

International Workshop on Network and Operating Systems Support for Digital Audio and Video

Sponsored by ACM SIGMultimedia

NOSSDAV 2006

Newport, Rhode Island

May 22-23, 2006



- Home Page
- Call For Papers
- Program Committee
- Local Information
- Author Instructions
- Workshop Program
- Registration

This page contains the titles and abstracts of each presentation.

Another page lists the NOSSDAV [schedule](#) of talks, including the Keynote speaker and Industry Panelists.

Previous NOSSDAVs

Title	Abstract
<p>GONE: an Infrastructure Overlay for Resilient, DoS-Limiting Networking</p> <p>Xiaoming Fu (University of Goettingen, DE); Jon Crowcroft (University of Cambridge, UK)</p>	<p>With today's penetration in volume and variety of information flowing across the Internet, data and services are experiencing various issues with the TCP/IP infrastructure, most notably availability, reliability, and mobility. So far the proposed approaches have focused on applying application-layer routing and path monitoring for reliability and enforcing stateful packet filters in hosts or network to protect against Denial of Service (DoS) attacks. Each of them solves its own aspect of the problem, trading scalability for availability and reliability among a relatively small set of nodes, yet there is no single overall solution addressing these issues in a large scale available. We propose an alternative overlay network architecture by introducing a set of generic functions in network edges and end hosts. We conjecture that the network edge constitutes a major source of DoS, resilience and mobility issues to the network, and propose a new solution to this problem, namely the General Internet Signaling Transport (GIST) Overlay Networking Extension, or GONE. The basic idea of GONE is to create a half-permanent overlay mesh consisting of GONE-enabled edge routers, which employs capability-based DoS prevention and forwards end-to-end user traffic using the GIST messaging associations. GONE's use of GIST on top of SCTP allows multi-homing, multi-streaming and partial reliability, while only a limited overhead for maintaining the messaging association is introduced. In addition, upon the services provided by GONE overlays, hosts are identified by their unique host identifiers independent of their topologies location, and simply require (de-)multiplexing instead of the traditional connection management and other complex functionality in the transport layer. As a result, this approach offers a number of advantages for upper layer end-to-end applications, including intrinsic provisioning of</p>

	<p>resilience and DoS prevention in a dynamic and nomadic environment.</p>
<p>NSYNC: Network Synchronization for peer-to-peer streaming overlay construction</p> <p>Hongbo Jiang (Case Western Reserve University, US); Shudong Jin (Case Western Reserve University, US)</p>	<p>In peer-to-peer streaming applications such as IP television and live shows, a key problem is how to construct an overlay network to provide high-quality, almost real-time media relay in an efficient and scalable manner. Much work has focused on the construction of tree and graph network topology, often based on the inference of network characteristics (such as delay and bandwidth). Less attention has been paid to improving the liveness of media delivery, and to exploiting the flexibility of applications to construct better overlay networks. We propose the NSYNC, an ongoing work on constructing low-latency overlay networks for live streaming. It aims at solving the following problems. In typical applications, peers must buffer a portion of a real-time event, e.g., for at least a few seconds. Thus, it introduces both (1) delay, especially long delay for peers that are many hops away from the origin servers, and (2) partial ordering between the peers. With NSYNC, the users (such as media players) can slightly increase or decrease the speed of playing (displaying) media. Thus, the peers in a network can be synchronized to achieve two effects. First, late peers can catch early peers (and the origin server) such that the entire peer networks improve liveness. Second, the client/server roles between a pair of neighboring peers can be reversed, allowing opportunities for constructing more efficient overlay networks. NSYNC can be used in various peer-to-peer streaming systems. Our preliminary results show NSYNC can be effective.</p>
<p>Optimistic Load Balancing in a Distributed Virtual Environment</p> <p>Roman Chertov (Purdue University, US); Sonia Fahmy (Purdue University, US)</p>	<p>Distributed virtual environments such as massive multi-player games require support of multiple servers to balance computational load. This paper investigates the architecture of a unified environment where the virtual online world is not partitioned according to rigid boundaries, but according to an adaptive paradigm. Since it is difficult to develop an efficient load balancing algorithm for a unified environment, we propose an optimistic scheme that quickly converges. The cost of frequent migrations is minimized by using stateless servers/connections, and following a push/push data model. We analyze the computational time costs of such a system and give simulation results to gauge its performance. The simulation results confirm that our load balancing scheme is efficient and can support large numbers of clients.</p>
<p>Harmonic Interleaving: File System Support for Scalable Streaming of Layer Encoded Object</p>	<p>We propose a novel file system technique to efficiently support scalable streaming of layer encoded content. From file system's point of view, scalable streaming introduces another dimension of complexity in disk scheduling. This is particularly because when a subset of information is retrieved, the playback does not necessarily coincides with the sequential access. The current file structure and the file system organization for real-time multimedia application leaves much to be desired to efficiently support scalable streaming service. We propose a new file organization technique called "Harmonic Interleaving" whose basic idea is to cluster the frequently accessed layers together. Harmonic Interleaving scheme is conceptually very simple but which layers need to be clustered together is yet to be answered. Further, proving its effectiveness in practical situation requires rather serious treatment. We successfully develop elaborate</p>

<p>Sooyong Kang (Hanyang University, KR); Youjip Won (Hanyang University, KR); Seunghyun Roh (Hanyang University, KR)</p>	<p>analytical models for a number of different file organization techniques, which are Progressive Placement, Interleaving Placement and Harmonic Interleaving Placement, and perform in depth investigation of various file organization techniques under varying workload conditions. The models developed in this work enables us to predict the performance and efficiency of the storage system in scalable streaming environment. We find that the Harmonic Interleaving yields the most promising performance in scalable streaming environment. The result of our work provides not only the analytical techniques to predict the performance of a given file organization technique but also the useful guidelines for ISP or real-time multimedia contents providers on how to organize their contents in physical system.</p>
<p>Evaluating dead reckoning variations with a multi-player game simulator</p> <p>Wladimir Palant (University of Oslo, NO); Carsten Griwodz (University of Oslo, NO); Pål Halvorsen (University of Oslo, NO)</p>	<p>One of the most difficult tasks when creating an online multi-player game is to provide the players with a consistent view of the virtual world despite the network delays. Most current games use prediction algorithms to achieve this, but usually it does not go beyond applying the DIS dead reckoning algorithm proposed in the mid-90s. In this paper we introduce a simulator called GLS that allows us to evaluate different aspects of DIS and its variations. We examine the impact of prediction and clock synchronization on game consistency. We also evaluate the convergence algorithm we introduce here. Furthermore we look into ways for compensating increasing delays to keep the player's view of the game state sufficiently consistent with other players.</p>
<p>A hybrid thin-client protocol for Multimedia Streaming and Interactive gaming applications</p> <p>Davy De Winter (Hogeschool Gent - Ghent University, BE); Pieter Simoens (University of Ghent, BE); Lien Deboosere (Ghent University, BE); Filip De Turck (Ghent University, BE); Joris Moreau (Hogeschool Gent - Industriële wetenschappen,</p>	<p>Despite the growing popularity and advantages of thin-client systems, they still have some important shortcomings. Current thin-client systems are ideally suited to be used with classic office-applications but as soon as multimedia and 3D gaming applications are used they require a large amount of bandwidth and processing power. Furthermore, most of multimedia and 3D applications heavily rely on the Graphical Processing Unit (GPU). Due to the architectural design of thin-client systems, these applications cannot profit from the GPU resulting in slow performance and bad image quality. In this paper, we propose a thin-client system which addresses these problems: we introduce a realtime desktopstreamer using a H264/AVC-based videocodec to stream the graphical output of applications after GPU-processing to a thin-client device, capable of decoding a videostream. We compare this approach to a number of popular classic thin-client systems in terms of bandwidth, delay and image-quality. The outcome is an architecture for a hybrid protocol, which can dynamically switch between a classic thin-client protocol and realtime desktopstreaming.</p>

<p>BE); Bart Dhoedt (Ghent University, BE); Piet Demeester (Ghent University, BE)</p>	
<p>AMTrac: Adaptive Meta-caching for Transcoding</p> <p>Dongyu Liu (George Mason University, US); Songqing Chen (George Mason University, US); Bo Shen (HP Labs, US)</p>	<p>The increase of aggregate Internet bandwidth and the rapid development of 3G wireless networks demand efficient delivery of multimedia objects to all types of wireless devices. To handle requests from wireless devices at runtime, the transcode-enabled caching proxy has been proposed and a lot of research has been conducted to study online transcoding. Since transcoding is a CPU-intensive task, the transcoded versions can be saved to reduce the CPU load for future requests. However, extensively caching all transcoded results can quickly exhaust cache space. Constrained by available CPU and storage, existing transcode-enabled caching schemes always selectively cache certain transcoded versions, expecting that many future requests can be served from the cache while leaving CPU cycles for online transcoding for other requests. But such schemes treat the transcoder as a black box, leaving no room for flexible control of joint resource management between CPU and storage. In this paper, we first introduce the idea of meta-caching by looking into a transcoding procedure. Instead of caching certain selected transcoded versions in full, meta-caching identifies intermediate transcoding steps from which certain intermediate results (called metadata) can be cached so that a fully transcoded version can be easily produced from the metadata with a small amount of CPU cycles. Achieving big saving in caching space with possibly small sacrifice on CPU load, the proposed meta-caching scheme provides a unique method to balance the utilization of CPU and storage resources at the proxy. We further construct a model to analyze the meta-caching scheme and propose AMTrac, Adaptive Meta-caching for TRAnsCoding, which adaptively applies meta-caching based on the client request patterns and available resources. Experimental results show that our proposed AMTrac can significantly improve the system throughput over existing approaches.</p>
<p>On the Effects of Multi-hop Buffering and Adaptation for Video-Based Sensor Networking Applications</p> <p>Jie Huang (Portland State University, US); Wu-chi Feng (Portland State University, US); Wu-chang Feng</p>	<p>As video sensor networks become more widely deployed, mechanisms for buffering and adaptively transmitting video data within the network are necessary because of their generally large resource requirements compared to their scalar counterparts. In video sensor networks, many video sensors can inject video into the network at the same time as well as store and forward data for other sensor nodes. The effects of video buffering and adaptation in such multi-hop networks are not that well understood. In this paper, we propose a “bucket-brigade-style” multi-hop buffering and adaptation framework for video-based sensor networking applications. We explore several approaches in this framework and compare their performance with traditional IP-based video streaming technologies. Our experiments show that the proposed framework outperforms traditional technologies in bandwidth efficiency, video quality,</p>

<p>(Portland State University, US); David Romano (Intel, US)</p>	<p>bandwidth sharing fairness, and energy efficiency.</p>
<p>The Fun of using TCP for an MMORPG</p> <p>Carsten Griwodz (University of Oslo, NO); Pål Halvorsen (University of Oslo, NO)</p>	<p>Massive multi-player online games have become a popular, fast growing, multi-million industry with a very high user mass supporting hundreds or thousands of concurrent players. In many cases, these games are centralized and every player communicates with the central server through a time-critical unicast event stream. Funcom's Anarchy Online is one of these; it is based on TCP. We find that its kind of traffic has some interesting properties that inspire changes to protocol or architecture. In these game streams, using TCP is not TCP-friendly, using TCP does not have to be slower than using UDP, and only repeated timeouts ruin game experience. Improving the latter in the sender implementation does not impose any remarkable penalty on the network. Alternatively, a proxy architecture for multiplexing could save about 40% resources at the server, allow congestion control to work and also reduce the lag of the game.</p>
<p>Assigning Game Server Roles in Mobile Ad-hoc Networks</p> <p>Oliver Wellnitz (Technische Universitaet Braunschweig, DE); Lars Wolf (TU Braunschweig, IBR, DE)</p>	<p>Over the last couple of years, multi-player games have become more and more popular. Additionally, new mobile devices now have sufficient resources to play these multi-player games in mobile and wireless networks. As the classic centralized game server design is unsuited for mobile ad-hoc networks, a group of nodes can take the role of a distributed game server. This group of nodes can provide the necessary redundancy which is needed in the dynamic environment of a mobile ad-hoc network. In this paper we present a modified dominating set algorithm for server selection which uses local information to determine well-suited nodes from the group of players. Our algorithm uses three phases (discovery, determination and marking) to calculate an initial server set and adjusts to network changes during the game. We implemented our server selection algorithm in NS-2 and evaluated its behaviour in two different realistic scenarios for mobile games (schoolyard and train) as well as in an artificial stress scenario.</p>
<p>A Platform for Dynamic Microcell Redeployment in Massively Multiplayer Online Games</p> <p>Bruno Van Den Bossche (Ghent University, BE); Tom Verdickt (Ghent University, BE); Bart De Vleschauwer (Ghent University, BE);</p>	<p>As Massively Multiplayer Online Games enjoy a huge popularity and are played by tens of thousands of players simultaneously, an efficient software architecture is needed to cope with the dynamically changing loads at the server side. In this paper we discuss a novel way to support this kind of application by dividing the virtual world into several parts, called microcells. Every server is assigned a number of microcells and by dynamically redeploying these</p>

<p>Stein De Smet (Ghent University, BE); Stijn De Mulder (Ghent University, BE); Filip De Turck (Ghent University, BE); Bart Dhoedt (Ghent University, BE); Piet Demeester (Ghent University, BE)</p>	<p>microcells when the load in a region of the world suddenly increases, the platform is able to adapt to changing load distributions. The software architecture for this system is described and we also provide some evaluation results that indicate the performance of our platform.</p>
<p>Intra-Stream Encoding for Multiple Depth Streams</p> <p>Sang-Uok Kum (Univ. of North Carolina - Chapel Hill, US); Ketan Mayer-Patel (University of North Carolina, US)</p>	<p>Using multiple cameras, a dynamic real world scene can be captured, represented, and streamed for realistic interaction in a 3D immersive system. One method to represent the dynamic environment is to use multiple depth streams -- streams with color and depth information. However the captured data is too large to be streamed uncompressed. Fortunately, these data exhibit spatial and temporal coherence that can be used for compression. In this paper, we examine the initial step in compressing multiple depth streams. To use spatial coherence between streams to compress multiple depth streams, base reference streams -- intra-streams -- have to be encoded without the use of other depth streams. We investigate encoding of this base depth stream using only temporal coherence. The proposed algorithm extends video stream encoding techniques to encode color and depth. We show that for encoding a depth stream which require high quality, such as base reference streams, using separate motion vectors to encode color and depth performs better than using a single motion vector.</p>
<p>Design and Evaluation of 3D Video System Based on H.264 View Coding</p> <p>Kalva Hari (Florida Atlantic University, US); Liam Mayron (Florida Atlantic University, US); Lakis Cristodoulou (Florida Atlantic University, US); Oge Marques (Florida Atlantic University, US); Borko Furht</p>	<p>Recent advances in video compression and 3D displays have necessitated a further understanding and development of 3D video coding algorithms. The emergence of low cost autostereoscopic displays is expected to drive the growth of 3DTV services. This paper discusses key issues that affect the quality of 3D video experience on autostereoscopic displays. The stereo views can be encoded at different quality without affecting the perceptual quality of the 3D video. We examine the bounds of this asymmetric stereo view compression with a user study. The H.264/AVC video coding algorithm was used to compress each view. This paper also presents the design and development of a modular video player with stereoscopic and multi-view capabilities including a discussion of useful tools for accelerating development and enhancing flexibility. The player implements the 3D image reconstruction caused by binocular fusion. A discussion of experimental results and recommendations for improvements of stereoscopic image quality in terms of 3D video enhancement, processing and acquisition, deployed video algorithms, as well as smart encoders and decoders are also presented in this paper.</p>

<p>(Florida Atlantic Univ., US)</p>	
<p>FGS-MR: MPEG4 Fine Grained Scalable Multi-Resolution Video Encoding for Adaptive Video Streaming</p> <p>Siddhartha Chattopadhyay (University of Georgia, US); Kang Li (University of Georgia, US); Suchendra Bhandarkar (University of Georgia, US)</p>	<p>The MPEG-4 Fine Grained Scalability (FGS) profile aims at scalable video encoding, in order to ensure efficient video streaming in networks with fluctuating bandwidths. In order to allow very low bit rate streaming, the Base Layer of a FGS video is encoded at a very low bit rate, consequently, at very low video quality. In this paper, we propose FGS-MR, which uses content aware multi-resolution video frames to obtain better video quality for a target bit rate, compared to MPEG-4 FGS Base Layer video encoding. MPEG-4 FGS Base Layer video encoding is optimized using Adaptive Quantization (FGS-AQ). Given a quantitative quality of video (e.g. PSNR value) desired, FGS-MR leads to lower bit rate video encoding compared to FGS-AQ. In addition, FGS-MR is an integrated approach that requires only encoding side modification, and is transparent to the decoder, as opposed to FGS-AQ. We have shown various potential mobile applications which can benefit significantly by using FGS-MR, such as video conferencing applications.</p>
<p>Video Authentication for H.264/AVC using Digital Signature Standard and Secure Hash Algorithm</p> <p>Nandakishore Ramaswamy (Qualcomm Inc, US); K. R. Rao (University of Texas at Arlington, US)</p>	<p>This paper proposes a hard video authentication and sender verification scheme for video sequences compressed using H.264/AVC Main Profile by using digital signatures and cryptographic hash. Features from the transform domain are used as the authentication data for a macroblock. The algorithm can detect the cause of authentication failure and point out the location of the tampered frames in case of frame tampering. Results demonstrate that the proposed scheme is highly robust to temporal and spatial manipulations</p>
<p>Redundancy-Controllable Adaptive Retransmission Timeout Estimation for Packet Video</p> <p>Ali Begen (Georgia Institute of Technology, US); Yucel</p>	<p>Time-constrained error recovery is an integral component of reliable low-delay video applications. Regardless of the error-control method adopted by the application, unacknowledged or missing packets must be quickly identified as lost or delayed, so that necessary timely actions can be taken by the server/client. Historically, this problem has been referred to as retransmission timeout (RTO) estimation. Earlier studies show that existing RTO estimators suffer either from long loss detection times or large number of pre-mature timeouts. The goal of this study is to address these problems by developing an adaptive RTO estimator for high-bitrate low-delay video applications. By exploiting the temporal dependence between consecutive delay samples, we propose an adaptive linear delay predictor. This way, our RTO estimator configures itself based on the video characteristics and varying network conditions. Our approach also features a controller</p>

<p>Altunbasak (Georgia Institute of Technology, US)</p>	<p>that optimally manages the trade-off between the amount of overwaiting and redundant retransmission rate. The skeleton implementation shows that the proposed RTO estimator discriminates lost packets from excessively-delayed packets faster and more accurately than its rivals, which consequently enables the applications to recover more packets under stringent delay requirements.</p>
<p>Understanding Mesh-based Peer-to-Peer Streaming</p> <p>Nazanin Magharei (University of Oregon, US); Reza Rejaie (University of Oregon, US)</p>	<p>Existing peer-to-peer (P2P) streaming mechanisms often use tree-based overlays. These mechanisms are unable to effectively utilize outgoing bandwidth of participating peers, and therefore they are not self-scaling. In contrast, swarm-like content delivery mechanisms such as BitTorrent, exhibit self-scaling property but incorporating them into P2P streaming application are challenging due to both in-time requirement of delivery and the limited availability of future content in P2P streaming applications. In this paper, we illustrate the fundamental design issues and tradeoffs in incorporating swarm-like delivery into mesh-based P2P streaming applications. Leveraging an un-tangled view of the overlay, we present a global pattern for streaming content over a mesh-based overlay that can effectively utilize the outgoing bandwidth of most participating peers. Understanding the proper pattern for streaming over a mesh sheds an insightful light on design and evaluation of P2P streaming mechanisms. In particular, we show that for a given scenario, there is a sweet range for node degree in the overlay that maximizes delivered quality to individual peers.</p>
<p>On Combining Temporal Scaling and Quality Scaling for Streaming MPEG</p> <p>Huahui Wu (Worcester Polytechnic Institute, US); Mark Claypool (Worcester Polytechnic Institute, US); Robert Kinicki (Worcester Polytechnic Institute, US)</p>	<p>Temporal scaling and quality scaling are both widely-used mechanisms to reduce data rates of streaming video. However, combinations of temporal and quality scaling have not been systematically studied. This research extends previous work to provide a model for combining temporal and quality scaling, and uses an optimization algorithm to provide a systematic analysis of their combination over a range of network conditions and video content. We find: 1) quality scaling typically performs better than temporal scaling, with performance differences correlated with the motion characteristics of the video. In fact, when the network capacity is moderate and the loss rate is low, quality scaling performs nearly as well as the optimal combination of quality and temporal scaling; 2) when the network capacity is low and the packet loss rate is high, quality scaling alone is ineffective, but a combination of quality and temporal scaling can provide reasonable video quality; 3) adjusting the amount of forward error correction (FEC) provides significantly better performance than video streaming without FEC or video streaming with a fixed amount of FEC.</p>
<p>AC/DC: an Algorithm for Cheating Detection by Cheating</p>	<p>Time cheats represent some of the most delicate issues in online gaming. Since they act on timing properties of generated game events, these malicious schemes are particularly difficult to thwart when distributed games are deployed over peer-to-peer architectures. Indeed, the absence of a global clock shared among peers enables cheaters to see into the future by waiting for events generated by other peers before generating its own ones (lookahead cheat). This may give an unfair advantage to the cheater. In this work, we consider a version of lookahead cheat generalized in the context of real-time (i.e., not round-based) games. To face this time cheat, we</p>

<p>Stefano Ferretti (University of Bologna, IT); Marco Rocchetti (University of Bologna, IT)</p>	<p>present an algorithm that enables to detect cheaters based on monitoring of network latencies. The basic idea is that of conducting against each suspected peer a sort of cheating activity, by delaying events before notify them to the (hypothetic) cheater. This permits to notice whether that peer waits for these events before generating its own ones. Our claim is that an approach based on the monitoring of communication patterns among peers, allows cheat detection without affecting the performances of the game.</p>
<p>A Multi-stream Adaptation Framework for Bandwidth Management in 3D Tele-immersion</p> <p>Zhenyu Yang (University of Illinois at Urbana-Champaign, US); Klara Nahrstedt (University of Illinois at Urbana-Champaign, US); Bin Yu (University of Illinois at Urbana-Champaign, US); Ruzena Bajcsy (University of California, US)</p>	<p>Tele-immersive environments will improve the state of collaboration among distributed participants. However, along with the promise a new set of challenges have emerged including the real-time acquiring, streaming and rendering of 3D scenes to convey a realistic sense of immersive spaces. Unlike 2D video conferencing, a 3D tele-immersive environment employs multiple 3D cameras to cover a much wider field of view, thus generating a very large volume of data that need to be carefully coordinated, organized, and synchronized for Internet transmission, rendering and display. This is a very challenging task and a dynamic bandwidth management must be in place. To achieve this goal, we propose a multi-stream adaptation framework for bandwidth management in 3D tele-immersion. The adaptation framework relies on the hierarchy of mechanisms and services that exploits the semantic link of multiple 3D video streams in the tele-immersive environment. We implement a prototype of the adaptation framework that integrates semantically-linked stream selection, content adaptation, and 3D data compression with user preference. The experimental results have demonstrated that the framework shows good quality of 3D rendering in case of sufficient bandwidth, while it adapts 3D video streams in a coordinated and user-oriented fashion, and yields graceful quality degradation in case of low bandwidth availability.</p>
<p>Revisiting multimedia streaming in mobile ad-hoc networks</p> <p>Peng Xue (University of Notre Dame, US); Surendar Chandra (University of Notre Dame, US)</p>	<p>Mobile ad hoc networks have been the subject of active research for a number of years. In this paper, we investigate the feasibility of using such networks for transmitting multimedia streams. Our target scenario uses relatively resource rich laptops. We observed that even modest streams place inordinate resource demands on the intermediate nodes for a wide variety of devices and operating systems. Programmed IO operations on the intermediate node dominates in the amount CPU consumption. This raises an important design decision for the ad hoc community, should we a) demand that ad hoc routers support <code>\textit{server-like}</code> hardware with DMA support, b) force a global resource management scheme that cooperatively reduces the network flow such that the intermediate nodes would not see enough traffic to overwhelm them or c) can the local nodes protect themselves from transit traffic without resorting to making specific demands out of hardware. In this paper, we explore mechanisms that can provide some control over the resource consumed without a major revamp of existing operating systems or</p>

	<p>requiring special hardware. We implement our mechanism in the network driver. We present encouraging preliminary results.</p>
<p>Optimizing Overlay Topology by Reducing Cut Vertices</p> <p>Xiaomei Liu (Michigan State University, US); Li Xiao (Michigan State University, US); Andrew Kreling (Michigan State University, US); Yunhao Liu (Hong Kong University of Science and Technolgy, CN)</p>	<p>Most overlay networks are highly decentralized and self-organized. Therefore, cut vertices may exist in such systems due to the lack of centralized management. A cut vertex is defined as a network node whose removal increases the number of network components. Failure of these nodes can break an overlay into a large number of disconnected components and greatly downgrade the network services. In this paper, we propose a distributed mechanism which efficiently detects the cut vertices before they fail and neutralizes them into normal overlay nodes with slight overhead so that the possibility of network decomposition is minimized after they fail. We prove the correctness of this algorithm and evaluate the performance of our design through trace driven simulations.</p>
<p>Off-line Economies for Digital Media</p> <p>Darko Kirovski (Microsoft Research, US); Kamal Jain (Microsoft Research, US)</p>	<p>We propose a novel platform for building off-line markets for digital content. The key objective is to enable an arbitrary user of specific digital content to resell it to other users in an off-line peer-to-peer manner so that part of the proceeds goes to content's copyright holder. Most importantly, one part of the revenues is retained by the seller as an incentive for participating in the distributed economy. To address this objective, a transaction is finalized and incentives distributed to the seller on-line using a client-server architecture. Technologically, such systems can be readily created, for example, by adding a communication tool such as Bluetooth to a portable media player such as the iPod. In this paper, we present a threat model for this system and devise a novel protocol that relies on traditional public-key cryptography to ensure secure and efficient off-line transactions of arbitrary digital content. As a consequence, in our system copyright holders can control the pricing and recruit a powerful marketing and sales force with marginal investment and via various types of incentives, users are offered the ability to sell or purchase content they like anywhere, anytime, and to/from anyone.</p>