

**Review Form: 1<sup>st</sup> International Workshop on  
Services and Infrastructure for the Ubiquitous and Mobile Internet (SIUMI'05)**



**SIUMI 2005**

**WEB MINDS**

Columbus, Ohio,  
USA, June 6<sup>th</sup>, 2005

In conjunction with the 25th Int. Conference on Distributed Computing Systems (**ICDCS'05**)

Paper Number: 14

Paper Title: The last hop of global notification delivery to mobile users: matching preferences, context, and device constraints

Authors: Dimitrii Zagorodnov and Dag Johansen

---

**Reviewer1:**

<b>Familiarity</b> Rate your familiarity with the topic	1	2	3	4	
	Novice	Some knowledge	<b><u>Familiar</u></b>	Expert	
<b>Significance</b> Technical relevance and practicality of ideas in the paper	1	2	3		
	Not significant	Somewhat significant	<b><u>Highly significant</u></b>		
<b>Novelty</b> How original the problem and/or solution method is	1	2	3		
	Not novel	<b><u>Somewhat novel</u></b>	Highly novel		
<b>Quality of Presentation</b> Writing and presentation style/accuracy	1	2	3		
	Poorly written	Could be improved	<b><u>Well written</u></b>		
<b>Overall Recommendation</b>	1	2	3	4	5
	Strong reject	Weak reject	Weak accept	<b><u>Accept</u></b>	Strong accept

**Contributions**

After a description of their publish/subscribe mechanism, the authors focus on optimizing the delivery of messages from the next-to-last hop of a network to the mobile device used by users. Issues addressed include optimizing the bandwidth utilized and power and and storage capacity used by mobile devices. Some simulations and their results are described and an algorithm for the spooling mechanism is presented.

**Strengths and weaknesses**

(Strengths) Paper is relevant to the conference and the problem domain and architecture are interesting. Results from simulations are presented and described

(Weaknesses) Some of the results are obvious

### Detailed public comments

The paper addresses an important problem – efficiently supporting publish/subscribe infrastructures on the internet, especially the last hop. However, the title of the paper is more ambitious than the content of the paper. In particular the issue of user context is really not addressed in the paper (at least, from the ubiquitous computing perspective of context).

Also, you mention the goal of filtering spam, but that is not addressed in the paper either.

---

### Reviewer2:

<b>Familiarity 3</b> Rate your familiarity with the topic	1	2	3	4	
	Novice	Some knowledge	Familiar	Expert	
<b>Significance 2</b> Technical relevance and practicality of ideas in the paper	1	2	3		
	Not significant	Somewhat significant	Highly significant		
<b>Novelty 3</b> How original the problem and/or solution method is	1	2	3		
	Not novel	Somewhat novel	Highly novel		
<b>Quality of Presentation 3</b> Writing and presentation style/accuracy	1	2	3		
	Poorly written	Could be improved	Well written		
<b>Overall Recommendation 4</b>	1	2	3	4	5
	Strong reject	Weak reject	Weak accept	Accept	Strong accept

### Contributions

The paper analyses the issue of delivering events in the “last hop” to mobile users with constrained devices. It analyses the problems, proposes a somewhat sound solution, and thoroughly evaluate it.

### Strengths and weaknesses

The paper is good, well written, well motivated, and the proposal is evaluated in a very satisfactory way. Although I am not quite sure how the problem is really relevant in the context of pub-sub systems with respect to the problem of, say, delivering streams to mobile users in resource constrained devices.

### Detailed public comments

### Reviewer3:

<b>Familiarity</b> Rate your familiarity with the topic	1	2	3	4	
	Novice	Some knowledge	Familiar	Expert	
<b>Significance</b> Technical relevance and practicality of ideas in the paper	1	2	3		
	Not significant	Somewhat significant	Highly significant		
<b>Novelty</b> How original the problem and/or solution method is	1	2	3		
	Not novel	Somewhat novel	Highly novel		
<b>Quality of Presentation</b> Writing and presentation style/accuracy	1	2	3		
	Poorly written	Could be improved	Well written		
<b>Overall Recommendation</b>	1	2	3	4	5
	Strong reject	Weak reject	Weak accept	Accept	Strong accept

### **Contributions**

The paper presents and discusses various issues concerning the design of a publish-subscribe system for pervasive computing. The authors discuss different aspects related to the optimisation of the notification forwarding mechanism to resource-constrained devices and propose some possible solutions. The approach is not completely original, but the design choices are well motivated.

The authors provides a series of simulation results supporting their claims and design choices. In other words, they present the preliminary phase of their research project that will lead to the full implementation of the simulated publish/subscribe system.

### **Strengths and weaknesses**

The paper presents an interesting discussion and evaluation of the problem of the optimisation of the notification forwarding mechanisms to mobile hosts of a publish/subscribe system, taking into consideration possible different design choices. The presentation of the overall architecture of the system is rather obscure and the paper is not well-organised and, in some aspects, not well-focused, however the authors provide several convincing arguments and quantitative simulation results supporting their claims.

The approach of the authors is novel and interesting, however the data and input parameters used in the simulation seem not derived from real applications. Therefore, the conclusions reported by the authors may be wrong or not directly applicable to real situations. The simulations are carried out with sufficient scientific rigour and, in general, the analysis of the results is comprehensive and convincing.

The solutions proposed by the authors are not original or ground-breaking, however this paper provides interesting insights for the design of publish/subscribe systems for pervasive environments.

### **Detailed public comments**

The paper provides interesting insights for the design of publish/subscribe systems for mobile computing. Even if, the proposed solutions are not original, the motivations provided by the authors are convincing and supported by simulation results.

The approach proposed by the authors is presented in sufficient detail, but, in general, the paper is not well organised. For example, the description of the architecture is not clear. More specifically, the authors should provide a more detailed description of the infrastructure of the system (brokers, etc.). Another point that

should be better clarified is the description of the different operations performed by the the infrastructure and by the mobile devices.

Furthermore, the authors do not discuss the scalability of their solution. In fact, the computational load of the spoiling server may also be not negligible. In particular, it seems to me that if the forwarding mechanism is optimised and tailored for every single user, the required computational load is really consistent.

Sections 2 is misleading, since the remainder of the paper does not describe the concepts presented in this section especially with respect to code mobility. Section 3 is rather uncorrelated with the previous and following section.

In Section 3.3 the authors discusses the limitations related to device constraints; however it seems to me that in the remainder of the paper this aspect is not discussed anymore (how this context information is used for the optimisation of the forwarding mechanism?)

The description of the simulation test-bed is insufficient. The input parameters used in the simulations are not derived by measurements so they may lead to design choices that are not the best ones for real settings.

The need of changing the ranks of the notifications presented in Section 4.4 should be better motivated. It is not also clear to me if this is common case in real deployment scenarios or not (what about using different types of notifications?)

As far the discussion about the expirations is concerned, I think that this aspect is really application-dependent and therefore making any assumption in this direction in mobile environments is often impossible and, in general, quite difficult.

The pseudo code in Figure 7 should be commented (it appears uncorrelated with the concepts presented in the paper).

The authors do not provide any discussions about related work in the area. I suggest to have a look at other academic prototypes presented recently in the literature such as STEAM (Trinity College Dublin) and Pronto (Cambridge University) and commercial systems such as Broadbeam ExpressQ, IBM Websphere EveryPlace and iBus by Softwired.