ADAPTIVE SYSTEMS - PROGRESS REPORT
JULY 2004

RESEARCH PROGRESS

The Adaptive Systems Project team has conducted research and accomplished design subprojects which all targetted a highly adaptive communication network. With this underlying target in mind the related sub-groups in the Adaptive Systems project have reached new milestones in achieving adaptivity in various levels of the OSI layer of a general communication network, including application layer, transport layer, link layer and physical layer. The following paraghpraphs goes into more detail with an effort to describe the research progress achieved in the last 12 months.

As part of our work, we have looked into enabling new and interesting multi-access services brought about by the simultaneous use of multiple wireless interfaces. These services constitute Bandwidth Aggregation, Mobility/Reliability Support, Resource Sharing and Data-Control Plane Separation. As a first step towards realizing in practice the above mentioned services, we developed a network layer architecture that supports multiple communication paths. The architecture operates at the network layer and introduces minimal changes to the infrastructure and is totally transparent to the applications. We also identified and implemented on an experimental test bed, the various functional components that make up this architecture.

While the architecture can support many different services, we explored in depth one such service provided by the architecture: Bandwidth Aggregation (BAG) in the context of Video and TCP applications. Aggregating bandwidth across multiple interfaces can be used to improve the raw throughput of the client's applications. However, it introduces challenges in the form of packet reordering which can result in excess delay for video applications and congestion control invocation lowering throughputs for TCP applications.

An important aspect of the architecture when providing BAG services is the scheduling algorithm that partitions the traffic onto different wireless interfaces such that the application requirements are met. We proposed one such algorithm Earliest Delivery Path First (EDPF), for video applications. This algorithm ensures that packets meet their playback deadlines by scheduling packets based on the estimated delivery time of the packets. When enabling BAG services for TCP applications, a similar scheduling algorithm that minimizes reordering is needed, however it needs to operate under best-effort conditions. We proposed one such algorithm Packet Pair based Earliest-Delivery-Path-First Scheduling Algorithm for TCP applications (PET). The algorithm estimates the available bandwidth through Packet Pair technique and minimizes reordering by sending packet pairs on the path that introduces the least amount of delay. A buffer management policy (BMP) is also introduced at the client to hide any residual reordering from TCP. Our evaluations show that the proposed scheduling algorithms achieve good bandwidth aggregation under a variety of network conditions.

Another aspect we have looked into in the context of video applications is selective frame discard to cope with cases when adequate bandwidth cannot be reserved on the interfaces. Our proposed content adaptation algorithm, Min Cost Drop (MC-DROP) drops frames based on the impact the frame drop has on meeting future frame deadlines and hence on overall quality of the video. Trace driven simulations using our algorithm show that it outperforms by a good margin other considered approaches in terms of suitably defined metrics that capture overall video quality.

In this section, we present our research progress towards enabling adaptive wireless applications in the past one year. First, we exploit the content properties of applications to enable efficient resource management in a wireless cellular network. Specifically, we develop adaptive scheduling/error control techniques that exploit the diversity in the application data to optimize the system performance. Next, we develop low-cost scalable network monitoring techniques to enable adaptive applications for wireless networks.

A major bottleneck in satisfying the increasing demand for wireless multimedia access is the dynamic error conditions requiring a very high error protection overhead in terms of energy consumption and access delay. We explored ways to use application-level data properties to reduce the error protection cost. We observed that the objects constituting a wireless web-based access differ in terms of the inherent resiliency to residual error, and the object's importance to the overall content quality. Additionally, since the cost of achieving a target residual error level depends on the experienced channel conditions, the access cost can be affected by changing the transmission schedule of objects according to the channel conditions. Based on the above observations, we developed a framework that exploits the application properties, namely content importance and error resiliency, to reduce the cost of error protection by scheduling objects for transmission and adapting the lower-layer error control mechanisms based on the data properties as well as wireless channel conditions. The approach is evaluated under diverse channel conditions using a MATLAB-based simulation environment. Our experimental evaluation show that the proposed approach leads to significant savings in access cost (average 37.4%) with minimal degradation in application quality.

One of the bottlenecks in enabling efficient multimedia applications over wireless is the variability in the wireless network conditions. The dynamic variability in network condition is caused by the inherent nature of wireless channel as well as inter/intra-handoffs between different access technologies based on the service availability. In order to cope with the variations in wireless network, dynamic adaptation techniques are being developed to adjust the multimedia delivery according to the current network condition. Hence, in such adaptive applications, it is critical to estimate/measure the current network condition accurately to enable effective adaptation. In this project, we have started to develop low-cost scalable techniques to estimate/measure current available bandwidth that can be used by the network adaptive applications.

Another adaptivity aspect we explored was the video coders. Standard video coders often use the immediate past frame as a reference frame with motion compensation for video encoding. This mode of coding, called inter coding, achieves substantial compression gains as it exploits the temporal redundancy present in the video sequence. We obtained improvements in video quality obtained by using more than a single reference frame in the inter-coding mode at the cost of increased encoder complexity. A specific case of multiple reference frame video encoding is the dual frame motion compensation, where two reference frames are used. We use dual reference frame motion compensation which uses the immediate past and a long-term past frame as reference frames.

We used the dual frame technique over the Always Best Connected (ABC) network scenario. The ABC service provides a seamless connection between different wireless/wired technologies. For the ABC system, the bandwidth of the connection can vary from 11 Mbps (wireless LAN) to 16 kbps (GPRS). The mobile user switches to a different network if his existing network does not provide service at his current location. For video streams running over such a network, it is important that the video quality is reasonable for the various wireless access technologies. In this context, the most demanding condition for video quality is a switch from Wireless LAN (WLAN) to GPRS, as this is liable to degrade the perceptual video quality tremendously. By retaining a frame when the video encoder operated under high bandwidth and using this as the long term frame in the dual frame technique, it is possible to increase the video quality as measured in Peak Signal to Noise Ratio (PSNR).

With no bandwidth transitions occuring as in an ABC system, a video stream that experiences low bandwidth can still take advantage of the dual frame buffer technique. By allocating bits unevenly among frames, we periodically create a high-quality frame which is then retained as the long-term reference frame for some time. This outperforms a normal dual frame motion compensation scheme in which the long-term reference frames are regular quality frames, and are not allocated any extra rate.

We have shown gains in the average PSNR over 300 frames for the News, Container, Foreman, Claire and Akiyo video sequences.

Energy efficient operation is of paramount importance for battery-powered wireless nodes. In an effort to conserve energy, standard protocols for WLANs have the provision for wireless nodes to ``sleep'' periodically. We considered the problem of optimizing the timing and duration of sleep states for a wireless node fed by a stream of packets with the objective of minimizing power consumption with respect to a QoS constraint. The QoS parameter we have focused on is average delay. We assume a memoryless Bernoulli model for the arriving packets as well as fixed energy cost for waking up and going to sleep. Using a dynamic programming formulation, coupled with a duality argument, we solved the optimization problem numerically. The solutions from the numerical calculations strongly suggest that the optimal policy (that which minimizes average power consumption subject to an average delay constraint) is such that the wireless node should transit to the "sleep" state only when there are no data packets remaining for it to serve. Using a branching process analysis, we were able to derive closed form expressions for the optimal sleep duration, as well as the associated minimal rate of power consumption.

We studied the same problem as mentioned above, for a two-user case. We considered the problem of optimizing the timing and duration of sleep states for two wireless nodes fed by a stream of packets with the objective of minimizing the total power consumption with respect to a QoS constraint namely their average delay. We assume a memoryless Bernoulli model for the arriving packets as well as fixed energy cost for waking up and going to sleep. In the two user case also, we use a dynamic programming formulation to solve the optimization problem numerically. However, unlike the one user case, the solutions from the numerical calculations for the two user case were hard to analyse. The policy as suggested by the numerical calculations was not tractable as well.

We are currently working in developing a heuristic for the two user case. More precisely, we seek a centralised policy, deployable at the access point, which shall govern the sleep/wake-up schedule of the wireless nodes, such that their power consumption is minimised while their packets meet their delay constraints. The aim of the heuristic is to mimic the optimal policy (as suggested by the numeric calculations of the DP) as closely as possible. However, we want it to be scalable as well (unlike the DP approach). We would then compare the performance of our heuristic with the optimal policy obtained from the numeric solutions of the DP.

In another subproject we considered adaptive scheduling and power control algorithms for wireless cellular neworks. We introduced a novel power allocation and scheduling in a DS-CDMA (Direct Sequence – Code Division Multiple Access) wireless communication system with real-time connections. Genarally in a CDMA system all users share the same bandwidth which makes the sytem, interefence limited. Because of this nature of such a system, power control is crucial especially in uplink. Traditionally, for a real-time traffic power control is aimed to have equal received powers at the base station to overcome the near-far effect. We relaxed this design objective and explored a different power allocation and scheduling scheme to achieve maximum total throughput in the system. We formulated the optimization problem which achieves the optimum power allocation and scheduling scheme, and then reduced the optimization problem to a linear programming problem and solved it. We used our results to discuss and compare our scheme with the traditional perfect power controlled CDMA system, to show that the introduced scheme outperforms traditional perfect power controlled CDMA. The analysis was carried out for both uplink and downlink traffic. Mathematica based simulations were used to justify the theorethical findings.

In another part of the project, we have characterized workloads in Public-Area Wireless Networks (PAWNs), and have shown that: (i) user loads are often time varying and location-dependent; (ii) user load is often unevenly distributed across access points (APs); and (iii) the load on the APs at any given time is not well correlated with the number of users associated with those APs. Administrators in such networks thus have to address the challenge of unbalanced network utilization resulting from unbalanced user load, and also guarantee its users a minimum level of quality of service (e.g., sufficient wireless bandwidth).

We have addressed the challenges of improving PAWN utilization and user bandwidth allocation using a common solution: dynamic, location-aware adaptation. By adaptively varying the bandwidth allocated to users in the wireless hop within certain bounds, coupled with admission control at each AP, the network can accommodate more users as its capacity changes with time. Further, by adaptively selecting the AP that users associate with, the network can relieve sporadic user congestion at popular locations and increase the likelihood of admitting users at pre-negotiated service levels.

We present the problem of first-hop wireless bandwidth allocation as a special case of the well-known online load balancing problem, and have proved that the general online problem of finding an optimal assignment of users to APs in an arbitrary network with arbitrarily sized user bandwidth requests is NP-complete. We therefore developed three online heuristic algorithms for first-hop bandwidth allocation. We describe how these algorithms enable the network to transparently adapt to user demands and balance load across its access points.

Finally, we have evaluated the effectiveness of these algorithms on improving user service rates and network utilization via simulation, incorporating real workloads from campus, conference, and corporate environments. We show that our algorithms improve the degree of balance in the system by over 45% and allocate over 30% more bandwidth to users in comparison to existing schemes that offer little or no load balancing.

The performance of popular Internet Web services is governed by a complex combination of server behavior, network characteristics and client workload -- all interacting through the actions of the underlying transport control protocol (TCP). Consequently, even small changes to TCP or to the network infrastructure can have significant impact on end-to-end performance, yet at the same time it is challenging for service administrators to predict what that impact will be.

In this project we designed and implemented a tool called "Monkey" that helps address such questions. Monkey collects live TCP trace data near a server, distills key aspects of each connection (e.g., network delay, bottleneck bandwidth, server delays, etc), and then is able to faithfully replay the client workload in a new setting. Using Monkey, one can easily evaluate the effects of different network implementations or protocol optimizations in a controlled fashion, without the limitations of synthetic workloads or the lack of reproducibility of live user traffic. Using realistic network traces from the Google search site, we showed that Monkey is able to replay traces with a high degree of accuracy and can be used to predict the impact of changes to the TCP stack.