economy, where network elements can purchase and sell services, optimizing their utility measures and maximizing their revenues. The resources traded include both physical resources such as CPU cycles, storage, bandwidth, I/O devices or sensors as well as higher-level services such as file storage, name service, database or web service. Access to resources is done through currency, carrying unique currency Identifiers (IDs) that establish traceable liability in the use of resources. The network economy built by MarketNet is a collection of mechanisms that regulate the allocation of a set of resources among competing demands.

Several salient features of economic mechanisms are of particular importance in addressing the challenges of survivability of large-scale information networks. First, local decentralized decisions and protocols are used to accomplish global behaviors and emerging global behaviors can be quantified as competitive equilibrium between supply and demand.

Second, competitive equilibrium reflects various measures of optimal allocation of resources relative to individual preferences, QoS requirements, and budgets of consumers. This in turn provides a framework for optimal local behaviors, and can be highly stable with respect to local perturbations. Further it can adapt to dynamic changes in supply and demand reflecting the availability of resources and providing the opportunity to automate load balancing and load redistribution.

Third, currency and prices provide a uniform global measure to quantify value and availability of resources, as well as consumer demands and priorities, and a mechanism for unified, scalable, resource-independent access control to resources. Consumers may gain different levels of resource access (QoS) depending on their purchasing power. Access of resources through money in MarketNet provides resource managers with a uniform instrumentation to quantify the net-present-value of resource loss. Resource managers invest current revenues to balance

against loss by replicating resources to maximize future availability. Thus, the services most valued by their clients are provided with the highest redundancy.

Fourth, prices of resources and services provide a mechanism to control access to resources and a natural selection process of restricting access to high-priority customers in case of loss. Loss of resources due to failures or attacks results in optimal load redistribution. Clients trying to optimize their purchase of resources within their available budget, redirect access to available replicas. Rising prices reflecting the reduction of supply due to the loss, limits access to high-priority (high budget) clients.

3. Optimal Survivability of Resources and Services

Dynamic replication of popular or vital services and resources can be used to improve the performance and the reliability of the system. Part of the difficulty is deciding which services and resources to replicate and how many replicas of each to maintain. In order to achieve better performance characteristics, the number of replicas of each service and resource should change to reflect variations in the user access patterns. To achieve better survivability, the number of replicas of services and resources should also change with the "importance" of each service and resource. Thus, it is of great benefit and importance to replicate services and resources *dynamically*, as the importance or the demand for them changes.

Replication decisions are taken by resource managers, which desire to maximize revenues. Value is associated with resources and services in MarketNet. Dynamically adjustable prices reflect this value and provide resource managers with the necessary instrumentation for uniform quantification of loss. Resource managers estimate the net-present-value of resource loss taking into account the demand and the importance of the service of the resources they manage. Managers use these estimates to hedge against loss, by investing current revenues for future availability.

Dynamic replication could result in significant system overheads, e.g., in the form of additional storage and communication load. In designing market-based replication mechanisms in MarketNet we are building on efficient replication mechanisms that do not incur significant system overheads. We are addressing the issue of appropriate trigger mechanisms for initiating replication or deletion of services and resources. A poor choice of triggers can cause the system to perform unnecessary replication or deletion which can lead to a waste of system resources. For example, deleting the replica of a service that will be replicated again in the "near" future, leads to an unnecessary increase in the replication cost. Our investigation is targeted to *threshold-based* trigger methods with *hysteresis* behavior.

Threshold-based approaches can result in a cost-controlled creation and deletion of replicas, according to the changes in the access patterns. The main motivation for using them is that many systems incur significant server setup, usage, and removal costs. Under light loads it is not desirable to operate unnecessarily many service replicas, due to the incurred setup and usage cost. On the other hand, it is also not desirable for a system to exhibit very long delays, which can result from lack of services under heavy loads. Threshold-based approaches allow to maintain high levels of performance of a service and at the same time maintain an acceptable cost for operating the service, by dynamically adding or removing services, depending on the system load. Picking appropriate threshold values is not a trivial matter. The values for these thresholds depend on many factors, such as the values of service setup, usage, and removal costs as well as characteristics of the user access patterns and the service demand.

The service characteristics that trigger the replication process are prone to oscillations which may result in high and unnecessary setup and removal costs associated with service replication. More specifically, it is desirable to add services only when a system is moving towards a heavily loaded operation region, and it is desirable to remove services only when a system is moving towards a light load operation region — it is not desirable to alter the number of services during "momentary" changes in the workload, i.e., during oscillations. Simple threshold systems are therefore not sufficient, but when adding hysteresis to the system such oscillation regions can be effectively avoided.

4. Optimal Survivability of Applications

In a similar fashion to the selfish operation of the service providers who try to maximize their revenues, the operation of customers is governed by their desire to optimize purchasing of resources within the available budget. Customers choose the cheapest service that satisfies their preferences and their QoS requirements.

The important question we are investigating is that of the utility functions that applications and users use to represent their resource preferences. The type of eco-

nomic models considered in MarketNet depend strongly on the assumptions on these utility functions. Unfortunately, little is known about these utility functions. Clearly, however, many are strongly non-linear. For example, for interactive packet audio, utility drops to zero once bandwidth drops below a few kb/s for speech and about 56 kb/s for music, or delay reaches the limit of interactivity. Utility also strongly depends on the use a particular resource is put to, e.g., whether the audio is used for distributed game playing, background music, or a long-scheduled business meeting.

5. Supply-demand Equilibrium and the Role of Prices in MarketNet.

In the previous sections we have described the operation of the customers and service providers in the MarketNet economy. Here we describe the economic mechanisms underlying their interaction. This interaction is governed by competitive equilibrium, which captures various measures of optimal allocation of resources relative to individual preferences and budgets of consumers. It also provides a framework for optimal local behaviors, can be highly stable with respect to local perturbations, and can adapt to dynamic changes in resource supply and demand.

MarketNet develops technology to accomplish global competitive equilibrium of resource supply and demand through decentralized, and local mechanisms and decisions. Such mechanisms include resource manager mechanisms and algorithms to adjust prices to reflect resource availability and balance supply and demand, client mechanisms and algorithms to optimize clients' preferences, and dynamic directory services to post resource services and prices. We are further investigating mechanisms to accomplish high stability of competitive equilibrium relative the dynamic perturbations of supply and demand, and fast convergence to a new equilibrium under major changes in resource availability and demands. The MarketNet directory services addresses two main issues. First, current directory services are not designed to support dynamic, frequent, and secure updates of the information they advertise. Second, they cannot scale to the numbers of services expected to be advertised in future network economies or to the number of accesses for retrieving information they store. We are exploring enhancements of the Domain Name Service (DNS) and of the Resource Reservation Protocol (RSVP) to support advertising of resource prices. These enhancements lead to non-trivial design issues. For example, DNS presently does not offer protocol primitives to change records dynamically. Thus, changes in prices cannot be dynamically recorded in DNS databases. While versions of DNS being standardized now include some ability to change records dynamically, this leads to complex security issues. Furthermore, DNS caching and replication mechanisms create inefficiencies in accessing records that change rapidly. Price records for services require specialized handling and adaptation of DNS mechanisms to assure that neither inefficiencies nor security risks are created through changes. RSVP currently does not provide protocol support for conveying prices or budgets. Enhancements of RSVP should be able to negotiate prices and purchase resources. This poses particular difficulties since failure of an upstream negotiation may force cancellation of earlier "contracts". The complexity of negotiations and the desire of consumers not to reveal their utility functions make the protocol design difficult. One can define an RSVP payload consisting of a Java program that computes the user's trade-off between price and amount of resources to be reserved and tracks the amount of money spent as the reservation progresses through the network.

Competitive equilibrium achieved through the pricing mechanism significantly assists in system reliability as it provides automatic load balancing under normal operation and automatic demand redistribution when loss of resources or services occurs. Consider the following example: a route or a path of an information network is owned by a manager (supplier). The resources at each link along the path are buffer and bandwidth, which are priced by the manager. The consumers (user traffic classes) buy resources such that their QoS needs are satisfied. The network provider prices resources based on demand for resources from the consumers. Consumers have preferences for packet loss, average delay and throughput. The system is dynamic, as user sessions within a traffic class arrive and leave at arbitrary times, causing fluctuations in demand. When congestion along the path occurs, the manager rises the prices reflecting unavailability of appropriate supply to handle the current demand. Customers switch to other managers offering similar paths. According to the mechanisms outlined in Section 3, the manager uses part of the current revenues to invest in creating replicated services (in this case more routes are acquired), to account for possible future loss. This increases service and drives prices down until revenues are balanced against costs. If there is a failure in one of the routes, the manager readjusts the existing sessions along other routes that are active. The failure results in an increase in the advertised prices of the replicated resources reflecting the decrease in availability. This creates a natural selection process where high-priority clients can apply their budget to continue and obtain high QoS, while low

priority clients are priced-out. Thus a loss results in graceful selective degradation of services that optimizes the balance between available resources and demands.

The previous example depicts a few important features of the role of pricing mechanism in MarketNet. First, prices combined with the budgets available to clients provide an effective mechanism for establishing dynamically tunable access control to resources and services. High prices can reflect limited availability of resources (e.g., due to loss, congestion or attacks), or the intention of resource managers to limit access to resources only to a limited set of clients. High-priority clients in the first case can be provided with the appropriate budget to access the service even when its availability is reduced. Qualified clients in the second case, can be the only ones provided with the appropriate budget to access the service.

Second, rising prices help reduce the duration and damages caused by faulty or malicious clients, forcing them to exhaust their available budget at an increasing rate to sustain the attack.

Third, rise of prices leads to automatic load redistribution, since customers aiming to optimize their resource utilization subject to their budget will redirect their request to replicas of the service or similar services offered at lower prices.

Fourth, rise of prices will trigger the mechanisms to investigate whether a resource is under attack. Tracing back to the owner of the currency that is used to access the resources (through the currency IDs of the currency used for access) will identify and isolate the attack source.

6. Hedging to Reduce Network Uncertainties

Financial markets have created risk reduction instruments such as futures and options, used to hedge intrinsic uncertainties. We are developing analogous instruments to reduce intrinsic network uncertainties. Specifically, we concentrate on two types of uncertainties. The first is the uncertainty associated with the loss of a service (e.g., failure of critical servers or loss of communications bandwidth). This uncertainty exposes network processes to the risk of future service unavailability. The second uncertainty of interest is that associated with competing demand for service access. This uncertainty exposes the processes to the risk of inadequate QoS. We are developing instruments and pricing mechanisms that allow processes to purchase protection and reduce the risk of availability and QoS delivery.

The testbed for the use of financial instruments to reduce network uncertainties uses instruments to protect against unavailability of bandwidth due to competing demand. The bandwidth provider offers two types of services: guaranteed and best-effort. Users hedge against unavailability of best-effort service by buying the right to use guaranteed service should the need to do so arise. Guaranteed service is traded in the *primary market*, where the service provider sells forwards contracts that guarantee use of bandwidth slots in the future, and in the *secondary market*, where users trade options on forwards contracts along with forwards contracts themselves. Clients with strict QoS constraints buy forwards contracts in advance even if uncertain about their traffic characteristics. Once they discover they will not need to exercise some of the forwards contracts they have bought, they trade them in the secondary market, where low-budget customers purchase them to hedge against congestion for best-effort service due to competing demand.

The goal of this effort is to demonstrate hedging as an appropriate mechanism to provide cheap protection against uncertainties in networks, and furthermore to show that the level of uncertainty and guarantees can be optimally tailored to the user's needs and available budget.

7. Currency Domains Encapsulate High Level Protection

The mechanisms presented so far deal with increasing reliability and availability of resources in the face of loss of resources. Another important goal in MarketNet is limiting the spread of faults and attacks that could severely impact the reliability of the system and the availability of resources.

In MarketNet resources and clients are organized in currency domains. Domains use special currency to provide unified, scalable access to their services. The currency of a domain gives the holder the right to access resources in that domain. Therefore, to gain access to the resources in a domain, clients first have to use part of their budget and exchange it for the desired currency. Budgets of network entities (such as domains and clients) are hierarchically enforced by secure, trusted banks. Enforcement of budgets is a very powerful tool for limiting attacks and damages in large-scale networks. The access and damage a client can do is limited by their available budget expressed in terms of the currency of the particular domain they want to attack.

Currency domains encapsulate domain-level protection policies set by the domains. Specifically, domains

control who can acquire their currency, along with the total currency outflow, the rate of currency outflow and other parameters. This imposes strict, domain-controlled limits on the access and attack power of any entity wishing to access the domain resources. Clients purchase currency of a domain which provides unified, resource-independent access to resources. Furthermore, as mentioned earlier, currency is tagged with unique currency IDs that are used to establish liability in the use of resources. Monitoring of currency flows is used to detect intrusions, and currency IDs are used to identify and isolates attack sources. Monitoring can be done in much the same fashion as is common in transaction processing systems. Currency flows provide a good way to model temporal behaviors of clients and patterns of resource access to classify activities into those that are legitimate and those that seem suspicious and hence warrant further inspection and authorization. Once an attack has been identified, currency IDs can be used to identify and isolate the source of the attack, without affecting the operation of users in legal domains. In case a whole currency has been conquered and the breach is detected, the currency of that domain can be declared invalid, limiting the spread of faults or attacks, until appropriate action is taken to restore normal operation.

8. Conclusions

Application of economic and market principles may provide effective solutions to the key survivability challenges in large-scale information systems.

Quantification of the value of resources and services through currency enables the deployment of efficient, rational, dynamic replication of services according to the interest of users and the importance of services. Unified automatic monitoring and identification of attacks, along with identification and isolation of attack sources, are made possible through a collection of mechanisms based on the organization of resources in currency domains issuing traceable currency and establishing liability in resource usage. Supply-demand equilibrium reflected in prices of services leads to automatic load balancing and redirection to replicated resources. Optimal resource availability is achieved through investment of current revenues to replicate resources and services. The incentive for the investment is the loss of revenues associated with resource unavailability. Instruments similar to those in financial markets can be used to purchase protection against intrinsic network uncertainties such as loss of a service or resource due to failure, or reduced availability due to competing demands.

References

[1] Clearwater, S., editor. *"Market-based Control of Distributed Systems,"* World Scientific Press, 1996

[2] Hull, J. C. *"Options, Futures, and Other Derivatives,"* third edition, Prentice Hall.

[3] Kurose, J., M. Schwartz, and Y. Yemini *"A Microeconomic Approach to Optimization of Channel Access Policies in Multiaccess Networks,"* Proc. Of the 5th International Conference on Distributed Computer Systems, Denver, Colorado, 1995.

[4] MacKie-Mason, J., and H. Varian *"Pricing the Internet,"* in B. Kahin and J. Keller, editors, *Public Access to the Internet*, ACM, Boston, Massachusetts, May 1993.

[5] Sairamesh, J., D. Ferguson, and Y. Yemini *"An Approach to Pricing, Optimal Allocation and Quality of Service Provisioning in High-speed Packet Networks,"* in Proc. of the Conference on Computer Communications, Boston, Massachusetts, April 1995.

[6] Walsh, W., M. Wellman, P. Wurman, and J. MacKie-Mason *"Some Economics of Market-Based Distributed Scheduling,"* In Proc. of the 8th International Conference on Distributed Computing Systems (ICDCS-98), Amsterdam, the Netherlands, May 1998.

[7] Yemini, Y. *"Selfish Optimization in Computer Networks,"* Proc. of the 20th IEEE Conference on Decision and Control, pp. 281-285, San Diego, CA., Dec. 1981.

[8] Zhang, L., S. Deering, D. Estrin, S. Shenker, and D. Zappala *"RSVP: A New Resource ReSerVation Protocol,"* IEEE Network magazine, 7(5):8-18, September 1993.