Performance Evaluation of a Framework to Support Path Changes in IP-based Access Networks

Liesbeth Peters^{*}, Ingrid Moerman, Bart Dhoedt and Piet Demeester
Department of Information Technology (INTEC), Ghent University - IBBT - IMEC
Gaston Crommenlaan 8 bus 201, B-9050 Gent, Belgium

{Liesbeth.Peters, Ingrid.Moerman, Bart.Dhoedt, Piet.Demeester}@intec.UGent.be

ABSTRACT

The increasing use of wireless networks and the popularity of multimedia applications, leads to the need of Quality of Service support in a mobile IP-based environment. In [1], we investigated the reasons for path changes in the access network and we presented a framework to support these path changes regardless their cause and to avoid router inconsistencies. In this paper, the performance of this framework is thoroughly evaluated. To this end, we added the functionality of the framework to our Q-MEHROM handoff scheme [2]. However, the ideas and results presented in this paper are not restricted to this micromobility protocol. The performance with and without the defined functionalities is compared and tested under different network topologies, in terms of received service, packet loss and end-to-end delay.

Categories and Subject Descriptors: C.2.2 [Computer-Communication Networks]: Network Protocols—*routing protocols*.

General Terms: Performance.

Keywords: IP mobility, IP QoS, link failure support, micromobility, resource reservations.

1. INTRODUCTION

Today, wireless networks evolve towards IP-based infrastructures to allow a seamless integration between wired and wireless technologies. Mobile IP [3] is the best known routing protocol that supports host mobility. Mobile IP is used to support macromobility, as the use of Mobile IP in combination with frequent handoffs suffers from high handoff latency, packet loss and control overhead in the core network. Micromobility protocols, such as Cellular IP, HAWAII and MIPv4 Regional Registration [4], try to solve these weaknesses. Low latency handoff protocols [5], are developed to reduce the amount of configuration time in a new subnet

^{*}L. Peters is a Research Assistant of the Fund for Scientific Research - Flanders (F.W.O.-V., Belgium).

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

MSWIM '06, October 2-6, 2006, Torremolinos, Malaga, Spain.

Copyright 2006 ACM 1-59593-477-4/06/0010 ...\$5.00.

by using link layer triggers. Most research in the area of micromobility assumes that the access network has a tree or hierarchical structure.

In a mobile IP-based environment, users want to receive real-time applications with the same QoS (Quality of Service) as in a fixed environment. Current work within the IETF NSIS (Next Steps in Signaling) working group includes the standardizing of an IP signaling protocol with QoS as the first use case [6]. Other research to support QoS in the micromobility environment, resulted in several extensions to RSVP (Resource ReSerVation Protocol) [7]. All these extensions also assume that the structure of the access network is hierarchical and forms a tree topology.

An access network that has to deliver QoS, should be robust against link failures and able to balance the load on its links. Therefore, a meshed topology, instead of the assumed tree topology, is preferred. In our previous work, we developed the MEHROM (Micromobility support with Efficient Handoff and Route Optimization Mechanisms [8]) per-host handoff scheme. It shows a good performance, irrespective of the topology. This handoff mechanism was also extended to support resource reservations. In case of handoff, Q-MEHROM [2] updates the routing information and allocates the resources for the mobile host simultaneously. Resource reservations along the common part of the old and new path are reused.

In this paper, we consider the reservation of resources for data flows towards mobile hosts in a more general scenario. Not only handoffs, but also link failures or load changes, can be reasons why the setup of a new path is required. We shortly describe the framework, presented in [1], to support path changes regardless their reason. The performance of the different framework mechanisms is compared and tested under different network topologies. The rest of this paper is structured as follows. Section 2 investigates the causes of path changes in the access network. Section 3 describes the functionality that is required in the access network elements to support these path changes. The evaluation of this framework is presented in Section 4. The final Section 5 contains our concluding remarks.

2. REASONS FOR PATH CHANGES

Let us consider an IP-based access network and a mobile host (MH), residing in the area of one of the access routers (AR). A path for the MH between the domain gateway (GW) and the current AR is set up by a micromobility protocol and resources for the MH are allocated along this

path by a resource reservation mechanism. We distinguish the following three situations, in which the setup of a new path and the reservation of resources along this new path are necessary or preferred:

1. **Mobility** (Fig. 1.A): When the MH moves from an old access router (oAR) into the area of a new access router (nAR), a new path must be set up. The cross-over router (CR) is defined as the router, closest to the nAR, that is part of both the old and new path.

2. **Link or node failure** (Fig. 1.B): When a link or a node on the path for the MH fails (in what follows, we consider single link failures), a new path is needed. In the assumption that a new path exists, CR1 is defined as the last router on the common part that starts in the GW and CR2 is defined as the first router on the common part that ends in the AR.

3. **New resources** (Fig. 1.C): When, e.g. due to the movements of other MHs, enough resources become available, a MH that receives best-effort service can change paths to obtain QoS service. In this case, it is preferred to set up a new path and allocate the requested resources. In contrast to the case of a link failure, the new path can be exactly the same as the old path. We do not define any CR, as new reservations must be made along the entire path from the GW to the AR.

Figure 1: Illustration of the reasons why the setup of a new path is necessary or preferred.

3. FRAMEWORK

This section shortly describes the functionality that is required in the different access network elements to support the path changes regardless their cause. Functionalities, needed to avoid router inconsistencies in the access network, are also incorporated into the framework.

3.1 Mobile Host (MH)

The functionality of the MH is kept as simple as possible. When a MH receives a beacon (e.g. a Mobile IP Agent Advertisement) from an AR, it must only decide whether it sends back a registration request (REG-REQ).

- **Mobility support:** A REG-REQ is sent by the MH, every time it receives a beacon from a new AR and decides to perform handoff.

- **Link failure support and new resources usage:** A new path must be set up while the MH stays connected to the same AR. As the setup of a new path must be triggered by the MH, the MH also sends a REG-REQ after the receipt of a beacon from its current AR.

- **Avoid router inconsistencies:** The MH never sends a REG-REQ to its current AR, when it is expected to perform handoff soon. This decision can be made using cross-layer information, e.g. link layer triggers or location information. This avoids simultaneous path setups caused by a link failure or newly available resources and caused by handoff.

3.2 Access Router (AR)

Most functionality is concentrated in the ARs. Upon the receipt of a REG-REQ, the AR first detects the reason for the request and then reacts accordingly.

- **Mobility support:** If the REG-REQ is sent by a MH that is not yet connected to any AR, the power up mechanism is started. If the REG-REQ indicates that the MH's current AR is different from the AR receiving the REG-REQ, handoff is performed.

- **Avoid router inconsistencies:** If the AR already has a wireless entry for the MH, no new path is set up before the

receipt of an acknowledgment to finish the previous setup for that MH. Moreover, if any router inconsistency is detected by the AR, a new path setup does not rely on the path info in the REG-REQ.

- **Link failure support:** The AR detects whether the current path is broken and needs to be restored.

- **New resources usage:** If the MH receives best-effort service, the AR checks if new resources became available to switch to QoS service.

The ARs also perform admission control. If, at the time of connection setup, the access network has not enough resources to deliver the requested QoS, the request is rejected. If, at the time of a handoff or a link failure, the required resources for an existing connection are not available anymore, the delivered service is reduced to best-effort.

When an AR is informed about a link failure, it sends out a beacon to get a REG-REQ from the MH as fast as possible. However, a link failure is not predictable and can occur during an ongoing path setup. In order to avoid router inconsistencies, the AR also sends a CHECK message to the GW for each path setup that is indicated as not yet finished. When the GW or a network router on the path has no entry for the MH or the entry points to a broken link, this is reported to the AR that starts a new path setup.

3.3 Network Router (NR)

A NR must only decide whether it has the function of a cross-over router (CR). A NR detects itself that it has the function of a CR, if it receives a route update message, originated by the MH's new AR, on a different interface than the interface stored in its routing cache.

- **Mobility support:** As soon as the CR is updated, a new path is set up. The CR must not forward the route update further.

- **Link failure support and new resources usage:** In the case of a link failure, or when the service switches from best-effort to QoS or vice versa, it is necessary that the route update message is forwarded all the way to the GW.

4. PERFORMANCE EVALUATION

For the evaluation, the network simulator NS-2 [9] is used. Q-MEHROM is used as the coupled micromobility and resource reservation mechanism [2] and QOSPF is used as the routing agent [10]. In addition, the proposed framework functionalities of Section 3 are implemented. The following simulation parameter values are chosen:

- Access network: The wired links of the access network have a delay of 2 ms and a capacity of 1.5 Mbps. For the wireless links, IEEE 802.11 is used with a physical bit rate of 11 Mbps. The wired links have a very low bandwidth in order to simulate a highly loaded access network with a limited number of MHs. Doing so, we avoid that the wireless links become a bottleneck. The investigated topologies are illustrated by Figure 2. An AR broadcasts beacons at fixed time intervals of 1.0 s. The distance between two adjacent ARs is 200 m, with a cell overlap of 30 m.

- Mobility pattern: Each of 8 MHs moves at a speed of 20 m/s along a straight line between the center of the cells. A MH resides thus in an overlap region during 1.5 s. At the start of a simulation, each MH starts in a different cell. A MH randomly chooses a cell. Once it arrives at the center of the selected cell, a new destination is chosen immediately. A MH is assumed to perform handoff soon when it is located within the overlap region of two adjacent cells.

- Link failure scenario: Every 200 s, a randomly chosen link on Level 2 of the tree topology goes down for 25 s.

- Data traffic and background traffic: CBR (Constant Bit Rate) traffic patterns are used, with packets of 1500 bytes. One UDP connection is set up between the sender (a fixed host in the core network) and each MH. During one simulation, all 8 data traffic connections have the same data rate. In addition, one UDP connection is set up between the GW and each AR, resulting in 8 background traffic connections. These have also an identical data rate.

- Requested service: Background traffic always receives best-effort service. In order to receive QoS service, a MH requests that an amount of bandwidth, equal to the data traffic rate, is reserved along the wired links in the access network. A small part (1%) of the link bandwidth is reserved for control traffic. For the scheduler, a variant of CBQ (Class Based Queueing) is used.

- Simulations: The duration of a single simulation is 1800 s. For the given mobility pattern and link failure scenario, this results in a 180 handoffs/MH and a 9 link failures.

Figure 2: The tree (solid lines), meshed (solid and dashed lines) and random (solid, dashed and dotted lines) topology, used during the simulations.

4.1 Framework Operation

This subsection aims to clarify the operation of the different framework mechanisms. Therefore, all mechanisms are active and all reasons for path changes are present in the investigated scenario. The MHs move conform the mobility pattern, link failures occur according to the link failure scenario, and each MH requests QoS service. Background traffic connections with a data rate of 375 kbps are used.

The upper figure of Figure 3 shows how frequently each possible action, that can be taken by a MH upon receipt of a beacon, occurs. As the MHs move at constant speed, the amount of REG-REQs sent to a new AR to perform handoff and the times that no REG-REQ is sent because the MH is performing handoff soon is the same, irrespective of the topology and the requested amount of bandwidth.

The main difference between the topologies is found in the amount of REG-REQs sent to the current AR. A MH may receive a new beacon and send a new REG-REQ before it receives the registration reply, confirming path setup after handoff. In the case of the tree topology, this part of REG-REQs is caused by the link failures. Indeed, when a MH performs handoff to an AR that has no path to the GW due to a link failure, the AR can not send a registration reply. In the case of the meshed and random topology, this type of REG-REQs occurs only for high bandwidth values and is caused by rejected power up requests. This is explained by the fact that even for handoffs between AR11 and AR12, the resource reservations along the links closest to the GW can be reused, preventing other MHs to power up.

The lower figure of Figure 3 shows how frequently each action, taken by an AR upon receipt of a REG-REQ, occurs. The results depend more on the access network topology and the amount of requested bandwidth. In general, part of the actions is the setup of a new path after handoff. The percentage of this part decreases when the requested bandwidth increases. Indeed, for higher data rates, more MHs receive best-effort service, and the ARs check more frequently if it is possible to switch to QoS service. In the case of the tree topology, when the AR detects that the path for a MH is broken by a link failure, no other path is available. In the case of the meshed and random topology, the AR succeeds in the setup of a new path after a link failure. This percentage also lowers for higher data rates.

4.2 The Switching Mechanism

The mechanism to use newly available resources is also referred to as the switching mechanism, i.e. to switch from best-effort (BE) service to QoS service. In this subsection, the following scenario is considered: the MHs move according to the described mobility pattern, no link failures occur and each MH requests QoS service. The data rate of each background traffic connection is 375 kbps. Both framework mechanisms to support mobility and link failures are active. Figure 4 compares the performance with and without the defined switching mechanism.

In order to evaluate the impact of the switching mechanism on the **received service**, we look at the fraction of the total simulation time that a MH receives QoS service, calculated by $\frac{\sum_{i=1}^{H-1} (t_{i+1} - t_i) r_i + (t_{\text{end}} - t_H) r_H}{(t_{\text{end}} - t_0)}$. Here, t_0 indicates the start of the simulation, t_1 is the time of a successful power up and t_{end} is the moment that the simulation ends. H is the total number of path setups. At the end of each path

Figure 3: Analysis of the actions taken by a MH upon receipt of a beacon from an AR (upper figure), and by an AR upon receipt of a registration request from a MH (lower figure).

setup. i.e. at t_i with $1 \leq i \leq H$, the flag r_i indicates whether the requested resources were successfully reserved in the access network ($r_i = 0$ means BE, $r_i = 1$ means QoS). For the range [150kbps, 200kbps], even without the switching mechanism, QoS can be delivered 100% of the time. For the range [250kbps, 550kbps], the time that QoS service is received increases when the switching mechanism is used. The impact is most significant for the tree topology. This can be explained as follows: when a MH switches from AR11 to AR12 or vice versa in the tree topology, the resources on the links closest to the GW must be released. Other MHs receiving BE service can take advantage of this release and make reservations. In the case of the meshed and random topology, the resources on these links can be reused, preventing other MHs from switching to QoS service or even from powering up. However, the meshed and random topology still experience a slight improvement by the use of the switching mechanism: when the links closer to the ARs form a bottleneck, a MH is forced to switch to BE service and other MHs can use these newly available resources.

The simulated **packet loss** can have two possible causes: packet loss due to handoff and packet loss due to bandwidth limitations, i.e. buffer overflows. For the range [150kbps, 200kbps], the packet loss is very small and only caused by handoff. For the range [250kbps, 550kbps], the packet loss is mainly caused by the bandwidth limitations. The packet loss is smaller when the switching mechanism is used, as the switching mechanism results in more bandwidth reservations for the MHs, at the cost of higher packet losses for the background traffic.

Also the **end-to-end delay** of the data packets, i.e. the time needed by a data packet to travel from the sender to

the receiver, is influenced by two factors: the number of hops traversed by the data packets and the time that these data packets spend in the queues of the nodes along the path. These two factors have a different impact. In the case of BE service, the path has always the minimum number of hops between the GW and a given AR. However, the data packets are only routed if enough bandwidth is available and the packets may spend more time in the queues. In the case of QoS service, a path with enough available resources is selected and resources are reserved. As a result, the time spent in the queues is very small, but the path may count several additional hops. The use of the switching mechanism may thus result in a better or a worse performance in terms of end-to-end delay. In the tree topology, the number of hops is always minimal, regardless the offered service. As a result, the switching mechanism results in more QoS service, thus lower delays in the queues and a lower end-to-end delay. In the case of the meshed topology, the end-to-end delay is mostly higher than in the case of the tree topology, as the use of the mesh links to deliver QoS service results in an extra hop in the path. The random topology can have a lower end-to-end delay than the tree topology, as more than one path between the GW and an AR exists with the same minimum number of hops. In general, the end-to-end delay has an increasing tendency as a function of the requested bandwidth, due to the increasing load on the network, resulting in longer times spent in the queues of the routers.

4.3 Link Failure Support

In this subsection, a similar evaluation is performed for the impact of the use of the link failure support mechanism on the performance. The following scenario is considered: the MHs are fixed, with one host in each cell, link failures occur and each MH requests QoS service. To increase the impact of the link failures during a simulation, a link break occurs every 100 s instead of every 200 s. The data rate of each background traffic connection is 375 kbps. The mechanisms to support mobility and switching are both active. Figure 5 compares the results for the situation where the link failure support mechanism is used with the situation where the link failure support mechanism is not active.

The results for the tree topology are independent of the use of the link failure support mechanism. At the time of a link failure, no alternative path exists between the GW and the involved ARs. The link failure support mechanism is simply not able to set up a new path.

A remarkable result is that the time that a MH **receives QoS service** is the same for all topologies regardless the support of link failures. The reason is that only information, received by the MHs in the registration replies, is used for this calculation. The MHs are not aware of the occurrence of a link failure and only receive a new registration reply when a new path is successfully set up, by the link failure support mechanism or when the broken link is up again. As no mobility is involved and only links of Level 2 fail, reservations on the links of Level 1 are never released. As a result, the reservation flag in the replies has the same value as before the link failure.

The average **packet loss** per link failure decreases when the link failure support mechanism is active, as expected. However, for the range [400kbps, 550kbps], the packet loss is mainly caused by the bandwidth limitations. For the range [500kbps, 550kbps], the packet loss for the meshed and ran-

Figure 4: Impact of the use of the switching mechanism on the performance.

Figure 5: Impact of the use of the link failure support mechanism on the performance.

dom topology with link failure support remains even higher than the packet loss for the tree topology. This can be explained by the fact that the setup of a new path for the affected MHs and the rerouting of the affected background traffic (by OSPF), increases the bandwidth shortage in the rest of the access network. This, in its turn, influences the packet loss experienced by the other MHs.

The **end-to-end delay** increases when the link failure support mechanism is used in the meshed topology, because the new path will have additional hops. In the random topology, the new path can also count the minimum number of hops. As no mobility is involved and the resources are not released due to a link failure, in contrast to handoff, only a limited number of MHs can successfully power up during the simulation. As a result, the amount of data connections that are set up decreases as a function of the requested bandwidth, the load on the access network remains more or less the same, as well as the end-to-end delay.

5. CONCLUSIONS

This paper evaluated our framework [1], developed to support path changes in an IP-based access network regardless the reason of the path change. The operation of our framework was clarified by simulations and the performance was evaluated in terms of received service, packet loss and end-to-end delay, for different network topologies.

The simulations showed that the tree topology is the topology that takes most advantage of the switching mechanism to increase the received QoS service and to lower the packet loss and end-to-end delay. However, the tree topology is not able to take advantage of the link failure support mechanism as only a single path exists between the GW and each AR. By the meshed and random topology, a new path can be set up, which lowers the packet loss. The major difference between the latter topologies can be found in the end-to-end

delay. In general, the random topology shows a lower end-to-end delay as more than one path with minimum number of hops may exist between the GW and a given AR.

6. ACKNOWLEDGMENTS

Part of this research is funded by the Belgian Science Policy Office through the IAP (phase V) Contract No. IAPV/11, by the IWT in Flanders through the GBOU Contract 20152 "End-to-end QoS in an IP-based mobile network", and by the CELTIC-WISQUAS Project CP2-035 "Wireless systems providing high quality services".

7. REFERENCES

- [1] L. Peters, I. Moerman, B. Dhoedt, and P. Demeester. Support of path changes with resource reservations for mobile hosts in IP-based access networks. *GLOBECOM 2005, USA*, vol. 6, pp. 3559-3563.
- [2] L. Peters, I. Moerman, B. Dhoedt, and P. Demeester. Q-MEHROM: mobility support and resource reservations for mobile senders and receivers. *Computer Networks, Elsevier*, 2006, vol. 50, no. 8, pp. 1158-1175.
- [3] C. Perkins, Ed. IP mobility support for IPv4. *IETF RFC 3344*, Aug. 2002.
- [4] A. Campbell, J. Gomez, S. Kim, C. Wan, Z. Turanyi, and A. Valkó. Comparison of IP micromobility protocols. *IEEE Wireless Communications*, Feb. 2002, pp. 72-82.
- [5] R. Koodli, and C. Perkins. Mobile IPv4 fast handovers. *draft-ietf-mip4-fmip4-00.txt (work in progress)*, Feb. 2006.
- [6] R. Hancock, G. Karagiannis, J. Loughney, and S. van den Bosch. Next steps in signaling (NSIS): framework. *IETF RFC 4080*, June 2005.
- [7] B. Moon and A.H. Aghvami. Quality-of-service mechanisms in all-IP wireless access networks. *IEEE JSAC*, June 2004, vol. 22, no. 5, pp. 873-887.
- [8] L. Peters, I. Moerman, B. Dhoedt, and P. Demeester. MEHROM: micromobility support with efficient handoff and route optimization mechanisms. *16th ITC Specialist Seminar*, Belgium, 2004, pp. 269-278.
- [9] NS-2 Home Page, <http://www.isi.edu/nsnam/ns>
- [10] QoS in NS-2, <http://www.netlab.hut.fi/tutkimus/ironet/ns2/ns2.html>

Session 11: Cellular Nets

- **Location-based Admission Control for Differentiated Services in 3G Cellular Networks** 322
R. Núñez-Queija (*Eindhoven University of Technology*), H.-P. Tan (*University of Dublin*)
- **Analysis of Soft Handover Measurements in 3G Network** 330
K. Raivio (*Helsinki University of Technology*)
- **Towards Frequency Reuse 1 Cellular FDM/TDM Systems** 338
M. C. Necker (*University of Stuttgart*)
- **Towards the Fourth Generation of Cellular Networks: Improving Performance Using Distributed Negotiation** 347
R. Lin, D. Dor-Shifer, S. Rosenberg, S. Kraus (*Bar-Ilan University*), D. Sarne (*Harvard University*)

Session 12: Mobility and Routing

- **A Mobility Sensitive Approach for Efficient Routing in Ad Hoc Mobile Networks** 357
A. Bamis (*University of Patras*), A. Boukerche (*University of Ottawa*),
I. Chatzigiannakis, S. Nikolettseas (*CTI and University of Patras*)
- **Bandwidth Constrained Routing Problem in Multi-Hop Wireless Networks** 365
C.-Y. Chiu, Y.-L. Kuo (*National Taiwan University*), H.-K. Wu (*National Central University*)
- **Performance Evaluation of a Framework to Support Path Changes in IP-based Access Networks** 370
L. Peters, I. Moerman, B. Dhoedt, P. Demeester (*Ghent University*)
- **The QoS-RWP Mobility and User Behavior Model for Public Area Wireless Networks** 375
G. Resta, P. Santi (*IIT - CNR*)
- **Routing in Intermittent Network Topologies** 385
P. Mundur, S. Lee (*University of Maryland*),
M. Seligman (*Laboratory for Telecommunication Sciences*)

Author Index 390

**The Association for Computing Machinery
1515 Broadway
New York, New York 10036**

Copyright © 2006 by the Association for Computing Machinery, Inc. (ACM). Permission to make digital or hard copies of portions of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyright for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permission to republish from: Publications Dept., ACM, Inc. Fax +1 (212) 869-0481 or <permissions@acm.org>.

For other copying of articles that carry a code at the bottom of the first or last page, copying is permitted provided that the per-copy fee indicated in the code is paid through the Copyright Clearance Center, 222 Rosewood Drive, Danvers, MA 01923.

Notice to Past Authors of ACM-Published Articles

ACM intends to create a complete electronic archive of all articles and/or other material previously published by ACM. If you have written a work that has been previously published by ACM in any journal or conference proceedings prior to 1978, or any SIG Newsletter at any time, and you do NOT want this work to appear in the ACM Digital Library, please inform permissions@acm.org, stating the title of the work, the author(s), and where and when published.

ISBN: 1-59593-477-4

Additional copies may be ordered prepaid from:

ACM Order Department
PO Box 11405
New York, NY 10286-1405

Phone: 1-800-342-6626
(US and Canada)
+1-212-626-0500
(all other countries)
Fax: +1-212-944-1318
E-mail: acmhelp@acm.org

ACM Order Number 617061

Printed in the USA

ACM MSWiM 2006

Proceedings of the Ninth ACM Symposium on
**Modeling, Analysis and Simulation
of Wireless and Mobile Systems**

October 2-6, 2006 • Torremolinos, Malaga, Spain

Edited by

Enrique Alba & Carla-Fabiana Chiasserini (*General Co-Chairs*)

Nael Abu-Ghazaleh & Renato Lo Cigno (*Program Co-Chairs*)

Sponsored by

ACM SIGSIM & IEEE Computer Society