Resource Trading Agents for Adaptive Active Network Applications

Lidia Yamamoto – Guy Leduc

University of Liège
Research Unit in Networking
Institut Montefiore, B28
B−4000 Liège, Belgium

yamamoto@run.montefiore.ulg.ac.be
leduc@montefiore.ulg.ac.be

ABSTRACT: *Adaptive applications seem to be the only realistic answer to the increasing diversity and decentralisation in networks. In order to accommodate fluctuations in network conditions, adaptive applications need to obtain information about resource availability. Using mobile agent technology applied to active networks, new models for adaptive applications can be envisaged. In this article we review the state of the art in agent and active network techniques for adaptive applications, and describe our work on an agent model for trading resources inside an active network node. We apply the model to the case of an audio mixing application. Our simulation results indicate that the model allows the application to make efficient use of the available resources, and to share resources according to user preferences.*

RÉSUMÉ: *Les applications adaptatives semblent être la seule réponse réaliste à la diversité et à la décentralisation toujours croissante dans les réseaux. Afin de pouvoir s'adapter à des fluctuations dans les conditions du réseau, les applications adaptatives ont besoin d'information sur la disponibilité des ressources dans le réseau. En appliquant la technologie des agents mobiles aux réseaux actifs, de nouveaux modèles pour les applications adaptatives peuvent être envisagés. Dans cet article nous révisons l'état de l'art des techniques basées sur les agents et les réseaux actifs pour les applications adaptatives, et nous décrivons notre modèle basé sur des agents pour échanger des ressources dans un noeud actif. Nous avons appliqué le modèle à une application de mixage audio. Nos résultats de simulations indiquent que le modèle permet à l'application de faire un usage efficace des ressources, et de partager les ressources selon les préférences des utilisateurs.*

KEY WORDS: *active network, mobile agent, resource management*

MOTS−CLÉS: *réseaux actif, agent mobile, gestion de ressources*

1. Introduction

In the context of an increasingly decentralised and heterogeneous network such as the Internet today, it is very difficult for the applications to know how much quality they can really expect. A lot of research effort has been dedicated to techniques to offer QoS guarantees, but these techniques are only effective when deployed in all the nodes concerned by a communication, or at least in all the nodes where resource shortage may occur. However, it is very difficult to achieve global agreements such that these techniques can be deployed. They are feasible in a private enterprise network, ISP or isolated operator, but not at the global scale. Therefore the only realistic answer seems to be to rely on applications that are able to adapt to a wide range of network conditions in a dynamic way, and in particular to the amount of resources available when these resources cannot be reserved in advance.

An adaptive application must be able to make optimal use of the available resources, and be able to adapt itself to fluctuations in resource availability. But this kind of application faces the difficulty to obtain enough information about the network conditions, due to the Internet black–box model.

Clearly the use of more intelligent network elements in the network can allow applications to obtain the feedback they need to perform the adaptation functions more easily. Recently router support has been considered to assist adaptive applications to achieve better performance. But then the same question arises, how to add new functionality inside the network nodes without going through a long standardisation process and without having to face the difficulties of global scale deployment.

Mobile agents and active networks can become useful tools to help in the adaptation process, since it becomes possible to inject customised computations at optimal points in the network, and the deployment problem can be easily solved using dynamic code mobility. In the long run this can lead to self–configurable, auto–adaptive network elements that are intelligent enough to "learn" the protocols they need to use at a given moment, according to the devices available, services offered, operator policies, user demand, etc.

A special class of adaptation mechanisms is the so–called market–based control [CLE 96], for which a considerable amount of research results are available mainly in the agents field. It provides algorithms inspired by optimisation and economy theories for distributed control of resource usage, with many applications to computer and telecommunication networks. The benefit of such mechanisms is two–fold: on one side, optimal resource sharing configurations can be achieved in a decentralised way; on the other side, it becomes easier to quantify heterogeneity in terms of resource availability, to offer the users the opportunity to trade one type of resource for another.

However, relatively few results have been shown which directly apply such artificial economy models to the specifics of active networks, with special attention to highly adaptive applications. We address this issue in this article by providing a simple model which allows the active applications to make decisions

about the amount of resources to use, according to the network conditions found in the active nodes. Using such a model, an audio mixer is developed as an instance of adaptive active network application, which is able to trade bandwidth for memory according to the available prices of each resource.

The article is organised as follows: in Section 2 we review the state of the art in agent and active network techniques for adaptive applications. Section 3 presents our model to trade resources inside an active node. Section 4 applies the model to the case of an audio mixing application. Section 5 shows simulation results for the audio application, and Section 6 concludes the article.

2. Background

In this section we give a survey of the current research directions concerning agent and active network techniques applied to adaptive applications. We start with an overview of current techniques used in network and transport level adaptation protocols, and then discuss the potential of mobile agents and active networks for such applications, with a survey of current proposals in this area. After that we review market–based control research applied to resource sharing in computer and telecommunication networks. The use of such techniques in the active network context is the focus of our work.

2.1 Adaptive Applications

Adaptive applications can tolerate fluctuations in resource availability, and are necessary in a heterogeneous environment such as the Internet today, where different network technologies and user terminals are interconnected together, and over which a multitude of services coexist. In the case of multimedia applications, a good survey can be found in [VAN 00].

The adaptation mechanisms can be implemented at several layers of the protocol stack, ranging from pure application layer techniques to network level protocols. For example, we can adapt to the available bandwidth using elastic traffic that reduces the data rate generated in presence of network congestion. Fluctuations in delay can be dealt with by using elastic buffers to adjust the play–out time. To deal with CPU and memory bottlenecks, some interaction with the operating system is necessary (see Section 5 of [VAN 00] for examples). Our focus in this paper is on network and transport level mechanisms for adaptation. At this level, the classical approaches are typically end–to–end based and adapt to the available network bandwidth (congestion control). The best known example is the TCP protocol, which adjusts its window size in response to the current level of congestion. Since TCP is so widespread today, but does not support the emerging multimedia applications, a lot of research is in progress in order to obtain TCP–friendly protocols for real–time and streaming applications, for unicast as well as multicast traffic [FLO 00] [KHA 00].

It is well known that classical end–to–end approaches often present drawbacks such as annoying quality fluctuations, sub–optimal resource utilisation and sharing, and slow convergence. Part of these problems come from the fact that the applications try to adapt in a blind way, without having enough information about the network conditions, since the Internet uses a black–box model, which hides all information from the end hosts.

Recently some router support has been considered to assist adaptive applications to achieve better performance, e.g. [GOP 00]. However, how to have such schemes widely accepted and deployed is still a question mark.

We believe that code mobility as provided by mobile agents and active networks is the only generic solution to allow adaptive software to be incrementally deployed and evolve through usage experience. In Section 2.2 we give a brief introduction to the active network concepts used throughout the article, and in Section 2.3 we discuss the relationship between mobile agents and active networks. In Section 2.4 we review the current research efforts towards adaptation protocols that can benefit from code mobility.

2.2 Active Networks

Active networks (AN) allow the network managers or users to program the network nodes according to their needs, offering a great amount of flexibility. The nodes of an active network [TEN 97] are capable not only of forwarding packets as usual but also of loading and executing mobile code. The code can be transported out–of–band, within specialised signalling channels (programmable networks) or in–band, within special packets called "capsules" (active networks). Capsules might contain the code itself (such as in [HIC 99]) or a reference to it, such that it can be downloaded when the first capsule containing the reference arrives at a given node (such as in [WET 98]).

If the distinction between active and programmable networks seemed at some point in time clear [TEN 97] [CHE 98], the tendency today seems to be towards an integration of the two concepts [ALE 99], since both are forms of achieving open programmability in networks [CAM 99], and special flavours in between or combining both approaches are also possible [ALE 99] [HJA 00]. In this paper we focus on the capsule model in order to be more generic and avoid programmability restrictions.

There are basically two AN architectural lines, derived from the two original communities on out–of–band and in–band programmability. The first one is the IEEE P1520 reference model [DEN 99], which offers a set of standard interfaces to program IP routers, ATM switches and other network devices. The P1520 model is the first standardisation effort towards open network programming interfaces, and is likely to be one of the first AN interfaces to become commercially available.

 The second architectural line is the framework for an active node architecture [CAL 99] which is being proposed within the DARPA AN research community. It includes a supporting operating system (the NodeOS), one or more execution

environments (EE), and the active applications (AA). These components are outlined in Figure 1. The NodeOS is responsible for managing local resources such as CPU processing time, link bandwidth and memory storage. On top of the NodeOS, a number of EEs can be installed. On top of each EE, various AAs can be dynamically loaded and executed. The EE is responsible for controlling the access from the AAs to local resources, and limiting resource usage depending on specified policies.
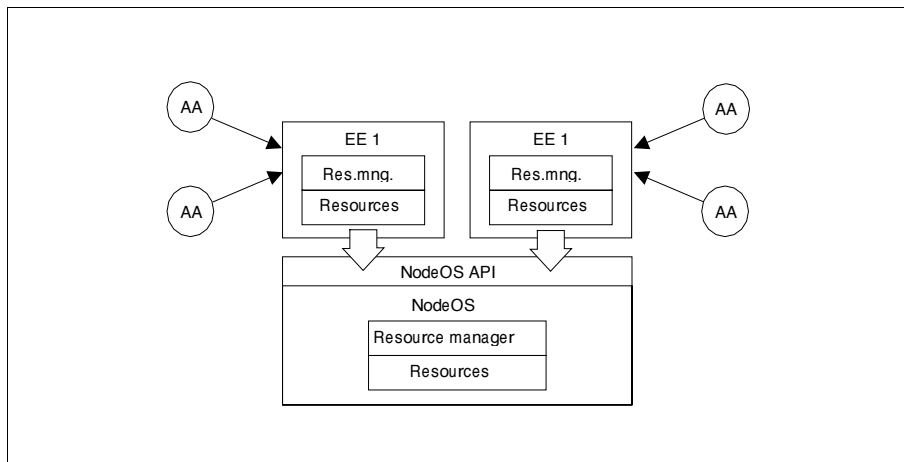

Figure 1. Main components of the DARPA AN architecture.

The NodeOS plays a crucial role in providing access to local node resources, as well as information about resource availability. A NodeOS API is currently being defined [PET 00]. At the moment this API treats four types of resources: computation, memory, communication, and persistent storage. The communication resource is handled through the channel abstraction, which when ready should include QoS support, as well as access to link information such as bandwidth, queue length, and other properties and statistics.

2.3 Mobile Agents and Active Networks

Mobile agents are autonomous pieces of mobile code that travel through the network acting on behalf of their owners. There is a vast amount of literature on the subject, applied to several areas such as manufacturing, e–commerce, network management [HAY 99a] [PAP 00].

The intersection between mobile agent technology and active network technology is the use of mobile code. The capsules of an active network can be

seen as subclasses of mobile agents, specialised for network–related operations. The similarity between active networking and mobile agent concepts is briefly discussed in [HAY 99b] [BRE 99].

Actually, we can consider that the pioneer AN platform actually came from a mobile agent framework: the MØ platform [TSC 93]. MØ is based on the concept of messengers as mobile computational entities that are able to perform any network service. In [TSC 93] it is shown that any protocol based on the classical PDU paradigm – which is the case of basically all the network protocols in operation so far – can be implemented using the messengers paradigm. On the other hand, it is also pointed out that not all the protocols that can be implemented with messengers can be implemented using simple PDUs. The typical case is the one of protocols that are able to evolve their own code on the fly. The author called such protocols "genetic protocols" but admitted to be unable to come up with concrete examples where such protocols would really be necessary in practice. Since then, the search for such protocols still persist in the AN and mobile agent worlds [BOL 00].

If conceptually the border line between mobile agents and AN capsules or messengers is blur, in practice many differences subsist between the two approaches. Mobile agents concentrate mainly on application–level or network management duties, and can typically accomplish much more complex tasks, with richer functionality, than what is generally allowed to capsules or even to out–of–band active code. Moreover, while capsules can be designed to be autonomous and mobile, downloadable modules for programmable networking are typically like plug–ins, and have no autonomy nor mobility once they are installed at the target network element. The architectures for mobile agent platforms and active network platforms differ in the kind of support for code mobility that is offered. Mobile agent platforms tend to concentrate on application level or network value added services, while active network platforms are optimised for transport rather than processing of information.

It is interesting to note that the simplicity of AN platforms compared to agent platforms may enlarge their applicability beyond the network domain they were initially designed for. An example of that can be found in [RUM 00], where the concept of "active network calls" is introduced, which uses AN capsules to support distributed computations. The authors report that efficiency was the reason for the choice of AN capsules instead of mobile agents.

On the other hand, mobile agents can also be used as enabling platform for active networks as pointed out in [BRE 99] [SUG 99]. In this case, instead of the usual plug–ins, it is possible to benefit from full mobile agent functionality in an AN environment.

Finally, it is important to remember that both mobile agent and AN technologies face the same challenges which are mainly security, performance, resource management and interoperability.

2.4 Resource Management

One of the main difficulties encountered in classical adaptation approaches is how to obtain the required information about resource availability, mainly when this information is hidden in a black box network and has to be inferred using only some indirect indications that are observed at the end systems. Using code mobility as provided by mobile agents and active networks, new models for adaptive applications could be envisaged, which can benefit from the possibility to send capsules or agents to certain elements inside the network. These agents can be in charge of collecting information about network conditions, without having to rely on indirect indications or on heavy signalling protocols. Indeed, the idea of sending small pieces of code directly to where the data needs to be treated, instead of exchanging a large amount of data, is one of the main motivations of mobile agent technology, and it can also be applied to mobile code in the case of active networks.

Actually many adaptation mechanisms come from the world of mobile agents. In [JUN 00] an adaptive QoS scheme for MPEG client–server video applications is described. It is based on intelligent agents that reserve network bandwidth and local CPU cycles, and adjust the video stream appropriately. Many agent–based adaptation schemes use artificial market mechanisms [CLE 96]. In [YAM 96] a market model to allocate QoS is applied to a conferencing tool targeted at casual meetings where sudden variations in bandwidth availability require an adaptive QoS control strategy. In [TSC 97] an open resource allocation scheme based on market models is applied to the case of memory allocation for mobile code. [GIB 99] describes a market–based mechanism to set up circuit switching paths with resource reservation. In [BRE 00] the problem of budget planning for mobile agents is addressed, such that they can successfully complete their tasks given their limited budget constraints.

In the domain of end–to–end congestion control in networks, optimisation schemes (see [LOW 99] for an overview) are receiving a lot of attention recently. The algorithms derived from such studies also use a price measure to indicate the level of congestion, and utility functions to quantify users' share of bandwidth. They are therefore closely related to artificial economy systems, but they are restricted to the specific case where bandwidth is the only scarce resource.

Active networks can more easily benefit from the various adaptation mechanisms described, when compared to classical networks, since their code can be dynamically deployed. However, most of the current AN architectures still offer little support for such mechanisms to be implemented in a straightforward manner. For example, ANTS capsules carry a resource limit field that is decremented for the consumption of resources [WET 98]. When it reaches zero the capsule is discarded. However, no mechanisms are specified to manage this field, or to quantify the amount to be decremented.

This situation is rapidly changing though. Recently, [ANA 00] presented a very promising market–based resource management infrastructure for active networks.

It includes a distributed trust–management system which ensures that the market–based policies can be properly enforced in a scalable way.

Another step towards market–based models for AN is the cost model proposed in [NAJ 00], which expresses the trade–off between different types of resources in a quantitative way. However, the recursive approach adopted makes its use more appropriate in the context of reservation–based applications, instead of highly adaptive ones.

A crucial issue for resource management in active networks is the support for incremental deployment. The success of active networks, protocols and services will depend on their ability to complement and interoperate with existing networks in a transparent way, such that active nodes and active functionality can be incrementally deployed. In a hybrid network where only a few nodes might be active, it is important to be able to estimate the resource availability outside an active node. In [SIV 00] an *equivalent link abstraction* is proposed as part of the Protean architecture. Using this abstraction, it is possible to consider a set of non–active nodes as a single link from the point of view of the active nodes involved. Such a virtual link presents changing properties which must be discovered in real–time, such as average rate, delay and packet loss probability. The use of this type of abstraction in the context of artificial market models, as well as its extension to multi–access links, to the case of asymmetric and/or changing routes, etc., are still open issues.

3. Resource Trading Model

In previous work [YAM 00a] we proposed a model for trading resources inside an active node, and applied it to an audio application. This paper is an extended, corrected and updated version of that work. The model was derived from our earlier work on a layered multicast protocol [YAM 00b], therefore it is not a mere abstract proposal but is motivated by concrete application needs.

Our model aims at offering a generic communication abstraction between active network agents, such that different adaptive applications and different resource management policies can be implemented. We are mainly concerned with adaptation to available resources when resources cannot be reserved in advance, either because the router itself does not support reservations, or because the network is heterogeneous and some of the routers along the path offer no QoS support.

In the model, two types of agents communicate to seek an equilibrium. Each agent tries to optimise its own benefits: on one side resource manager agents have the goal of maximising resource usage while maintaining a good performance level. On the other side, user agents try to obtain a better quality/price relation for the resources consumed, and to efficiently manage their own budgets avoiding waste. Both types of agents are implemented as AAs with different privileges, and they communicate such that the resource managers can "sell" resources to the user agents at a price that varies as a function of the demand for the resource.

A currency is introduced into the system to allow for trading of different resource types. This is the basic requirement for artificial economy models such as [BRE 00] [FER 96] [TSC 97] to develop in active nodes, and is also an essential feedback parameter for most algorithms based on optimisation (e.g. [LOW 99]).

The idea is to enable auto–configurable applications and resource managers, such that the code from both types of agents can be dynamically loaded, in order to make them evolve to adapt to new conditions. Figure 2 shows the model as it could be implemented over the DARPA AN architecture. The main implementation difficulty at the moment seems to be the definition of interfaces to local information concerning resource availability, which need to be exported to the active applications. In this matter, we can learn from the mobile agents field in order to model the interactions between resource manager agents and user agents. Software agent communication paradigms such as Agent Communication Languages can be helpful, but still need to be specialised to the AN context.
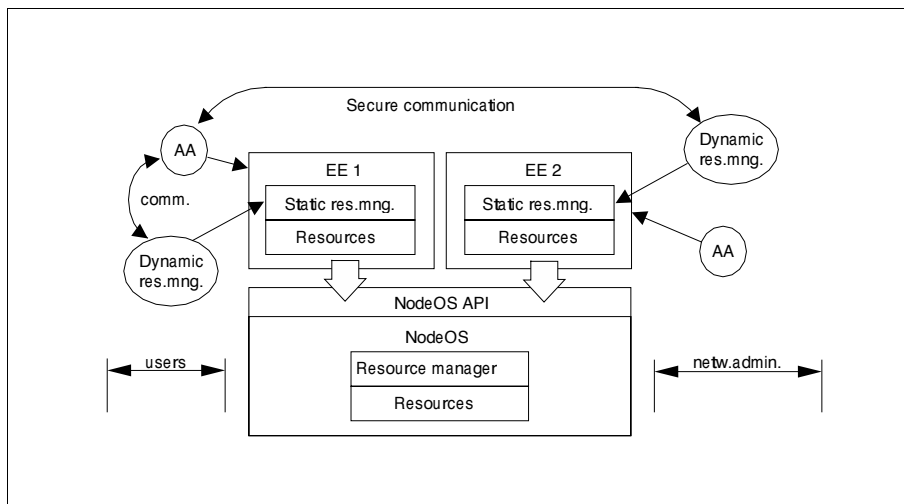

Figure 2. *Resource managers and active applications. Hypothetic placement over the DARPA AN architecture*

Such a model is per se not entirely new, and is in fact a mere simplification of existing artificial economy models which have been mainly applied to the agents domain. For example, the scheme in [GIB 99] is similar to ours, although more complex. The main difference with respect to our model is that we try to adapt it to the specifics of adaptive applications over active networks (no resource reservations), in which reaction time is critical, therefore precluding the use of complex transactions.

The AN market–based infrastructure proposed in [ANA 00] allows the trading of resources between producing agents and consuming agents, which can be mediated by service broker agents. The currency used takes the form of a resource access right. This architecture also bears many similarities with our model. The main difference is that we focus on AAs that can trade resources, while [ANA 00] focuses on the underlying architecture and security mechanisms to support such AAs. Our research efforts are therefore complementary, and we plan to implement our ideas over the platform developed by [ANA 00].

3.1 Resource Manager Agents

Resource managers export resource prices which are a function of the resource utilisation. The utilisation is related to the load, and to the demand for a resource. The function or algorithm used to calculate prices can be shaped to implement desired policies, such as to achieve high utilisation, but also to offer good quality to the users. Resource managers may contain dynamic and/or static code. Some lower level functions which are especially time–critical might be implemented using static code (or even in hardware) while the mobile part would be used to implement more complex policies and to select from a set of pre–existent lower level functions. It is necessary to have the possibility to use dynamic code in order to be able to improve strategies, that is, make them evolve over time, in an active network. Here is one of the places where the alliance between AN and mobile agents can become a must: resource managers can be deployed using mobile agents that are sent by the network manager in order to install new policies. Classical mobile agents for network management can be used for this purpose.

Resource managers are implemented as AAs injected by the network provider, which are executed with network administrator's privileges. There is one class of resource manager for each type of resource concerned. The most relevant classes are: link managers, CPU managers, and memory managers. We will give some more attention to link managers, since they play a crucial role in network congestion control.

Link manager agents also allow abstractions to be made which enable the user applications to adapt to a wide variety of environments in a transparent way. By exporting prices instead of the link internal state information directly, it is possible to hide the specific details of link characteristics while at the same time offering a proper congestion indication to adaptive applications. For example, the price function for a classical point–to–point link would be different from that of a multiple access link, where the local interface load is not a good indication of the actual link utilisation.

An important feature of the link manager price abstraction is that it allows the active applications to deal with non–active nodes in a transparent way. For example, a link manager that implements the equivalent link abstraction [SIV 00] could export a price which would be a function of the estimated average rate, average delay and packet loss probability, which are changing properties in the case of an Internet virtual link. As discussed earlier, this type of abstraction is

crucial for the success of active networks, since the deployment of active nodes will depend on their ability to complement and interwork with existing technologies.

A lot of research is still to be done on how to adjust prices in artificial agent–based economies. An interesting analysis can be found in [MIZ 99]. For this paper, however, this will not be our focus. We will rather concentrate on the user side, assuming that the resource manager agents in the active nodes are able to implement suitable pricing algorithms.

3.2 Capsules as Reactive User Agents

The second type of agents in the model are the active packets themselves, modelled as simple reactive mobile agents. Actually capsules can be regarded as a specialised subclass of mobile agents. They travel to network nodes where they decide when to continue or stop the trip (e.g. stop due to congestion), when to fork new capsules (e.g. in a multicast branch), and which amount of resources to use at each node in order to complete their (generally simplified) tasks.

Capsules have the properties of autonomy and mobility, but need to be simple enough to be executed at the network layer, where performance is often critical. They need to take fast decisions using little data and reasoning, therefore they must be extremely reactive. They can therefore be classified in the reactive agents category. This does not exclude from the model the possibility of using more intelligent mobile agents, jointly with capsules. However, in this paper we focus only on capsules.

Capsules generally represent the user interests, with the goal of attaining the best possible quality at the lowest possible cost. This means that a capsule must be able to make rational decisions on how to spend its limited budget, after consulting resource manager agents for information about the prices of the various resources needed.

Capsules carry a budget that allows them to afford resources in the active nodes, as the resource limit field in ANTS [WET 98]. It looks like a TTL (Time–To–Live) field which is decremented by a number of units for the consumption of a certain amount of resources. When it reaches zero the capsule is discarded. As explained in [WET 98], the budget field must be protected such that the capsules themselves cannot modify it. Again, since capsules are reactive agents, the transactions must be kept simple enough. Instead of the auction mechanisms frequently used by agents [CLE 96], the capsules will most of the time either accept or refuse to pay a given price. Refusing might imply that the capsule simply disappears from the system, because it does not have enough budget to proceed its journey to its destination. Again, we do not preclude the usage of more complex mechanisms implemented by intelligent mobile agents. These agents can also have access to the resource manager interfaces, therefore they can also benefit from the model, however they are not our focus as pointed out earlier in this section.

At the end systems (hosts) conventional programs or intelligent agents can be used to spawn capsules. Since these hosts can dedicate significant amount of resources to information processing, sophisticated strategies can be used to decide which capsules to spawn and when, and how much of the total user's income can be assigned to each capsule (planning). These tasks are then not delegated to the capsules themselves, but can be delegated to more generic mobile agents such as the ones described in [BRE 00].

Here we face the problem of how the budget should be distributed between capsules for one user, and among different users, and also of how to make purchase decisions according to budget constraints. This can be done with the help of utility functions, which quantify the level of user satisfaction for receiving a certain amount of a good, or more generally a combination of goods (market basket) [PIN 98].

According to economics, a typical behaviour of a rational consumer agent is to try to maximise its utility subject to the budget constraints given by its limited income.

For an example where only two goods are involved, the user optimisation problem is then typically expressed as:

Maximise U(x,y)

subject to: $p_x \cdot x + p_y \cdot y = I$

where $U(x,y)$ is the utility function, x and y are the quantities of two goods in their respective units, p_x is the price per unit of x, p_y is the price per unit of y, and I is the user's income.

Such a maximisation process will lead to an equilibrium if the utility function satisfies some properties such as being strictly concave increasing, i.e. the increase in satisfaction is smaller the more items of one good are consumed. The increase in satisfaction that a user obtains from consuming one more item of a given good is called the marginal utility. Within the rational consumer assumption, the marginal utility is a decreasing function of the number of items of a given good. This is called the diminishing marginal utility assumption, and it is directly related to the demand function of a given user for a given good. It means that the more one has from something, the less it is willing to pay to obtain more of it.

A typical utility function is the Cobb–Douglas utility function, given by:

$$U(x,y) = a \cdot \log(x) + (1-a) \cdot \log(y)$$

where $0 \leq a \leq 1$ is a constant which represents the importance the user assigns to x with respect to y.

Applying the method of Lagrange multipliers, the solution of the user optimisation problem with the Cobb–Douglas utility function is given by [PIN 98]:

$$x(p_x) = \frac{a \cdot I}{p_x} \qquad\qquad [1]$$

$$y(p_y) = \frac{(1-a) \cdot I}{p_y} \qquad\qquad [2]$$

The functions $x(p_x)$ and $y(p_y)$ are the demand functions for goods x and y given their current prices per unit p_x and p_y respectively. With demand functions shaped like these, it is possible to chose the quantity of a given resource to consume in order to maximise user satisfaction, given the current price for the resource and the agent's income. Since the price information for each type of resource is available in the active nodes, it is possible for an agent to calculate the amount of resources it can consume at each node, given only an upper bound on the budget per hop it has planned to spend. In the case of capsules, the planning decisions to calculate this upper bound could be performed either at the hosts or by more intelligent mobile agents that would be sent to the active nodes less frequently than capsules. The techniques for such planning are out of scope of this article, and the interested reader can refer to [BRE 00] for more information.

The same reasoning can be easily generalised to an arbitrary number of resources [FER 96]. However, for the purpose of our study this will be sufficient, since we will restrict ourselves to link and memory resources.

Another important characteristic of the Cobb–Douglas function is that it allows us to easily quantify the preferences of different consumers towards one good. For example, given two user agents u_1 and u_2 with the same income I, if agent u_1 has the weight a_1 on its utility function for good x, and agent u_2 has weight $a_2 = n \cdot a_1$ for the same good, we have:

$$x_1(p_x) = \frac{a_1 \cdot I}{p_x}$$

$$x_2(p_x) = \frac{a_2 \cdot I}{p_x} = \frac{n \cdot a_1 \cdot I}{p_x} = n \cdot x_1(p_x)$$

Thus for a given price p_x we have:

$$a_2 = n \cdot a_1 \Rightarrow x_2 = n \cdot x_1$$

It means that if we know that agent u_2 values resource x twice as much as agent u_1, then when faced with the same price, u_2 will get twice as much of x as u_1. This

is an interesting tool for computer network applications, since it provides a quantitative way to provide service differentiation according to the preferences of users. This capability is already known from literature, e.g. [LOW 99] [FER 96]. For example, when trading bandwidth for memory storage, a time–constrained application such as an audio–conference will certainly prefer bandwidth to storage if prices are the same, while a bulk transfer application would probably prefer to store as much information as possible when the links are congested, in order to avoid losses and retransmissions.

In real world economies it is difficult to quantify utilities, but in artificial economies this might be less difficult. Without having to relate artificial currencies to real ones, we could imagine that the maximum budget per unit of time is controlled by a policy server from the network provider that the user is subscribed to, in order to guarantee that users will employ such budget rationally and prevent malicious users from grabbing most of the resources by marking all their capsules with a high budget. This can be compared to the IETF diffserv policies. However, in diffserv only a few predefined classes of service are available, while with such artificial economies, a whole range of classes could appear (and eventually die out), defined by their particular utility characteristics.

4. Congestion Control for a Concast Audio Mixer

We illustrate the use of the trading model through a congestion control scheme for a many–to–one (concast) service. The concast example shows capsules that trade bandwidth for memory when there is congestion.

The term "concast" has been defined in [CAL 00] as a many–to–one service, in opposition to multicast (one–to–many). Figure 3 illustrates this concept. While multicast copies information from one source to many destinations, concast merges information from several sources to one destination. A concast service can be used, for instance, to aggregate feedback in a reliable multicast service, to transmit reception statistics in a multimedia session, to merge information coming from several sources in an auction or tele–voting application, or to combine several real–time streams into one, e.g. to perform audio mixing from several audio sources.

The concast service is faced with the feedback implosion problem, that is, multiple simultaneous sources might congest the path to the single destination, if no congestion control is performed. This problem is aggravated by the fact that many concast flows are used as signalling to support a more robust protocol such as the aggregation of NACK feedback messages for a reliable multicast protocol. Flow control for such signalling messages is often neglected or oversimplified, since the signalling traffic is assumed to be kept small enough when compared to the data traffic. This might lead to poor performance when there is congestion in such a signalling path. We argue that for future sophisticated active services congestion control will be equally important on the signalling and data paths, since the distinction between both tends to become naturally blurred as we approach new

network service composition frameworks which are not simply stack–based as the classical OSI model.

For the purpose of this study, we focus on the case of an audio mixing application. However, the same ideas are applicable to any application making use of the concast service, differing only in the way that packets are combined (merge semantics).


Figure 3. *Multicast (a) versus Concast (b) service abstractions*

In our audio mixing application, several sources generate audio streams in a session. The streams are collected at a single node, which can either record them, play them back or redistribute them to a multicast group. If the receiving node collects all the data and mixes it locally, it might end up with the implosion problem. It is possible to perform mixing operations at every active node, so that the receiver gets a single stream already mixed. However this might cost too much processing and/or memory space in the active routers, so it might become too expensive. Also note that mixing streams delays them, since it is necessary to wait until a minimum number of packets arrive in order to sum up their audio payloads.

The mixing operation is a physical sum of audio samples. Therefore the resulting audio payload has the same number of bytes as each of the original payloads. We assume that there is a maximum amount of signals that can be added up without saturation. Another assumption is that a constant bit rate codec is used, with no silence suppression, such that the signals are generated at a constant rate from the beginning to the end of the session. This initial rate will be altered along the path, as capsules are mixed according to the network conditions.

Note that the mixing operation requires a list of addresses, in order to identify the list of sources already mixed. One might argue that such a list might also occupy bandwidth, since the packet size has to increase in order to accommodate it. But if an estimation of the group size is available, a fixed–size bitstream can be used to hold the list of sources. Adding or removing a source becomes as simple as

setting or resetting a bit in the bitstream. The intersection and union operations can also be implemented as AND and OR binary operations respectively.

We propose to trade bandwidth for processing and memory space to achieve a compromise solution which uses resources efficiently and therefore is able to control congestion. The idea is that when there is congestion at an outgoing interface, the consequently high bandwidth prices will push the users to save bandwidth by performing mixing of data. When the congestion clears up, the users can again benefit from the available bandwidth to avoid the extra delays imposed by the mixing operation.

When a capsule arrives at a node, the first thing it does is to check whether another payload coming from the same source is already buffered. In this case, the arriving capsule cannot mix its payload to the buffered one, which carries earlier samples. Note that this check is in fact an intersection operation to check whether one of the sources mixed in the current capsule is already buffered. If that is the case, the arriving capsule immediately dispatches the buffered payload by creating a new data capsule and injecting it into the local execution environment. This procedure also prevents misordering of packets. As a consequence, at most one packet payload per session is buffered in an active node. We assume that all payloads have the same size.

After that, the capsule can take a decision to either proceed, buffer its payload (for mixing), or discard itself. This decision depends on the budget it carries, and on the prices of the memory and link resources it needs. A number of alternative decision strategies can be envisaged:

Null strategy: Corresponds to the trivial case when no congestion control is performed. In this case the capsule always goes intact to the outgoing interface.

First strategy: If there are other audio payloads from the same group waiting in the memory buffer, it adds its own payload to the buffered one and adds its list of sources to the list of sources already buffered (union operation). Then it terminates execution. If no other audio payloads from the same group are buffered yet, it decides for the cheapest resource: if the price of memory (to buffer the payload for future mixing) is currently lower than the transmission price for the capsule, then it decides to buffer itself; otherwise it decides to move on to the next hop.

Second strategy: The arriving capsule mixes its own data with the buffered one (if existent), then it chooses the cheapest resource, either memory or bandwidth.

The difference between the first strategy and the second strategy is that the first strategy always decides for storage when another payload from the same session is already buffered, while the second strategy always chooses the cheapest resource, independent on the fact that another payload is already buffered or not.

Although the first and second strategies are very naive, they already give quite reasonable results as we will see below. However, their behaviour is sub–optimal and they do not take into account the different preferences for resources.

Third strategy: It first mixes its own data with the buffered one (if existent), as in the second strategy. It obtains a new payload that combines samples from *n* sources (*n* is known from the list of source addresses). It also knows *N*, the maximum number of sources that can be mixed together without saturation. Then it calculates the amounts of link and memory resources according to equations [1] (for link) and [2] (for memory), where the *a* parameter is also carried in the capsule, and expresses its preferences for link resources with respect to memory. It then tries to keep the resources in the proportions obtained, as follows:

If $\quad \dfrac{N-n}{n} > \dfrac{x}{y}\quad$ then store, else move on.

Where *x* is the demand for link resources according to equation [1], and *y* is the demand for memory resources according to equation [2].

Since we have:

$$\frac{x}{y} = \frac{a \cdot p_y}{(1-a) \cdot p_x}$$

the decision is independent on the capsule's budget *I*. The *a* parameter will play a role in the proportion of link resources used with respect to memory resources. The higher *a*, the higher the amount of link resources used, and therefore less capsules will be mixed together, for given memory and link prices.

Note that in all cases, capsules that run out of budget are automatically discarded by the resource managers, thus there is no need to explicitly indicate this operation.

5. Simulations

The audio mixing AA has been simulated with the help of an AN module that we developed for the NS simulator [NS 00]. This module implements a simplified AN architecture consisting of a NodeOS, an EE, and some resource managers. The simulated EE executes capsules written in TCL language.

The topology for the simulations is shown in Figure 4. It consists of *n* sessions of *m* sources and one receiver each. The sessions traverse a bottleneck (link L), so that the capsules in active node N must decide to mix or to proceed intact to the receiver node, according to the prices of link or memory resources available from the resource managers.
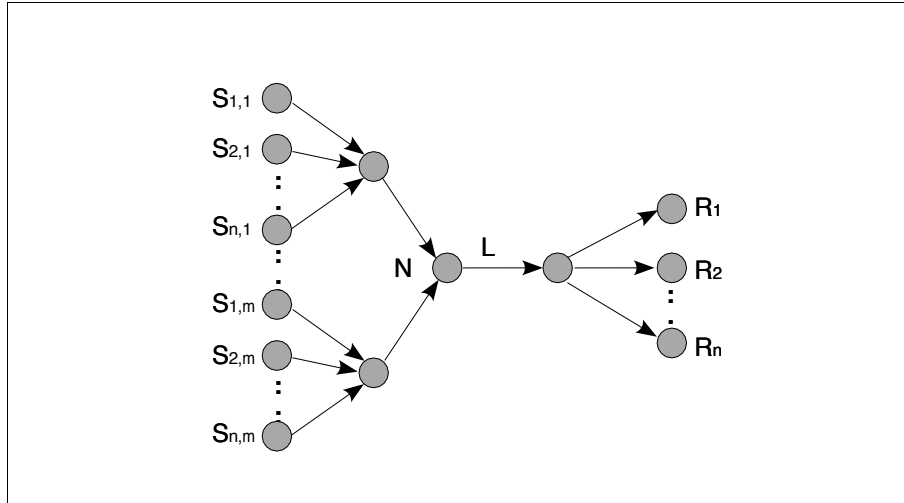
Figure 4. *Topology used in the simulations*

The price function used is based on the one in [TSC 97]:

$$price \ = \ 1000 \cdot \sqrt{\frac{1.01}{1.01 - load}} \ - \ 1000$$

This function is a practical implementation of a convex increasing function that would go to infinite as the load approaches 100\%. It forces the price to rise sharply as we approach high loads, which discourages applications from using a resource when its load is too high. This gives the applications a clear indication of the "dangerous zone" to avoid, while at the same time encouraging a relatively high utilisation.

In the case of the memory manager, the load is given by the ratio between the average number of memory units occupied, and the total number of memory units available to user capsules (which is of course assumed to be much smaller than the actual amount of memory available). The average is calculated using an exponential weighted moving average (EWMA [FLO 93]).

For the link manager, the load is given by the average queue occupancy ratio at the outgoing link interface. This average ratio is obtaining by calculating the average queue length as an EWMA, and then dividing by the maximum queue size. The resulting congestion indication is a bit similar to RED, except that here the binary feedback is replaced by an explicit price indication to the arriving capsules. Note that the actual usage of bandwidth is not taken into account in the

price function, if it does not cause queues to build up. This is therefore a very simplified version of a link manager, but it already serves the purpose of controlling congestion.

We first run an example where 2 sessions are active. There are 5 sources per session, each sending an audio stream of 100kbps to a single receiver, resulting in a total of 1Mbps of traffic arriving at N. The capacity of link L is set to 500kbps.
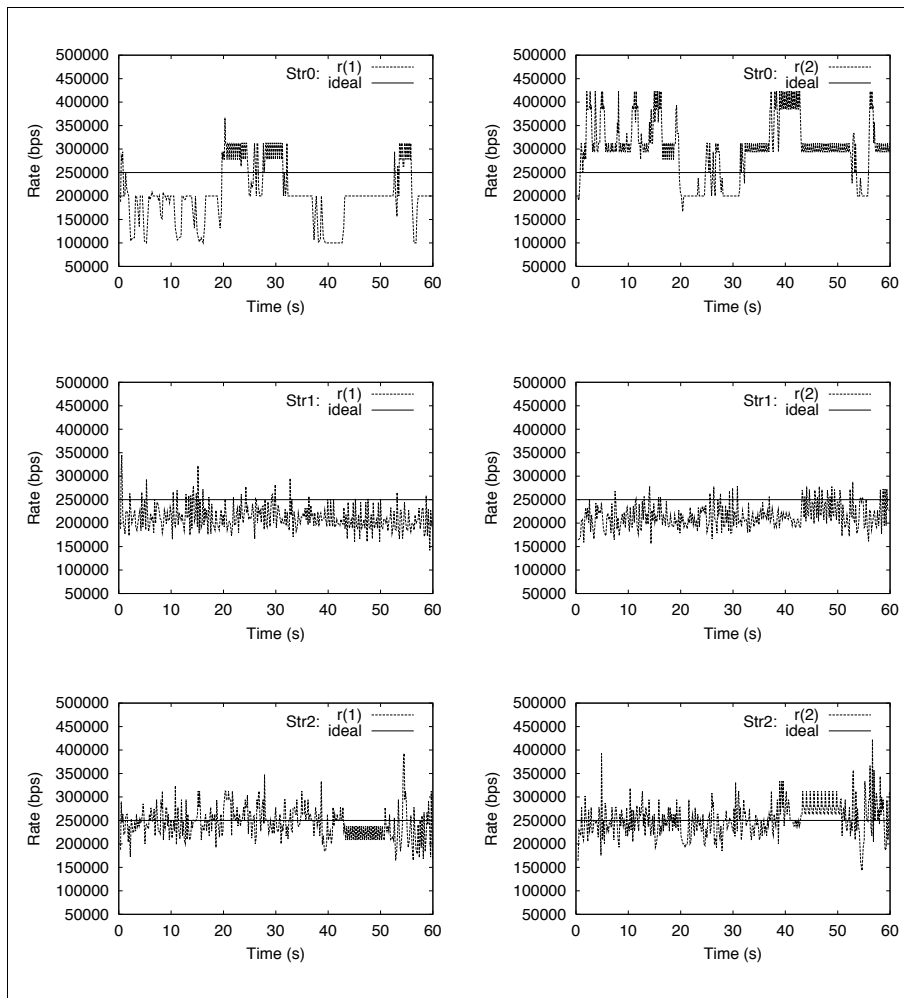

Figure 5. *Evolution of rates in time for 2 sessions, as perceived by their respective receivers. Left: first session (receiver r_1). Right: second session (receiver r_2). Top: no congestion control. Middle: first strategy. Bottom: second strategy*

Figure 5 (top) shows what happens in the trivial case when no congestion control is used. In this case, link overflow occurs at L, and half of the packets are dropped. The remaining packets receive an unequal share of the bottleneck link as we can see in the figure.

Figure 5 (middle) shows what happens when the first strategy is used. First we can notice that the two sessions (left and right) get approximately the same share of the bottleneck. Additionally, no packet losses were observed during the simulation. However, the link is underutilised. This can be explained by the fact that this strategy always favours memory when there is already an item in memory.

The second strategy is a bit more clever (Figure 5, bottom). It always chooses the cheapest resource, either memory or bandwidth. Therefore it is able to grab any bandwidth when it becomes available. Here again, no packet losses occur. However, with this strategy it is not possible to specify different weights for each resource.
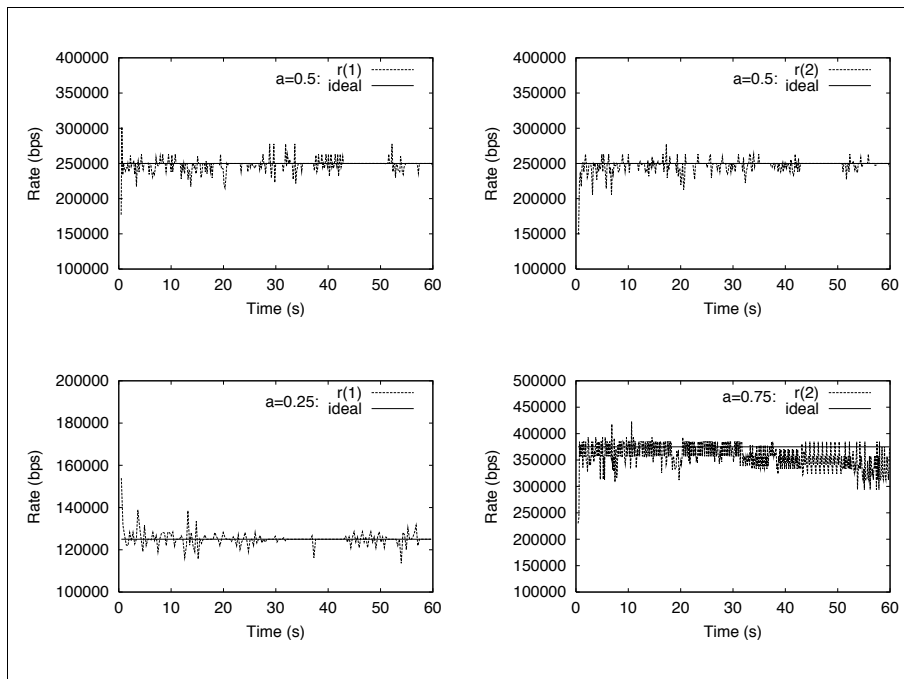

Figure 6. *Evolution of rates in time for 2 sessions using the third strategy. Top: $a_1 = a_2 = 0.5$ for link resources. Bottom: $a_1 = 0.25$, $a_2 = 0.75$*

Now let us look at the third strategy, which allows us to specify utility weights through the a parameter. Figure 6 (top) shows its behaviour when the weights of the two sessions are the same. We see that this strategy is able to share the bandwidth efficiently. It also leads to more stable rates when compared to the previous strategies. The same happens when the weights are different (bottom side of Figure 6), but in this case each session receives link resources in proportion to its respective weight, expressed by the a parameter value.

Until this point only the rates have been shown, since our main goal is to achieve congestion control. Table 1 shows average memory and link parameters taken over the complete duration of the simulations shown in figures 5 and 6. The null strategy (Str.0, no congestion control) occupies most of the link resources and causes the link price to rise, since the link queue is most of the time full. Strategy 1 reduces link utilisation by using memory for mixing, however it does that in an inefficient way when compared to strategy 2, which is able to use more bandwidth while decreasing both link buffer occupancy and memory utilisation. The result is a decrease in link and memory prices.

Strategy 3, when $a_1=a_2$ (Str.3a in Table 1), improves further by keeping the average link queue occupancy at a very low level, in spite of a high bandwidth utilisation. This can be explained by the fact that this strategy tries to find an optimum balance between the usage of memory (which delays packets) and the usage of link resources. It therefore waits for the good moment to send a packet over the link, the moment when link prices are favourable (which in our case corresponds to low queue occupancy).

Column Str.3b in Table 1 shows the average resource parameters when $a_1 \neq a_2$, corresponding to the situation illustrated at the bottom side of Figure 6. There are no significant changes in the parameters at node N, with respect to Str.3a, but the prices are higher, which is an unexpected result, as we would expect the two sessions to concentrate each one on its respective preferred resource, making the load equally distributed. We are still investigating the reason for this discrepancy.

Parameter (average)	Str.0	Str.1	Str.2	Str.3a	Str.3b
memory utilisation (%)	0	7.18	5.14	4.98	5.33
bandwidth utilisation (%)	99.91	84.54	98.18	98.93	96.37
link queue occupancy (%)	93.58	5.83	4.97	1.61	2.98
memory price ($)	0	37.71	26.57	25.62	27.37
link price ($)	2914.6	32.35	27.71	10.17	17.22

Table 1. *Memory and link usage parameters for different strategies*

In order to have a better insight on the mixing procedure, we now look at the number of audio sources carried by each capsule that arrives at its destination. Table 2 shows the percentage of packets that arrived at both receivers, over the total number of packets sent, which is the mix of a given number of sources. The first row ("0 (lost)") represents lost packets. The second row (one source) represents the percentage of packets that arrive intact from the source. The third row (two sources) represents packets that mix samples from two different sources, and so on for the rest of the rows. The columns represent the strategies used. For the null strategy, roughly half of the packets are lost, and the remaining packets arrive intact (no mixing). Strategy 1 is an all−or−nothing strategy: two thirds of the arriving packets contain data from only one source, while one third contains data from all the sources for a given session. The second strategy distributes the mixing effort more evenly.

# sources	Str.0	Str.1	Str.2	Str.3a	Str.3b(r1)	Str.3b(r2)
0 (lost)	49.61	0	0	0	0	0
1	50.39	66.22	39.87	1.54	0.05	60.86
2	0	0.02	33.53	96.16	0.48	39.14
3	0	0	15.80	2.30	1.93	0
4	0	0.43	6.26	0	97.54	0
5	0	33.34	4.53	0	0	0

Table 2. *Average percentage of packets carrying the sum of samples from a given number of sources*

As for the third strategy, column Str.3a of Table 2 shows the results when $a_1=a_2$, corresponding to the simulation result shown at the top side of Figure 6. Columns Str.3b(r1) and Str.3b(r2) show the results for receivers r_1 and r_2 respectively, when $a_1 \neq a_2$ (Figure 6, bottom). We can see that when $a_1=a_2$, most of the packets arrive at the receivers containing samples from two sources mixed together, while for $a_1 = 0.25$ (r_1), most of the packets contain samples from four sources, and for $a_2=0.75$ (r_2), more than one third of the arriving packets contain a mix of only two sources, and the rest only one source (no mixing). This shows that strategy 3 tries to stabilise at a target mixing level, which is characterised by the a parameter. The lowest the a parameter value is for a given session, the highest the mixing level which is achieved.
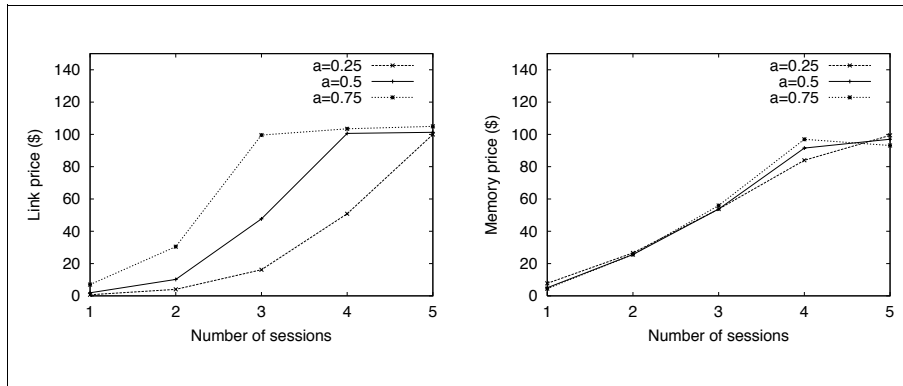
Figure 7. *Average link and memory prices when varying the number of sessions in parallel*

Finally, we vary the number of sessions in parallel using the third strategy, in three separate runs: during the first run all sessions have a=0.25, during the second, a=0.5, and the third, a=0.75. The total average prices for memory and link buffer occupancy are depicted in Figure 7. We can see that the a parameter has a clear influence on the link prices, that increase with a for a given number of sessions, as expected. However it has little influence on the memory prices. This can probably be explained by the fact that, although the application avoids using the memory during a too long period due to delay constraints, at any time each session has at most one packet stored in memory. In the simulations shown, memory does not become a bottleneck, therefore the small impact on prices.

6. Conclusions and future work

We have presented a survey of current research on agent and active network techniques applied to adaptive applications, with special attention to optimisation and market–based approaches. We have also described a model for trading resources inside an active network node, which draws many elements from agent technology. We have applied the model to a concast audio mixing application which trades off link resources against memory in the presence of bottleneck links. The concast application is able to take congestion control decisions locally at each active node, such that no closed loop feedback between source and destination is needed. Using simulations, we studied three different strategies to make a decision on the amount of resources to use: two naive strategies based on the cheapest price, and a strategy that makes use of utility function weights. The results indicate that the first two strategies are already able to make improvements over the case when no congestion control is used, but they use resources inefficiently. The third strategy gives better results, achieving a stable and efficient sharing of resources.

We have several research directions to pursue: the most immediate one is to perform more complex simulations involving multiple node and link types, resource manager types, active and non–active nodes, different user strategies, etc. An implementation over a real active networking platform is also envisaged for the near future. We also plan to investigate the issues of dynamic resource manager upgrade with the help of mobile agents. The precise communication abstractions among the various kinds of agents need further attention too.

7. Acknowledgements

8. References

[ALE 99] D. S. ALEXANDER, J. M. SMITH, "The Architecture of ALIEN", LNCS 1653, Proceedings of IWAN'99, Berlin, Germany, June/July 1999, p. 1–12.

[ANA 00] K. G. ANAGNOSTAKIS ET AL., "Scalable Resource Control in Active Networks", LNCS 1942, Proceedings of IWAN 2000, Tokyo, Japan, October 2000, p.343–357.

[BOL 00] L. BÖLÖNI, D. C.MARINESCU, "Agent surgery: The case for mutable agents", Proceedings of the Third Workshop on Bio–Inspired Solutions to Parallel Processing Problems (BioSP3), Cancun, Mexico, May 2000.

[BRE 00] J. BREDIN ET AL., "A Game–Theoretic Formulation of Multi–Agent Resource Allocation", Proceedings of the 2000 International Conference on Autonomous Agents, Barcelona, Spain, June 2000.

[BRE 99] M. BREUGST ET AL., "Grasshopper – An Agent Platform for Mobile Agent–Based Services in Fixed and Mobile Telecommunications Environments", In [HAY 99a], Chapter 14, p. 326–357.

[CAL 99] K. L. CALVERT (ED) ET AL., "Architectural Framework for Active Networks", (DARPA) AN Working Group, draft version 1.0, July 1999, work in progress.

[CAL 00] K. CALVERT, "Toward an Active Internet", Active and Programmable Networks Mini–conference, Networking 2000, Paris, France, May 2000.

[CAM 99] A. T. CAMPBELL, ET AL., "A Survey of Programmable Networks", ACM SIGCOMM Computer Communication Review, April 1999, p.7–23.

[CHE 98] T. M. CHEN, A. W. JACKSON, "Active and Programmable Networks", Guest Editorial, IEEE Network, May/June 1998, p.10–11.

[CLE 96] S. H. CLEARWATER (ED.), "Market–Based Control – A Paradigm for Distributed Resource Allocation", World Scientific Press, 1996.

[DEN 99] S. DENAZIS ET AL., "Designing Interfaces for Open Programmable Routers", LNCS 1653, Proceedings of IWAN'99, Berlin, Germany, June/July 1999, p. 13–24.

[FER 96] D. F. FERGUSON, C. NICKOLAOU, J. SAIRAMESH, Y. YEMINI,"Economic Models for Allocating Resources in Computer Systems", in [CLE 96], Chapter 7, p.156–183.

[FLO 93] S. FLOYD, V. JACOBSON, "Random Early Detection Gateways for Congestion Avoidance", IEEE/ACM Transactions on Networking, August 1993.

[FLO 00] S. FLOYD ET AL., "Equation–Based Congestion Control for Unicast Applications", Proceedings of ACM SIGCOMM 2000, Stockholm, Sweden, August 2000.

[GIB 99] M.A. GIBNEY, N.J. VRIEND, J.M. GRIFFTHS, "Market–Based Call Routing in Telecommunication Networks Using Adaptive Pricing and Real Bidding", LNAI 1699, Proceedings of the IATA'99 Workshop, Stockholm, Sweden, August 1999.

[GOP 00] R. GOPALAKRISHNAN ET AL., "A Simple Loss Differentiation Approach to Layered Multicast", Proceedings of IEEE INFOCOM 2000, Tel–Aviv, Israel, March 2000.

[HAY 99a] A. L. G. HAYZELDEN, J. BIGHAM (EDS.), "Software Agents for Future Communication Systems", Springer–Verlag, 1999.

[HAY 99b] A. L. G. HAYZELDEN ET AL., "Future Communication Networks using Software Agents", In [HAY 99a], Chapter 1, p.1–57.

[HIC 99] M. HICKS ET AL., "PLANet: An Active Internetwork", Proceedings of IEEE INFOCOM'99, New York, 1999.

[HJA 00] G. HJÁLMTÝSSON, "The Pronto Platform: A Flexible Toolkit for Programming Networks using a Commodity Operating System", Proceedings of IEEE OPENARCH 2000, Tel–Aviv, Israel, March 2000, p. 98–107.

[JUN 00] K. JUN, L. BÖLÖNI, D. YAU, D.C. MARINESCU, "Intelligent QoS Support for an Adaptive Video Service", To appear in the Proceedings of IRMA 2000.

[KHA 00] I. EL KHAYAT, G. LEDUC, "Contrôle de congestion pour la transmission multipoints en couches", JDIR'2000, 4ièmes Journées Doctorales Informatique et Réseaux, Paris, France, November 2000.

[LOW 99] S. LOW, D.E. LAPSLEY, "Optimization Flow Control, I: Basic Algorithm and Convergence", IEEE/ACM Transactions on Networking, 1999.

[MIZ 99] H. MIZUTA, K. STEIGLITZ, E. LIROV, "Effects of Price Signal Choices on Market Stability", $4^{th.}$ Workshop on Economics with Heterogenous Interacting Agents, Genoa, June 1999.

[NAJ 00] K. NAJAFI, A. LEON–GARCIA, "A Novel Cost Model for Active Networks", Proc. of Int. Conf. on Communication Technologies, World Computer Congress 2000.

[NS 00] UCB/LBNL/VINT Network Simulator – ns (version 2), URL http://www–mash.cs.berkeley.edu/ns/

[PAP 00] T. PAPAIOANNOU, "On the Structuring of Distributed Systems: The Argument for Mobility, PhD thesis, Loughborough University, February 2000.

[PET 00] L. PETERSON (ED) ET AL., "NodeOS Interface Specification", (DARPA) AN NodeOS Working Group, draft, January 2000, work in progress.

[PIN 98] R.S. PINDYCK, D.L. RUBINFELD, "Microeconomics", $4^{th.}$ Edition, Prentice Hall International Inc., 1998.

[RUM 00] R. Stainov, J. Dumont, "Distributed Computations by Active Network Calls", LNCS 1942, Proceedings of IWAN 2000, Tokyo, Japan, October 2000, p.45–56.

[SIV 00] R. Sivakumar, S. Han, V. Bharghavan,"A Scalable Architecture for Active Networks", Proceedings of IEEE OPENARCH 2000, Tel–Aviv, Israel, March 2000.

[SUG 99] K. Sugauchi et al., "Flexible Network Management Using Active Network Framework", LNCS 1653, Proceedings of IWAN'99, Berlin, Germany, June/July 1999, p. 241–248.

[TEN 97] D. L. Tennenhouse et al., "A Survey of Active Network Research", IEEE Communications Magazine, Vol. 35, No. 1, pp80–86. January 1997.

[TSC 93] C. Tschudin, "On the Structuring of Computer Communications", PhD thesis, University of Geneva, Switzerland, 1993.

[TSC 97] C. Tschudin, "Open Resource Allocation for Mobile Code", Proceedings of the Mobile Agent'97 Workshop, Berlin, Germany, April 1997.

[VAN 00] B. Vandalore et al., "A Survey of Application Layer Techniques for Adaptive Streaming of Multimedia", to appear in the Journal of Real Time Systems, 2000.

[WET 98] D. J. Wetherall, J. V. Guttag, D. L. Tennenhouse, "ANTS: A Toolkit for Building and Dynamically Deploying Network Protocols", Proceedings of IEEE OPENARCH'98, San Francisco, USA, April 1998.

[YAM 96] H. Yamaki, M.P. Wellman, T. Ishida, "A market–based approach to allocating QoS for multimedia applications", Proceedings of ICMAS'96, Kyoto, Japan, December 1996.

[YAM 00a] L. Yamamoto, G. Leduc, "An Agent–Inspired Active Network Resource Trading Model Applied to Congestion Control", LNCS 1931, Proceedings of MATA 2000, Paris, France, September 2000.

[YAM 00b] L. Yamamoto, G. Leduc, "An Active Layered Multicast Adaptation Protocol", LNCS 1942, Proceedings of IWAN 2000, Tokyo, Japan, October 2000, p.180–194.