

Middleware for Network Eccentric and Mobile Applications (MiNEMA) 2nd Closed Workshop*

Extended Program

November, 30th and December, 1st 2004
Lancaster, U.K.

Foreword. This document collects the abstracts of the 23 presentations at the 2nd MiNEMA programme closed workshop.

A fundamental goal of the MiNEMA programme is to foster the collaboration among the European research community in middleware for mobile and ad hoc networks. To achieve this goal, MiNEMA is committed to attract all relevant European research groups to participate in the programme activities. In addition to the original 11 members of the Steering Committee, appointed by the funding agencies, MiNEMA gathers today 17 other research groups.

For the 2nd MiNEMA programme closed workshop, to be held in Lancaster, U.K., from November, 30th to December, 1st 2004, MiNEMA has received 26 requests for participation from 22 members of the programme. The workshop is organized into 7 technical sessions and includes 23 research presentations.

The technical program was coordinated by Luís Rodrigues and Hugo Miranda, Universidade de Lisboa. Local arrangements were coordinated by Paul Grace, University of Lancaster. The organizers hope that the presentations will stimulate the discussion, making the workshop a fruitful experience to all participants.

*This programme is sponsored by European Science Foundation (ESF)

Session 1November, 30th 9h00-10h30*Monitoring Wireless Sensor Networks using Epidemic Protocols*

Daniela Gavidia, Vrije Universiteit Amsterdam, The Netherlands

This talk will focus on monitoring the state of the nodes in a wireless sensor network where the nodes are unaware of the location of the monitoring node(s). For that purpose, we have designed an epidemic protocol to propagate the state of the nodes in a simple, yet effective manner. The nodes communicate their state through local interaction with their neighbours. By utilizing the same protocol for communication, the monitoring node(s) can keep track of the state of the nodes in the network. Simulation results will be presented to illustrate the performance of the epidemic protocol on a network with thousands of nodes.

Adapting the Communication Infrastructure to Context

José Mocito, Liliana Rosa, Nuno Almeida, Hugo Miranda and Luís Rodrigues, Universidade de Lisboa, Portugal

The environment can significantly influence the performance of communication protocols. The selection of the most adequate algorithm must consider factors like the network topology and the availability of wired nodes. Furthermore, conditions are permanently changing, what requires constant monitoring and possibly reconfiguration of the communication substrate. Therefore, an efficient middleware infrastructure supporting communication protocols must be able to: i) monitor the environment to detect relevant changes in the network conditions; ii) in face of the running conditions, evaluate and select the most adequate configuration of communication algorithms and iii) coordinate with remaining participants the installation of the selected configuration.

This talk will present the Morpheus middleware architecture. Morpheus is a modular framework that monitors the execution context and coordinates the installation of the most adequate configuration. The talk will also present evaluation results based on a early prototype of the framework and on one communication protocol, developed to evaluate Morpheus capabilities.

Session 2November, 30th 11h00-13h00*Mobility friendly publish/subscribe*

Sebastien Baehni, EPFL, Switzerland

We describe a publish/subscribe algorithm for mobile ad-hoc networks (MANETs). Our algorithm relies on the mobility of the processes to ensure the propagation and spread of the events. The algorithm takes into account the memory limitation of the mobile processes in providing an efficient garbage collection of old events. We present simulation results that show the advantages/drawbacks of this approach in different mobility models.

Distributed Mobile Event Systems

Sasu Tarkoma, University of Helsinki, Finland

We give a brief overview of current event systems for mobile computing and present the Fuego event system. The proposed system aims to support mobile applications by mobility transparency, fast and efficient subscription topology updates, and facilities for filter matching, covering, and merging. The distributed architecture is based on the event channel abstraction, and uses covering relations and merging to reduce processing overhead in the event routing core. Initial experimental results based on a prototype implementation are also briefly discussed. We highlight our experiences with two applications implemented for mobile phones: the mobile information ticker and mobile presence application.

Scalable Service Discovery for MANET

Francoise Sailhan, INRIA, France

Mobile Ad NETWORKS (MANETs) conveniently complement infrastructure-based networks, allowing mobile nodes to spontaneously form a network and share their services, including bridging with other networks, either infrastructure-based or ad hoc. However, distributed service provisioning over MANETs requires adequate support for service discovery and invocation, due to the network's dynamics and resource constraints of wireless nodes. While a number of existing service discovery protocols have shown to be effective for the wireless environment, these are mainly aimed at infrastructure-based and/or 1-hop ad hoc wireless networks. Some discovery protocols for MANETs have been proposed over the last couple of years but they induce significant traffic overhead, and are thus primarily suited for small-scale MANETs with few nodes. Building upon the evaluation of existing protocols, we introduce a scalable service discovery protocol for MANETs, which is based on the homogeneous and dynamic deployment of cooperating directories within the network. Scalability of our protocol comes from the generated traffic that is minimized, and the use of compact directory summaries that enable to efficiently locate the directory that most likely caches the description of a given service.

Location Identification Algorithms for Location Based Services

Yiannis Socratous, University of Cyprus, Cyprus

“Access of the right information anytime anywhere” is becoming the new driving force that drives the information technology revolution. The “right” information is information that is relevant to the user’s profile and his/her *current position* and/or *time*.

Location Based Services (LBS), also known as LoCation Services (LCS), is an innovative technology that provides information or making information available based on the geographical location of the user. LBS is much promising to ameliorate everyday life and to enhance productivity, since they penetrate in all relationships: “Business to Business”, “Consumer to Business” and vice-versa. Their penetration in the market depends on the accuracy and consistency of the location estimation, among other key variables. The Federal Communication Commission has mandated certain accuracy requirements for location pinpoint. However, the wireless channel imposes many difficulties in estimating the mobile position, and accurate location estimation is indeed a challenge.

Session 3

November, 30th 14h30-16h00

Real-time robotics and process control with Windows CE and Rotor/.NET

Andreas Rasche, Hasso-Plattner-Institute für Softwaresystemtechnik an der Universität Potsdam, Germany

This talk will present the Distributed Control Lab (DCL), which provides an open infrastructure for conducting robotics and control experiments from the Web. Our component-based infrastructure uses .NET technology on PCs as well as control programs written in C language on Windows CE-based devices and the Lego Mindstorms’ Hitachi/Renesas microcontrollers.

We will share experience gained from operation of the lab and its integration into an embedded systems lecture. In addition, we will discuss our approach for protecting the lab from malicious code, which might be potentially uploaded from the Internet.

Within the Lego.NET project, which is part of DCL, we are working towards making the .NET runtime available for small embedded devices. We have implemented a front-end for the GNU Compiler Collection (gcc) which is able to translate pristine IL-Code to native machine instructions, including x86 and Hitachi/Renesas H8-300 as used in the Lego Robots.

Session 4December, 1st 9h00-10h30*Middleware Services for Information Sharing in Mobile Ad-hoc Networks: Challenges and Approach*

Thomas Plagemann, University of Oslo, Norway

Information sharing is a mission critical key element in rescue and emergency operations. Mobile ad-hoc networks (MANETs) could provide a useful infrastructure to support information sharing, but appropriate applications are needed. To facilitate efficient application development for this type of infrastructure, middleware support is needed. In the Ad-Hoc InfoWare project, we are currently developing corresponding middleware services. In this paper, we discuss the application requirements that are imposed onto the middleware services, and we outline our technical approach to address the corresponding challenges. The architecture we propose comprises five main building blocks, namely knowledge management, a local and a distributed event notification service, resource management, and security and privacy management. We indicate design alternatives for these building blocks, identify open problems and relate our approach to the state-of-the-art.

Middleware Platform for Sentient Computing Applications

Thirunavukkarasu Sivaharan and Maomao Wu, Lancaster University, U.K.

Advances in sensor-based computing and wireless communication are leading the way for novel ubiquitous computing applications such as intelligent vehicles and smart buildings. These applications require special properties that traditional computing applications do not support such as context awareness, high decentralisation, autonomous behaviour, adaptivity and innate collaboration. This talk presents a new computational model and middleware platform that reflect support for the required properties. The sentient object model is proposed for the construction of the ubiquitous computing applications. A flexible and (run-time) reconfigurable component based middleware has been built to provide run-time support to engineer sentient object based applications. This talk focuses on particular novel application - 'cooperating sentient vehicle' which is related to intelligent transportation systems which utilise inter vehicle cooperation without human assistance to provide autonomous vehicle navigation from a given source to a predetermined destination. The resultant sentient vehicles are 'context-aware' autonomous cars that form cooperative 'flotillas' of peers using mobile ad hoc network environments (MANETs). We report on our experiences on a particular configuration of the middleware architecture designed to meet the challenges of the aforementioned sentient vehicle environments. We show how such a middleware can be used to engineer a proof of concept cooperating sentient vehicle application, and highlight the research challenges raised.

From Ad Hoc Networks to Ad Hoc Applications

Benoît Garbinato, Université de Lausanne, Switzerland

This talk will present the first result of an ongoing research on mobile peer-to-peer communication paradigm and on adhoc networks. This result is three-fold: (1) we propose a definition of adhoc applications that is independent of the notion of adhoc networks, (2) we advocate for the need of an adequate development framework for building adhoc applications, (3) we introduce such a framework, which is independent of the underlying network technology but which relies on existing standards. To illustrate our approach, we will also show how our framework was used to build a typical adhoc application known as the Ubiquitous Flea Market.

Session 5December, 1st 11h00-12h30*Using Stigmergy to Co-ordinate Pervasive Computing Environments*

Peter Barron, Trinity College Dublin, Ireland

Pervasive computing environments have proven difficult to develop in a form that supports the integration and organisation of devices and applications in a spontaneous and transparent manner. This is partly due to the highly dynamic and unpredictable nature of these types of environments, and is often further hampered by the limited resources found on devices. In this paper we present a highly decentralized method of organising the components of a pervasive computing environment, supporting spontaneous interaction between entities and providing robust system-wide behavior. Our inspiration for this work stems from nature and the observations made by the French biologist Grass on how social insects co-ordinate their actions using indirect communication via the environment, a phenomenon that has become known as stigmergy. In the stigmergic approach there are fewer dependences between entities allowing for the incremental construction and improvement of solutions without adversely affecting the rest of the pervasive computing environment. The approach is encapsulated in Cocoa, a framework that supports the use of stigmergy to build self-organising environments that promotes the autonomy of entities. Experience using Cocoa has shown that entities can be integrated into a pervasive computing environment in a spontaneous manner and that co-ordinated behavior can emerge.

Optimistic Consistency with Dynamic Version Vector Weighted Voting

João Barreto and Paulo Ferreira, Inesc-ID, Portugal

Mobile and loosely-coupled environments call for decentralized optimistic replication protocols that provide highly available access to shared objects. A fundamental property of most optimistic protocols is to guarantee an eventual consensus on a commit order among the set of tentatively issued updates.

In this paper we propose a replicated object protocol that employs a novel

epidemic weighted voting algorithm based on version vectors for achieving such goal. An epidemic voting strategy eliminates the single point of failure of primary commit approaches, while not imposing the simultaneous accessibility of a plurality quorum.

Our protocol introduces a significant optimization over basic epidemic weighted voting solutions by allowing multiple-update candidates through the use of version vectors. As a result, it is able to commit multiple, causally related updates at a single distributed election round. Complementarily, we describe how dynamic version maintenance can be easily incorporated into the voting protocol in order to reduce version vector size and avoid the need for complete knowledge of group membership.

We demonstrate that our proposed algorithm is especially advantageous when considering realistic, non-uniform update models. We support such assumptions by presenting comparison results obtained from side-by-side execution of reference protocols in a simulated environment.

Session 6

December, 1st 14h00-16h00

Adaptation as a key technology for multimedia delivery in networks with fluctuating QoS

Laszlo Böszörményi, ITEC-DS, Klagenfurt University, Austria

In an environment where clients are mobile in several senses (they change their location, their equipment, their roles and interests etc.) the communication infrastructure must also be able to dynamically adapt to the changing requirements. In the field of video delivery, a lot of research has been done on “reactive” and “defensive” adaptation of video streams. Much less research has been done into “offensive” and “proactive” adaptation.

The state-of-the art of (multimedia) content distribution is: (1) Traditional Client/Server Architectures, (2) Peer-to-Peer (P2P) computing, and (3) Content Delivery Networks (CDNs). Client/Server Architectures do not scale well; P2P computing is too fragile and also has serious limits in scalability; CDNs are too rigid and expensive.

We suggest an intermediate solution, based on dynamic overlay networks of proxies. We call this technology X2Y Streaming (streaming based on proxy-to-proxy communication). The basic idea is that multimedia content is dynamically migrated/replicated (fully or partially) onto logical networks of proxies supporting a small geographic region and a certain group of interests.

Middleware for Adaptive Programmable Networks

Nico Janssens, Wouter Joosen and Pierre Verbaeten. DistriNet, K.U.Leuven, Belgium

Middleware is needed to manage the self-adapting behavior of heterogeneous and dynamic mobile networks. This involves support for (1) monitoring the network, (2) deploying new services on the protocol stacks to deal with the detected

problems and (3) maintaining a consistent network service composition after deployment. The latter is particularly important in mobile ad-hoc networks, since frequent (dis-)connections as well as autonomous mobility of embedded devices result into a highly dynamic topology. In this talk, we mainly focus on the second requirement, i.e. how to provide in safe service deployment for programmable networks.

When examining the type of services that are covered by protocol stacks to boost network quality, it becomes obvious that many existing distributed network services (such as compression, fragmentation and reliability) conform to a point-to-point based distributed service model. This service model represents a pair of tightly-coupled collaborating entities (not necessarily situated on neighboring nodes) that need to cooperate so as to implement a distributed service. By consequence safe evolution is essential to prevent runtime adaptation of such services from jeopardizing the correct functioning of the network.

Programmable networks impose preconditions to service adaptations. Depending on the network service being deployed as well as on the characteristics of the underlying execution environment (i.e. the network itself), different non-functional requirements (such as minimizing performance penalties, reducing total deployment time or maintaining quality of service (QoS) during the adaptation) are to be satisfied. To meet these requirements, the deployment algorithm must be customizable, providing in tailored optimized scenarios. Hence, the development of point-to-point based network services for runtime adaptable networks is often complex and error-prone.

We will elaborate on a middleware that coordinates transparent and safe addition, replacement and removal of distributed services among programmable nodes. This middleware provides in an exhaustive set of deployment scenarios (with this reducing the openness of programmable nodes with respect to service evolution). In doing so, we aim at minimizing the performance penalty of safe distributed service evolution by speeding up the reconfiguration process if possible. By packaging this complexity on top of the operating system of programmable nodes, the developer or administrator of a distributed service is spared from implementing a service specific reconfiguration protocol. In this way, we aim at making the development of distributed services for programmable networks less error-prone.

Causal Cluster Consistency Management for P2P systems

Boris Koldehofer, Chalmers University, Sweden

Causal cluster consistency provides a consistency management scheme which is aimed for middleware supporting multi-peer collaborative applications where a large number processes share a large set of replicated objects. Many such applications, like peer-to-peer collaborative environments for educational, training or entertaining purposes, platforms for distributed monitoring and tuning of networks, rely on a fast propagation of updates on objects, however they also require a notion of consistent state update. Causal Cluster Consistency is an effort to cope with these requirements by also ensuring scalability. In a cluster consistency protocol a privileged dynamic set of processes, called coordinators, may concurrently propose updates to a subset of objects which form a cluster. The updates are applied in some order of interest by the coordinators of the

cluster. Moreover, any interested process can receive update messages referring to replicated objects, with an option for the updates to be delivered unordered or in the same order as to the coordinators.

This talk illustrates a protocol implementing causal cluster consistency, which provides a fault tolerant and dynamic membership algorithm to manage the cluster members. The membership algorithm also coordinates the dynamic assignment of process identifiers to vector clock entries. Hence, this protocol provides optimistic causal order in combination with any group communication protocol.

JazzEnsemble: A Group Communication Middleware for Ad-Hoc Networks
Roy Friedman, Technion, Israel

The success of ad-hoc networks largely depends on having applications developed to this environment. As in other domains, appropriate middleware is key to enabling expedited software development in this domain. Given the nature of ad-hoc networks, it is reasonable to assume that many applications in these environments will be group based. Therefore, group communication is a highly desired middleware service.

In this talk I will describe JazzEnsemble, which is an adaptation to mobile ad-hoc networks of the Ensemble group communication toolkit (originally developed for wired clustered environments). JazzEnsemble inherited the modular layered architecture of Ensemble, which makes it highly attractive framework for both development and experimentation. In order to deliver good performance in these highly dynamic environments, we have based JazzEnsemble on the notion of fuzzy membership, which allows the system to better react to temporary poor connectivity while providing strong semantics to the application. In particular, the system can even deliver strong virtual synchrony semantics, and implement a replicated state machine with good performance, despite mobility.

The system can be used in true MANETs, yet is also includes emulations layers, which allow it to be run on standard clusters. This enables development and testing of production code in the emulated environment. JazzEnsemble can accommodate all levels of MANET related protocols, implemented either as a layer of the middleware, or provided as an external service, ranging from discovery and routing to high level applications.

JazzEnsemble is offered as an open source toolkit (BSD style license), and is currently ported to Linux and Windows on x86 (PCs and laptops) and on iPAQ PDAs. In the talk I will describe the current system and our on-going research on extending it.