doi:10.1093/comjnl/bxl036

Capsule Reviews

Fairouz Kamareddinne

The Capsule Reviews are intended to provide a short succinct review of each paper in the issue in order to bring the content to a wider readership. The Capsule Reviews were compiled by Fairouz Kamareddine. Professor Kamareddine is an Associate Editor of *The Computer Journal* and is based in the Department of Mathematical and Computer Sciences at Heriot-Watt University, Edinburgh, UK.

Design and Implementation of Video Streaming Hot-plug between Wired and Wireless Networks Using SCTP. C.-M. Huang, C.-H. Tsai and M.-C. Tsai
When the network user moves and changes its attachment of the network, the service provided by the application is interrupted. For some network applications, e.g. video-on-demand, the temporary disconnection will bring great inconvenience to users, loss of time and even loss of content. For this reason, a connection persistence mechanism should be developed to keep the continuity of video playout when it is needed to switch the link for video streaming. The paper defines the term video streaming hot-plug as the mechanism for keeping video streaming continuous when a host is switched between wired and wireless networks (in any direction, i.e. wired-wireless, wireless-wireless, wireless-wired). The paper uses a so-called Stream Control Transmission Protocol (SCTP) to solve and implement the video streaming hot-plug mechanism. SCTP supports the use of multiple IP addresses in one connection and hence, multiple paths can be used to transmit data allowing the change of the primary path through which the video streaming data are transmitted before disconnecting the original link. An SCTP-based video streaming hot-plug system called MOVIDEO is implemented where three video streaming hot-plug mechanisms are designed to tackle three types of video streaming hot-plug (wired-wireless, wireless-wireless, wireless-wired), all of them accomplished automatically whereas the first can also be accomplished manually. This gives the user an uninterrupted video streaming service. To evaluate the proposed SCTP video streaming hot-plug mechanism, MOVIDEO is implemented on Linux RedHat 9.0 and possible solutions for video streaming hot-plug using SCTP are tested. In addition to the MOVIDEO system, three different system configurations (which vary in the degrees of multihoming and reliable transmissions) are adopted for comparison. The performance of MOVIDEO is compared with three other configurations using different patterns of primary path change and data transmissions.

An Efficient Distributed Algorithm to Identify and Traceback DDoS Traffic. T. Y. Wong, K. T. Law, J. C. S. Lui and M. H. Wong
The success of the internet attracts malicious attackers who abuse system resources and expose the inherent security problem. Distributed denial-of-service (DDoS) attack is an urgent problem which continues to grow in size, frequency and severity. Two requirements to counter DDoS attacks are the traceback property (which deals with tracing the location of the attackers) and the filtering property (which deals with filtering out the attacking traffic). Tracing is difficult since the attackers often use fake or spoofed IP source addresses. This paper presents a distributed approach to effectively traceback the location of potential flood-based attack sources from where the dominating flows of attack packets were sent. Participating routers can collaboratively collect traffic statistics to a victim site. The latter can use this information to accurately determine the intensities of the local traffic. Based on this information, the victim can determine a subset of attacking routers whose workload consumes a large percentage of the victim's resource. The keypoint, in all this, is how to measure and collect the traffic intensities from the routers participating in the DDoS attack traceback. The paper proposes an approach based on the snapshot algorithm which coordinates all the participating routers in the traffic measurement and the data collecting procedures. Correctness and efficiency are studied in detail where the authors show that the proposed distributed algorithm correctly records the state of each participating router and effectively determines the subset of routers whose local traffic consumes a large portion of the victim's resources. Implementation, applications and limitation issues are discussed.

Security Considerations and Key Negotiation Techniques for Power Constrained Sensor Networks. B. Doyle, S. Bell, A. F. Smeaton, K. McCusker and N. E. O Connor
Sensor networks are used in a wide variety of commercial, industrial and military applications. Although each sensor

node has a limited computational and sensing ability, the combination of sensors working together offers an accurate picture of the spatial region in which the sensors are placed. However, sensor networks pose unique constraints on their operations and hence, traditional security techniques used by conventional networks cannot be applied. Furthermore, wireless sensor networks are vulnerable to attacks. This paper reviews the key distribution attacks commonly used against networks and identifies various encryption schemes necessary to build a secure sensor network. In particular, the paper analyses the feasibility of using a Tate pairing-based approach to solve the key distribution problem. In order to complete the experiments, the paper also needed hardware and software to measure the power consumption of the Tate pairing code. The software and hardware are described and then experiments are conducted and conclusive results are presented.

An Abstract Interface for System Software on Large-Scale Clusters. J. FERNÁNDEZ, E. FRACHTENBERG, F. PETRINI AND J.-C. SANCHO

This paper deals with high-performance computing clusters and their system software. In particular, the paper explores the relationship between the system software requirements and the specialized hardware it runs on, especially with respect to the abstract interface that is generally used. The paper tries to answer the question 'what hardware features, and consequently which abstract interface, should the interconnection network provide the system software designers?'. The paper argues that an efficient and scalable hardware implementation of a small set of network primitives is crucial for scalable software systems. Using an implemented prototype resource-management system, the paper presents experimental results which show that a cluster operating system can be scalable, powerful and simple to implement. A number of case studies is carried out using this prototype resource management system on three different clusters and comparisons with other systems are drawn.

Computation of Elementary Siphons in Petri Nets For Deadlock Control. D. Y. CHAO

Flexible manufacturing systems increase productivity. However, their performance is affected by deadlocks. A number of approaches have been proposed to control deadlocks. One such approach builds a Petri Net model adding necessary control places and arcs to make it deadlock-free while another approach enforces liveness by adding a control place to each strict minimal siphon (SMS)—a siphon is strict if it has no taps. However, the number of SMSs can be large even for a moderate-sized model. This led to a number of works on controlling fewer SMSs or attempting to progress without finding all SMSs. Such work divided SMSs into two groups—elementary and dependent—but still faced the problem of computing all SMSs. Since the number of SMSs grows exponentially, this paper proposes an algorithm to compute elementary siphons based on a new characteristic T-vector without the knowledge of all SMSs. The approach proposed is algebraic and remains to be extended to deal with more complex nets.

Automated Configuration of Multiple Buffer Pools. P. MARTIN, W. POWLEY, X. XU AND W. TIAN

The increased functionality of database management systems (DBMSs) increases their complexity, making it harder to manually manage their performance via the direct adjustment of parameters. A self-tuning or autonomic DBMS solves this management problem. Since effective use of buffers enhances system performance, commercial database management systems like Oracle provide multiple buffer pools. However, these systems provide little support for configuration. This paper applies the self-tuning technology to the configuration of the buffer area used by a DBMS. Some DBMSs divide the buffer area into a number of independent subareas called buffer pools where page replacement is local to each buffer pool. The contribution of the paper is a novel autonomic and integrated approach to handling the configuration of multiple buffer pools in DBMSs especially in the mapping of database objects to buffer pools and appropriately sizing the buffer pools. An automatic solution to the buffer pool configuration problem needs to capture the Database Administrator's (DBA) knowledge and expertise and then to use it in a decision-making tool formulating algorithms to solve the DBA decision-making represented as an optimization problem. Clustering is one of the techniques used in this paper to characterize workload. The proposed automated approach, called BPCluster, first uses data-mining techniques to analyze and cluster performance data collected from the DBMS to characterize the activity of the various database objects for a given workload. A cost-based greedy algorithm determines the sizes of the buffer pools to maximize the throughput for the workload. The effectiveness of the approach is illustrated with a set of experiments conducted with an implementation of BPCluster. Performance evaluations show a number of advantages of this approach.