Active Electronic Mail

S. Karnouskos
Fraunhofer FOKUS
Kaiserin-Augusta-Allee 31, 10589 Berlin, Germany
karnouskos@fokus.fhg.de

A. Vasilakos
Institute of Computer Science FORTH,
1385 Heraklion, Crete, Greece
vasilako@ath.forthnet.gr

ABSTRACT — Network infrastructures have evolved tremendously over the last years, offering new capabilities to the applications in higher levels. Email is a widely used communication tool that could benefit of an intelligent and active underlying network in order to support sophisticated services. We explore in this paper an infrastructure based on intelligent mobile agents and active networks, and point out how and where advanced features can be introduced to our current passive email platform in order to make it more flexible, open, secure, intelligent, and ubiquitous as possible.

Keywords — intelligent mobile agents, active networks, computational intelligence, email.

1. INTRODUCTION

The most widely used computer communication tool today is without doubt electronic mail (email), also called electronic messaging. Email has been with us since the earlier days of computer networks and became one of the three “killer applications”, along with *telnet* and *ftp,* that boosted Internet. Its popularity and ubiquity have established it as a communication standard worldwide. Yet, the last years, computer networks and the technologies attached to them have evolved. The underlying infrastructure has more to offer as it is more intelligent and service oriented. Email however, does not generally take advantage of the new advances in those domains. We will explore in this paper how technologies like computational intelligence, active networking, intelligent and mobile agents can be explored in order to advance the email infrastructure.

Computational Intelligence techniques are a set of heuristics that have several interesting properties such as learning, improvement through learning, flexibility, adaptation and abstraction. The backbone of computational intelligence constitute three technologies: neural networks, fuzzy set technology and evolutionary computing. Of course combinations of the above are possible such as neurofuzzy computing [25].

Active and programmable networks [1][31] is a relatively new concept where the network is transformed from a passive carrier of bits to a more general computation model. Active nodes are able to compute of data flowing through them and even allow user-injected code to modify, store or redirect that data. Although similar efforts exist via the introduction of packet filtering applications, web proxies, multicast routers, video gateways, RSVP, VLANs, RTP, application layer rooting etc, the idea is to replace all these ad-hoc approaches with an open generic capability of network programming. Active networks accelerate slow and expensive processes like network innovation since they eliminate the need for formal standardization as they support basic network services selectable on per packet basis.

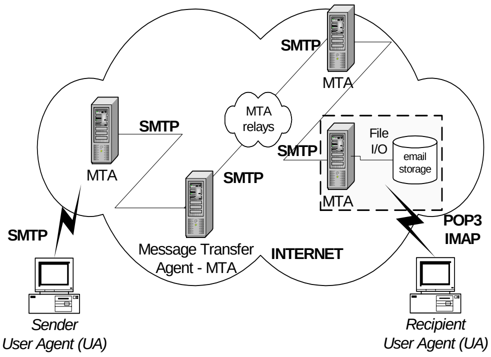

Figure 1 - Current email infrastructure

Software Agent technology [7] is a rapidly multi-developing area of research. Agents can be classified by their characteristics among which we are especially interested in mobility and intelligence. Mobile agent technology offers a new computing paradigm in which a program in the form of an software agent (intelligent or dumb) can suspend its execution on a host computer, transport itself to another agent-enabled host in the network and resume its execution in that host. The agents can act on behalf of a user and execute autonomously according to their internal goals. They provide robust networks as the hold time for connections is reduced only to the time required to move an agent. Furthermore they carry user’s credentials and therefore the connection is not tied to constant user authentication. Intelligent agents have emerged as a paradigm in which they represent a proxy for humans in interactions with computers. Standardization efforts for software agents exist with in the Object Management Group [32] and the Foundation for Intelligent Physical Agents [33].

2. MOTIVATION

Email is the defacto standard for user communication over the internet. It is used also to carry attachments, information content that can not be displayed in text, such as word processor files, video, audio etc. Reading emails while sitting in your office with a 100 Mbit connection has virtually no effect on the client-side of the service. However if one uses a slower dialup connection via a modem, then the frustration of downloading large attachments which at the end may end up to be useless, is not tolerable. Furthermore, many of the emails nowadays are Spam emails, from marketing companies that try to advertise their products over the Internet or even virus infected messages.

Reference: S. Karnouskos, A. Vasilakos, “Active Electronic Mail”, ACM Symposium on Applied Computing (SAC 2002), March 10-14, 2002, Madrid, Spain.

These unwanted email messages, that will be deleted anyway by the end user, introduce costs in the network side as they consume bandwidth, in the client side as they cost to be downloaded and jeopardize the security/safety of systems by propagating viruses. Even if the emails that reach the final user are useful in content, they might be useless if the user does not have the technical capability to fully display them to his device. Mobile users that use low capability devices such as mobile phones or PDAs might not be able to read the attachments because the format is simply not supported by their end-terminal. Here the need for converters comes into play. Additionally, since the end-devices are equipped with low processor/memory power it is desirable that this conversion and processing of user preferences is done within the network itself where high bandwidth links and huge processing power is a fact. Active networks promote computation within the network and this is exactly what we need. As also stated above we need to know the user's preferences and configure the network to react on behalf of the user. Therefore, computational intelligence is required. Intelligent agents can consult user's preferences and apply them to the network itself via computational intelligence techniques [25][26].

3. ACTIVE INFRASTRUCTURE

The current email infrastructure is depicted in Figure 1. Internet mail is a collection of protocols that are pulled together in client and server software. These protocols are mainly the POP3 [5], SMTP [3] and IMAP [6]. The basis of is a store and forward messaging model that does not rely on a persistent or even reliable network. The email is send from the *sender* to the node that hosts the *Mail Transfer Agent (MTA)*. That host finds via DNS (Domain Name System) queries the *IP address* of the recipient's host and attempts to deliver the email until this result to a success or the message times out. The recipient's host stores it to the *email storage* and awaits the *Mail User Agent (MUA)* to retrieve it.

The infrastructure we anticipate that the email platform will be running on, is compatible with the efforts today in the active networking area. Active networks, promote programmability of the nodes that compose the network infrastructure in various levels. Figure 2 depicts a general view of an active network node. The node hosts several Execution Environments (EE) such as ANTS [29], ALIEN [30], etc with different capabilities and architectures. The FAIN [28] project is currently developing an active network infrastructure similar to the one depicted here. For this paper, we focus our interest in the Mobile Agent execution environment where agents execute and offer their services to the applications. One of those applications is email. The future network infrastructure is a mosaic of active nodes and legacy nodes as we know them today. The active nodes will be able to compute on data they receive while the legacy ones will continue to pass on data as they do today. For a more detailed description of the infrastructure itself as well as its functionality, please check the previous work [2].

We anticipate the future email infrastructure as presented in Figure 3, which constitutes an evolution compatible with the current one as presented in Figure 1. The users are related to a *user context* where they have access to a variety of devices and are mobile. The nodes within the network are active nodes and host mobile agent execution platforms. For the nodes that are active but do not host a mobile agent EE, we assume that one of the other EEs offers the required services to the agents and can be remotely manipulated via e.g. RPC. The user end terminal can also be an active node or not. We also assume that there are various third party service providers whose service we take advantage of. Anti-spam, certification authority/security, component repository, user profile management, location and other internet-based services are supposed to exist in the future internet. The user must be able to exploit them for his own good. SMTP, POP3 and IMAP which are current default internet standards, are wrapped by agents which offer advanced extensions to these protocols or automatically fall back to the standardized mode if those extensions are not supported.

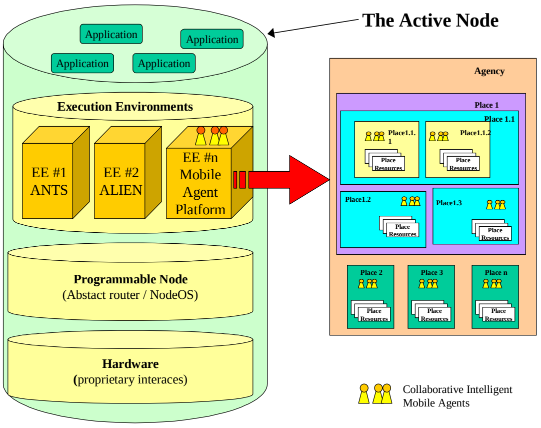

Figure 2 - The agent-based Active Network Node

4. ACTIVATING ELECTRONIC MAIL

Enhancements on the current email platform can be done on the server side, on the client side and also on the network components that rely between.

User Context Awareness. Today email is delivered to the user's email server. The mail client pulls the server, gets the email and then depending on its capabilities tries to present it to the user via a single device. However in the future, the emails will be pushed to the user according to his profile and his context. For instance lets suppose that you receive a multimedia email (with video and sound) but unfortunately your end-device is a PDA with no sound support. The network recognizes dynamically from your context that you also have a mobile phone. Therefore, it sends the video to your PDA and the sound to your mobile phone. The user's context include not only location but also capabilities of devices around the user that can be controlled. Therefore the network should be aware of the *user context* i.e. dynamically discover the user's position, preferences, devices and of course the capabilities of this specific context that specific moment.

The management and acquisition of this kind of info can be delegated to the agents. By querying localization devices such as active badges, GSM phones etc they are in position of estimating the user's position and maintain the dynamic changes to user's profile database. Based on the profile stored in this database, the agents residing on various nodes within the network are able to adapt their behavior while processing the emails directed to that specific user. The dynamic update of user's context is difficult, because as most matters that need human interference it will be neglected and outdated. Therefore, technology has to take as much responsibility as possible from the user and delegate it to intelligent agents which will act autonomously.

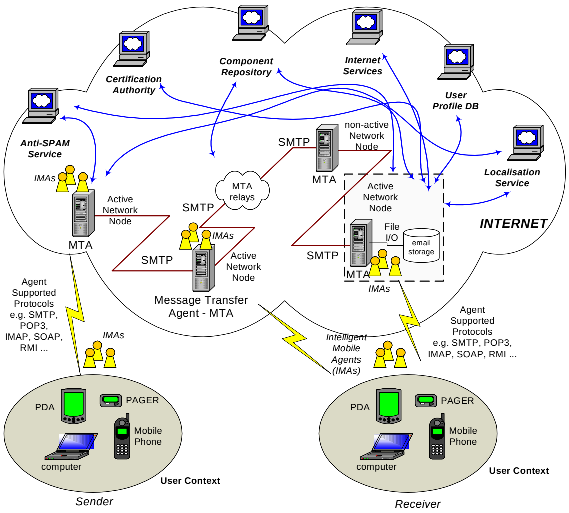

Figure 3 – Active Email Infrastructure

Distributed Anti-Spam. *Spam* emails are generally defined as unsolicited messages to many people. This type of unwanted communication has been flooding around the Internet the last years. Beyond reducing the bandwidth by filling up communications links, they also create financial damage to many companies by hindering corporate users from their daily tasks while they try to read and delete them. There are some tools that allow filtering of emails on the client side and also some server-based ones, however these tools are not personalized and are based mostly on keyword and IP address filtering. This is a start but it is awkward, non-intelligent and too complicated for the everyday user who just ignores it. As a result frustration is on a daily basis and this has several economic and social side-effects for companies. The problem is that users have little control over the email they receive and agent-based active networks are able to help with this matter. We can:

i) Filter emails that are passing through the node and drop the ones corresponding to well-known spammers and related sites. Real-time checking can be done today as there are sites that maintain blacklists of open mail relays, dialup lists and sites e.g. Dorkslayers [8], MAPS [9], ORBL [10], Osirusoft [11] etc.

ii) Filter emails based on user preferences, as they are defined on a profile database. The user may for personal reasons decide to deny emails coming from specific domains or people.

iii) Enable intelligent mail filtering and categorization based on advanced artificial intelligence techniques. More on this approach can be found below where we analyze email filtering in general for classification. Similar to that, emails recognized and classified as spam, trigger the appropriate actions e.g. delete, notification of the spammer's mail provider etc.

The combination of the above would be ideal, as it introduces customized black lists based on user preferences. These anti-spam tactics can be invoked on several places in the active email infrastructure. They can be invoked at the sender's MTA, but usually the spammers set up their own MTAs and large ISPs prefer loose enforcement of antispam techniques in the fear of accidentally losing emails due to false alarms. Therefore the best place would be within the network and in the other MTA relays. Even if some malicious MTAs allow spam the next "honest" MTA will stop the spam. Cases i) and ii) are the most appropriate here. Finally, the receiver's MTA is also a very good choice since it can apply the company's policy and honor the user's profile rules as stored in the profile database. Cases i), ii) and possibly iii) are most appropriate here. The gain we have by following our approach is that we are able to stop spam as close as possible to its source and remove offensive emails via official and customized rules, minimizing therefore their survivability.

Email filtering. This is done to some degree by most of the email clients today. However it is not flexible, requires technical skills and many people simply ignore it although they admit they need it. Filtering has to be more dynamic, based on the *user context* and of course be more intelligent and not only keyword based. In a dynamic scheme, I get emails when I want; from whom I want and to the devices I want, with minimal human intervention. This calls for intelligent techniques (e.g. a distributed neural net implemented by agents) that scan the email messages and, based on rules they possess, classify them. One sub-case is the classification of spam emails. However, the filtering of emails makes more sense when it is done on the MTA recipient server and not on the end user device, as the later may have inadequate computing resources. Furthermore, this helps nomadic users as they do not need to download non-relevant emails over the usually slow links via which they connect to company's intranet or can use web-based services e.g. web interfaces to read their emails from third party terminals.

Agent-based approaches like Maxims [14], Re:Agent [12], MailCat [13], and others [15][16], try to intelligently handle classification of text-based information such as emails. However this has to be done in a more flexible for the user way e.g. in natural language. For instance, System X [19] is able to reformulate questions in natural language into structured SQL queries and extract the data from a database. The combination of Natural Language processing, Multiagent systems and Computational Intelligence is still in a novel level but techniques could be imported from the internet information retrieval domain [18].

Mail Storage. The received email is usually stored as text. This makes tasks such as search, categorization etc very difficult and complicated. Almost everyone keeps emails in the incoming mailbox for future reference or as reminders of the tasks they have to fulfill. Therefore, it would be useful if the emails of the user could be used as a knowledgebase. Emails should be stored in a universally accepted format so that applications can have easier and better results while processing them. One way would be to store it as XML [35] files in a database where structured queries can be done. XML comes with a variety of tools and is driven by the vision of a universal homogeneous view of information.

Mail Notifications. Currently Email is passive. You have to check your email (pull-method) and see whether you have new messages. However, activeness is required. The email infrastructure should notify you when something happens (push-method) based on your profile. The agents can intelligently monitor your mailbox and based on your preferences send notifications or the email itself when it is needed. In this way, the interaction with the user is dynamic and may spawn agent-wrapped services e.g. Instant Messaging [36] or even automatic replies. Furthermore, this approach scales well for integration of multiple user mailboxes in various ISPs without user intervention.

Mobility and technologies that support it will play a key-role in the near future. With the emerge of mobile and palmtop computing, the support of nomads via intelligent routing and downloading of information such as email to pagers, cellular phones and palmtop computers is mandatory. People want to be offered seamless access to their emails from any location. Today's email infrastructure does not take care of mobility and its requirements e.g. location awareness, limited bandwidth, device dependencies and capabilities and this has to change. By using active and programmable nodes we are able to change the route of information for a user specific flow in various levels e.g. the application level. If an active network node does support queries to the user profile, it can consult it in order to redirect the email to the current location of the user, convert it to another media form e.g. fax or SMS (short message) or even drop it because this e-mail's usefulness was depended to a user specific geographic location. The reprogramming of the network nodes can be done via mobile agent technology in an asynchronous and flexible way.

Security has to be integrated in an easy-to-use fashion. A security infrastructure exists today, but it is cumbersome and as usual, if something requires some level of technical expertise or is a little bit complicated, the user will not bother to use it. The agents can automate tasks like signing of emails, signature verification, local certificate updates, search for a user's public key in various Certification Authorities' LDAP servers etc. For instance, for nomadic users to whom connect time and bandwidth matters, agents could verify online e.g. via OCSP [22] the signed emails they receive and pass them to the user with a verification status, therefore the user does not have to stay online while this procedure takes place. Furthermore, the agents can be used to implement wrappers for legacy services e.g. secure tunnels over untrusted infrastructures via SSL/TLS so that the emails that traverse the network are not sniffed by malicious users. Today acquiring emails and their contents is extremely easy with automated tools like ethereal [21] or dsnif [20]. We do not claim that by wrapping existing services and transparently securing them is a robust solution, but it minimizes significantly security and privacy risks over generally non-trusted infrastructures such as Internet.

Mail Integration / Openness. It is not cool to be different, at least not when Internet computing comes to the play. If I try to move my mail folders from one mail client to another, I always bump to walls of different-format representation and storage. There are homegrown isolated solutions (mostly a bunch of scripts) in the Internet community that offer client-to-client and between email clients and services integration. However, this is too complicated for non-technical users and that can change. The agents can implement the converters to a generally acceptable format e.g. in XML, or even negotiate about the capabilities and final representation of a service and its data [27].

Intelligent Mail handling. Today the email propagation is done hop-by-hop following the store-and-forward model between the servers. Processing of the content of the email is done sometimes at the server side and most usually at the client side. However active networking supports computation within the network which unveils new opportunities to process emails within the network itself and not only at its endpoint. The new model supported by active networking community, namely the store-compute-forward model must be exploited. Computational intelligence techniques for detecting and minimizing spamming, getting real-time info from the network, updating transparently email platform's components, rerouting information and notifying the users are some of the capabilities that now the network has and not just the endpoints. The applications are able to interact with the intermediate nodes within the network and customize them according to their requirements. This can result to optimization of the application itself as the network adapts to task-specific requests.

Mail Content Conversion. Since the user today does not use a standard terminal rather a set of devices with different capabilities both in software and hardware, we need to be able to display all email messages to that heterogeneous infrastructure. The ultimate goal is to realize "information any time, any place and in any form" as this is envisaged in the Virtual Home Environment [23] concept. This calls for conversion of communication media [17] (email content is a subclass) as well as dynamic application management. This is also a key-issue for expansion of multimedia email across the internet. Since active networks allow computation on data they receive, we can actually process the email content within the network depending on our preferences. Content transformation requires the existence of specific code on the active node that will perform this task. Mobile agents can be the vehicle that transports the specific piece of code to the node or can even wrap a remote service that offers this capability. If the agent on the active node that tries to deliver a multimedia email in your PDA recognizes from your profile database that your device is not capable of supporting color videos nor a specific format, it can convert it to another format, or download the necessary plugin to your PDA, and in parallel drop some frames from the video in order to reduce its size and minimize the download time (and cost) over slow links. Mobile agents do offer a promising asynchronous alternative to code deployment and handling.

Dynamic Content. Most of the content within an email is static (except from the links to external websites). However, parts of it are desired to by dynamic and actually project content, based on events. For instance, nowadays it is common to use vCard [4] to automate the exchange of personal information. If I check my mail archive and actually want to contact you, I can get your coordinates from the vCard which will be wrong if you have changed the company you work for. Therefore, the vCard should be dynamically composed and maybe provide a hint of the new email address/coordinates that you currently have. An agent can make a series of queries to certain directories and lookup alternative addresses before displaying the final message to the end-user. Furthermore in an active infrastructure this can be done prior to the user download of the emails, therefore reducing the communication overhead between the user and his home network (the overhead is transferred to the server side which we assume has enough computational resources). Dynamic content that depends on time, location, user or any other preferences is the next stage on the evolution of our email messages.

Minimizing complexity. The agents over active networks have increased capabilities. They can do tasks delegated by their user transparently, therefore decreasing the complexity of user-computer interaction. Reactive systems like the Channels [24] project issue a high degree of complexity in order to control communication via email. However many of these tasks can be handled automatically by an agent and this decreases the technical skills on the user site.

5. CONCLUSION

We have presented an advanced infrastructure that is based on active and programmable networks as well as intelligent mobile agent technology. Subsequently we have proposed enhancements that could be made to current email platform in order to take advantage of the underlying network. Areas like anti-spam techniques, filtering, security, mobility, content conversion and dynamicity can now be addressed within the network itself and not only at the end nodes. This brings advantages for the email platform itself and for the end users.

Active and programmable networks may still be an area where researchers experiment, but the interest to the capabilities it offers is increasing. Intelligent mobile agent technology on the other hand is with us quite some years, with an increasing popularity that will finally make it a tight part of the Internet [34]. We believe that the combination of both technologies is a powerful one that can offer to network applications, including email, the hooks they need to control the network itself for their own goals.

6. REFERENCES

[1] Active Networks at DARPA http://www.darpa.mil/ito/research/anets/

[2] S. Karnouskos, Realization of a Secure Active and Programmable Network Infrastructure via Mobile Agent Technology, Special Issue on Computational Intelligence in Telecommunications Networks, Computer Communications Journal, Volume 25, Issue 16, pp. 1465-1476, October 2002.

[3] Simple Mail Transfer Protocol, RFC 821

[4] vCard Specification, RFC2425 and RFC2426.

[5] Post Office Protocol - Version 3, RFC 1939

[6] Internet Message Access Protocol v4, RFC 2060

[7] Cetus Links on Mobile Agents : http://www.cetus-links.org/oo_mobile_agents.html

[8] Dorkslayers blacklist, http://www.dorkslayers.com

[9] Mail Abuse Prevention System (MAPS), http://www.mail-abuse.org

[10] Open Relay Black List (ORBL), http://www.www.orbl.org

[11] Osirusoft Confirmed Spam Sources blacklist, http://relays.osirusoft.com

[12] G. Boone, Concept Features in Re:Agent, an Intelligent Email Agent, Autonomous Agents 1998, Minneapolis USA.

[13] R. Segal J. Kephart, MailCat: An Intelligent Assistant for Organizing E-Mail, Autonomous Agents 1999, Seattle USA.

[14] P. Maes, Agents that reduce work and information overload, Communications of the ACM, July 1994

[15] W. Cohen, Learning rules that classify email, Proceedings of the 1996 AAAI Spring Symposium on Machine Learning and Information Access.

[16] T.R. Payne and P. Edwards, Interface agents that learn: An investigation of learning issues in a mail agent interface. Applied artificial intelligence 1997

[17] T. Pfeifer, Automatic Conversion of Communication Media, Ph.D. Thesis, TU-Berlin, GMD 2000.

[18] V. Keselj, Multi-agent systems for Internet information retrieval using natural language processing, Master Math Thesis 1998, University of Waterloo, Canada.

[19] N. Cercone, P. McFetridge, F. Popowich, D. Fass, C. Groeneboer, and G. Hall. The SystemX Natural Language Interface to Relational Databases. Sixth International Conference on Artificial Intelligence and Expert Systems Applications, Houston, TX, 1-2 December 1994.

[20] Dsnif, http://naughty.monkey.org/~dugsong/dsniff/

[21] Ethereal, http://www.ethereal.com/

[22] X.509 Internet Public Key Infrastructure Online Certificate Status Protocol – OCSP, RFC 2560

[23] ETSI Draft 22.70: Virtual Home Environments, 1997, Sophia Antipolis, France.

[24] R. J. Hall, How to avoid unwanted email, Communications of the ACM, March 1998.

[25] W.Pedrycz, A.Vasilakos, Computational Intelligence in Telecommunications Networks, CRC Press, Sept.2000

[26] A.Vasilakos, K.Anagnostakis, W.Pedrycz, Application of Computational Intelligence techniques in Active Networks, Soft Computing, vol 5,issue 4,2001,pp264-271.

[27] FIPA Agent Communication Language (ACL) http://www.fipa.org/specs/fipa00061/

[28] Future Active IP Networks (FAIN), http://www.ist-fain.org/

[29] D. J. Wetherall, J. Guttag and D. L. Tennenhouse, ANTS: A Toolkit for Building and Dynamically Deploying Network Protocols, IEEE OPENARCH'98, San Francisco CA, Apr. 1998.

[30] D. Scott Alexander, ALIEN: A Generalized Computing Model of Active Networks, Ph.D. Thesis, University of Pennsylvania, December 1998.

[31] S. Karnouskos, H. Guo and T. Becker, Trade-off or Invention: Experimental Integration of Active Networking and Programmable Networks, Special Issue on Programmable Switches and Routers, IEEE Journal of Communications and Networks, Volume 3, Number 1, pp 19-27, March 2001 (ISSN 1229-2370)

[32] OMG Web Site : http://www.omg.org/

[33] FIPA Web Site: http://www.fipa.org/

[34] D. Kotz, R. Gray, Mobile Agents and the Future of the Internet, ACM Operating Systems Review, Aug 1999.

[35] Xtensible Markup Language (XML) and other family technologies, http://www.w3.org/TR/

[36] Instant Messaging - Jabber, http://www.jabber.org/