By Cormac J. Sreenan and
Kobus van der Merwe,
Guest Editors

ENTERTAINMENT NETWORKING

Demanding new Internet applications have produced a radical and fundamental shift in the nature of entertainment.

The ongoing convergence of the Internet and telecommunications and entertainment networks continues to provide new and exciting opportunities for end users, application developers, content providers, and network operators. Traditional entertainment networks developed separately from networks used for communication and commerce. For example, until recently, cable TV networks were one-way broadcast infrastructures aimed solely at distributing analog TV signals. Today, much of this cable infrastructure has been upgraded to two-way all-digital facilities, providing one avenue for broadband access to the Internet. Among the many applications making use of broadband connections are telephony and streaming applications, some carrying content similar in concept, if not yet in quality, to the traditional telephone and TV networks.

ILLUSTRATION BY Lisa Haney

Convergence among networks is becoming a commercial reality, and the Internet Protocol (IP) is widely accepted by industry as the common underlying network protocol of choice. This is due to IP's fundamental ability to support a variety of applications along with its ubiquitous worldwide deployment.

In the past few years, users and vendors alike have started to use the Internet for entertainment and recreational purposes, driven and facilitated by greater penetration of the broadband access infrastructure. They're downloaded media content (such as music and video) for subsequent consumption, distributed live TV-style events, and delivered and supported multiplayer online/networked games. Using content-distribution networks to speed access to popular Web sites is itself popular, with several established companies maintaining thousands of edge servers worldwide. But when it comes to distributing and sharing content, peer-to-peer overlay networks are also gaining popularity, most notably for music distribution.

The need to support worldwide distribution of digital multimedia, especially for live events, is a major challenge for IP network designers especially over how to achieve scalable, cost-effective delivery. For example, Internet TV, commonly known as IPTV, has gained attention as a new service for broadband customers. However, many obstacles remain before the vision of access to any TV channel "anywhere, anytime" is realized. The key ones are how to support management of content, digital rights, and network/server resources on a global scale. Massive multiplayer online games are a commercial success in terms of revenue for their vendors and growth in numbers of subscribers, but their evolution hinges on a set of key research questions (such as how to support low-latency and high-bandwidth operation within global IP networks).

Many entertainment applications involve more stringent network requirements than traditional applications, especially in terms of tolerance for latency, a theme explored by several of the articles in this section. As access speeds increase and network support for entertainment applications improves, new applications will likely emerge and introduce other significant challenges for IP network designers; one oft-cited example is augmented virtual reality, which demands significantly greater bandwidth than is commonly available today through the public Internet. These observations and related emerging technologies motivated us to focus here on entertainment networking and the recreational use of IP networks. The articles cover four application areas: online multiplayer games; networked music performance; collaborative media streaming; and peer-to-peer media streaming. Collectively, they provide a broad perspective on state-of-the-art networked entertainment applications, identify important research challenges and opportunities, address selected technical problems, and describe prototype systems, along with illustrative experimental results.

Kuan-Ta Chen et al. start us off by explaining the relationship between network performance and game-player satisfaction. Using results from a study of a subscription-based massive multiplayer online game—Shen Zhou Online—they show that the duration of users' game play falls off in response to increased network latency, variation in that latency, and network loss rates. They also provide a model that helps predict user-play duration if network performance is known. This model is likely to be of special interest to network operators and game-hosting companies, as it allows trade-offs in network latency and loss rates to be performed with an understanding of their likely effect on user/customer behavior.

Mark Claypool and Kajal Claypool explore why the effects of network latency are dependent on the particular actions a player takes in an online game. They categorize player actions in terms of their deadlines and preferred precision, with actions (such as sniper shooting) in first-person shooter games falling into the most latency-sensitive class. This taxonomy is supported by experimental results that measure the relationship between latency and game-dependent quantitative performance metrics—for a first-person shooter game, a third-person action game, and a sports game. The article should be of particular interest to game designers seeking to understand user tolerance of latency, network designers assessing quality-of-service demands, and users deciding which broadband access package to use.

Jeremy Brun et al. consider how to design networked games to support users located around the world. In this context they explore distributed servers and the relationship between network latency and fairness in networked games. Differences in latency can lead to inconsistency in game state, giving certain users an unfair advantage over others. This problem is exacerbated in situations in which decisions concerning the game state are made on distributed servers, rather than through the more conventional centralized server architecture. The authors identify two techniques that are useful for mitigating these effects: trading inconsistencies and judiciously selecting the location of the distributed servers. The article should be of particular interest to companies interested in scaling-up game servers for use by a truly global subscriber base.

Zefir Kurtisi et al. offer a vision for interactive music performance in which individual musicians communicate over a wide-area network (such as the Internet). Focusing on one key issue—ensuring that the extremely tight delay bound (30msec) necessary for lis-

tening to music—they explore how it might be achieved in a wide-area network. They explore the anatomy of the processing and data paths that determine minimum processing and buffering delays, yielding a value for the maximum acceptable network latency. They use experimental results from several networks to demonstrate that the latency bounds necessary for networked music performance can be achieved in practice.

Reza Rejaie looks into a future in which everyone will be able to broadcast their own TV shows, along with interactive online multimedia forums, video blogs, and virtual experiences. Although peer-to-peer overlay networks involving the cooperation of large numbers of computers seem attractive as a vehicle for distributing live video over the Internet, several key research challenges must be solved first: how to deal with differences in the access bandwidth capabilities of users; how to mitigate the effect of users deciding to opt-out of the video distribution; and how to ensure that finite network bandwidth is used efficiently. Rejaie outlines a possible approach to their solution based on the concept of a mesh-based peer-to-peer network.

Finally, Verena Kahmann et al. challenge the conventional notion of media streaming as a passive activity involving lone users, introducing the idea of collaborative streaming. In collaborative streaming, several users interact as an integral part of the streaming experience; the illustrative examples are of a family watching a movie together and a distance-learning scenario. The related technology involves several challenges, especially how to share and control the session state. The authors describe the architecture they have developed for collaborative streaming based on the use of the standard Session Initiation Protocol, which is known mostly for its use in voice-over-IP.

Together, these articles provide a broad and timely perspective on the emerging area of entertainment networking, emphasizing exciting new applications and the research challenges faced by IP network designers. Over the next five years, as they are resolved by the research community and embraced by industry, we can expect a fundamental shift in the nature of all entertainment. The focus will move from users passively viewing TV broadcasts to users being able to interact as if they were in a global village, accessing a vast repository of TV and video. C

Cormac J. Sreenan (cjs@cs.ucc.ie) is a professor of computer science in the Department of Computer Science at University College Cork, Cork, Ireland.
Kobus van der Merwe (kobus@research.att.com) is a researcher at AT&T Labs-Research, Florham Park, NJ.