Security Issues in E-Home Network and Software Infrastructures

Almut Herzog, Nahid Shahmehri,
Andrzej Bednarski, Ioan Chisalita, Ulf Nordqvist, Levon Saldamli, Diana Szentivanyi, Måns Östring
Linköpings universitet
almhe@ida.liu.se

Abstract

Many people enjoy an Internet connection even at home. These private connections are becoming cheaper, faster and more stable resulting in more and more private PCs being connected at all times. A permanent Internet connection offers new possibilities for Internet services targeting the private home. But it also creates new risks for the household. The Internet is a new door into the house that can be exploited by hackers. They might learn about the behaviour of household members by eavesdropping on their Internet traffic or they can gain control over network nodes and access sensitive information in the house.

In this paper, we take a look at ways of securely connecting a private home to the Internet. We give an overview of existing technology for connecting more than just PCs to the Internet and describe proposed network infrastructures for the home. We also offer an outlook on the emerging area of electronic services (e-services), an area that differs substantially from the propagated e-business of today and gives a better reason for permanently connecting one's home than occasional web browsing or e-mailing.

This paper introduces network and security terminology as well as a prospect for future homes with requirements on the software platform needed for e-services.

1. Introduction

Today, many people have one or more PCs at home. They often use a dial-up connection to the Internet through a modem with a transmission rate of up to 56Kb/s or up to 1.5Mb/s with ISDN. Cable modems using cable TV lines – with up to 2 Mb/s – are evolving and offer permanent Internet connection for the home at a reasonable price. Still, a home Internet connection is mostly used for Web browsing, e-mail with colleagues and friends, newsgroup activities, some Web shopping or exchange of e.g. MP3-music. However, with the arrival of these faster and inexpensive permanent connections, the security awareness of their end users needs to increase. Previously, a dial-up connection to an Internet Service Provider (ISP) resulted in the assignment of a dynamic IP-address to the dialling computer. A hacker had to guess this IP-number and hack the connected computer before the connection is terminated by hanging up the modem. With a permanent connection, hackers have more time to act. While the risk of such a hacking attack might deter people from getting a permanent connection, such a decision would also exclude them from the next frontier in the world of Internet: e-services.

E-services (electronic services) rely on a permanent Internet connection and offer broader possibilities to their customers than simple web browsing. E-services can be utilities that enable home owners or tenants to remotely control their house or apartment by checking the alarm system or peeking into the refrigerator. In addition, e-service software is designed to allow controlled access or service to the home by a third-party company, a so-called service provider (SP). Such an SP can be responsible for e.g. meter reading, for uploading new software for particular devices or for communication with or surveillance of patients that are treated at home. E-services need not exclusively address the private home; they can as well be used to facilitate access between small offices and their central office or remote workers and their home office.

A complement to this development is mobile Internet which offers such services and more on mobile devices. An overview of technologies and applications for this are given in e.g. Tarasewich and Warkentin [1].

In this paper, we present requirements and prerequisites for connecting a small home or office network to the Internet in a secure way while enabling it for e-services. This comprises network and software requirements which we describe and exemplify.

2. Connecting the Home to the Internet

2.1 What to connect

It is certainly nice to do web surfing at home and at a reasonable speed, but there is more to the Internet than applications that require a PC at home.

Very soon all kinds of appliances will be network-enabled. There are some products already with Electrolux' Internet-enabled refrigerator Screenfridge [2] or Merloni's washing machine Margherita2000.com [3]. Set-top boxes are currently competing in offering Internet integration by e.g. allowing Web browsing with the TV remote control or download, replay and integrated payment of music or movies [4].

On the research side, MIT Medialab presents the coffee-machine Mister Java, the smart microwave PC Dinners, and the network enabled kitchen counter Counter Intelligence [5].

By using non-IP-based networks, such as CEBus [6], LonTalk [7] or the older and less convenient X-10 standard, it is possible to control electrical devices such as lamps or fans.

While appliances for the networked home are being developed, the home network itself is still very much under discussion [8]:

The A/V industry wants to distribute digital real time audio and video throughout the home needing a high-speed network. The computer industry foresees multiple PC's, shared PC peripherals and home electronics throughout the home connected on a preferably IP-based, moderate speed network, which should use existing cables like cable TV or telephone lines. The home systems industry dealing with security, home automation, HVAC (heating, ventilation, air condition), and appliances lobbies for highly reliable sense and control networks at low to moderate speeds (<1Mb/s) that control all electrical appliances in the house, preferably using the existing power lines.

However, in order to connect to the global Internet, the home network must be able to interface with the IP-based network of the Internet. Thus, we assume an IP-based or at least IP-compatible home network when we now discuss ways of connecting the home network to the Internet.

2.2 How to connect the home to the Internet

The first way one thinks of is to connect every device directly to the Internet (c.f. Figure 1a). The Internet Service Provider (ISP) of the household reserves a certain amount of IP addresses for the home that can be assigned to any IP-based device on the home network like the PC, set-top box or A/C controller in Figure 1a. The household uses a simple hub to multiply the network outlets for the house. A router at the ISP site forwards IP-packets to the addresses in the house and plays an important role by implementing e.g. security policies that apply to its connected households.

With the coming of IPv6 [9], which solves the shortage of the IP address space, this architecture is technically feasible. The major disadvantage is that network addresses in the house depend on the ISP. Should the homeowner want to change ISP, network addresses need to be reset for all networked devices in the house. Even though it is technically possible to keep a once assigned address, this would mean that exception addresses would have to be inserted in already strained Internet routing tables; something which is not good in the long run.

This problem is solved by a *private network* [10] within the house (c.f. Figure 1b). Generally, a private network means that a network with private IP-addresses – often an Intranet – can be connected to another network – usually the Internet – through a gateway that implements network address translation (NAT) [11, 12]. Such a gateway transparently translates private IP-addresses into public addresses and vice versa. According to [10], private IP addresses cover the address space from 10.0.0.0 to 10.255.255.255, from 172.16.0.0 to 172.31.255.255, and from 192.168.0.0 to 192.168.255.255 and can be used by anyone whose network is not directly connected to the Internet. A minimum of one public IP address is needed for a whole private network. It depends on the configuration of NAT and the number of public IP addresses assigned to the private network if nodes on the private network are directly accessible from the outside Internet or not, i.e. if inside addresses are at all visible from the outside. If they are not, the inside network cannot be detected by network scans and thus not easily be targeted by hackers.

The security of the network in Figure 1b depends very much on the configuration of the gateway including NAT that connects the home network with the Internet. Such a gateway must be set up with firewall properties that allow only certain types of traffic according to – sometimes complex – policies. A discussion of firewalls and policies can be found in [13].

Generally speaking, if devices are *directly* connected to the Internet (as in Figure 1a and possibly 1b), they have to implement firewall functionality that protects them and their network surroundings from unwanted traffic i.e. hacking attacks. In our homenet scenario this means, that if one device on the homenet is directly reachable from the Internet (by either a direct connection or a configuration of NAT), this means that this device must implement a firewall. If the firewall implementation or configuration of either the residential gateway or such a directly connected device is flawed or incomplete, this endangers not only the homenet but potentially other computers on the Internet: a successful hacker might choose not to cause havoc in the house but to make compromised machine(s) on the home network participate in an e.g. denial-of-service (DOS) attack on other Internet nodes.

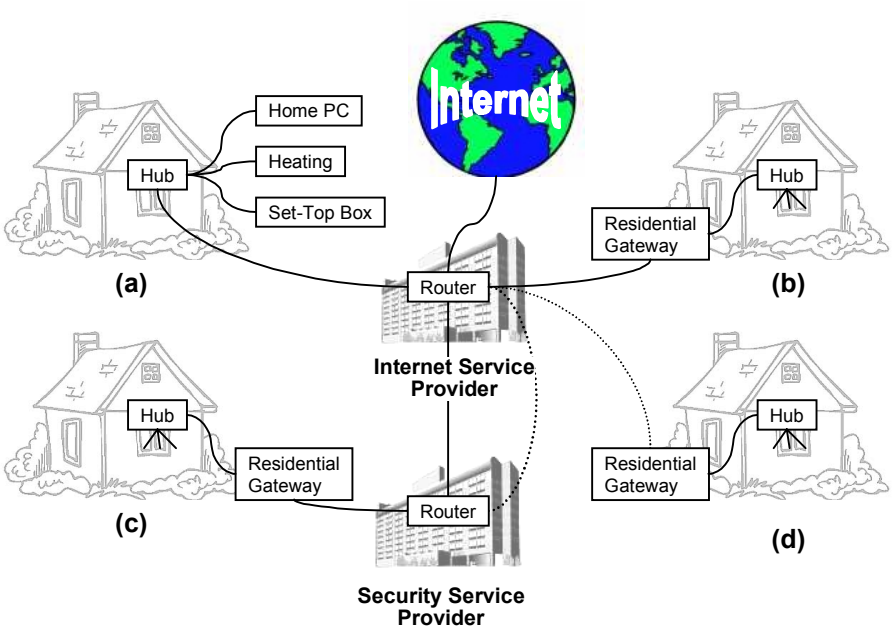

Figure 1: Four ways to connect the home to the Internet

(a) direct connection of all devices; (b) private network protected by residential gateway; (c) private network protected by residential gateway and security provider using proprietary network between security provider and home; (d) private network protected by residential gateway and security provider using VPN technology (dotted line).

As shown above, a correct set-up of the firewall on the gateway is of great importance. A business idea might be to relieve the homeowner from this difficult task. In Figure 1c a checkpoint is inserted between the private home network and the insecure Internet. Instead of directly connecting to the Internet, the home gateway connects to a service provider (SP) that manages the Internet access to and from the home. We call such an SP a *security service provider* (SSP) (c.f. Figure 1c and 1d) so as not to confuse it with an ISP, even though it might be reasonable for an ISP to offer services as an SSP. In fact, for the household of Figure 1c, the SSP acts as an ISP. One could say that such a security service provider performs another Network Address Translation between Internet and home network: not for the sake of saving IP addresses but for security reasons.

The network between the home and the SSP need not at all consist of an SP-owned line (as in Figure 1c). The access network might just as well be the Internet by e.g. using VPN technique (Virtual Private Network [14]) (c.f. Figure 1d) or by setting up the involved nodes' routing tables so that home networks are only reachable through an SSP router that does firewalling on behalf of the home.

Setting up firewall routers with NAT capabilities at an SSP site is a very strict approach, because it turns the SSP router(s) into bottlenecks and single points of failure. It also means putting all trust on this SSP that it does not audit the home users' private traffic.

A more convenient *hybrid solution* would force all traffic *from* the Internet *to* the home through the SP while home access to the Internet can be configured to be direct. This means that users wishing to access their home from remote have to authenticate at the SP site, probably be subject to auditing and then may access their applications at home. In contrast, web traffic from inside the house would not be subject to such restrictions. This can be realised using VPN technology and a smart configuration of routers and the home gateway.

2.3 Discussion

All of the proposed network infrastructures have their pros and cons. We will now discuss the three aspects of computer security – confidentiality, integrity, and availability [15] – for the four proposed models.

Confidentiality is defined as the "prevention of unauthorised disclosure of information" [15] and includes the aspect of privacy. An attack at confidentiality in an e-home environment would e.g. mean that an intruder learns when people leave the house, what is in the refrigerator, what files reside on the hard disk of the home PCs. The exposure to this risk depends on the way user authentication and access control is performed. In Figure 1a and 1b such a control is performed in the house. In 1a it might even be performed by each connected device. Figure 1c and 1d introduce a barrier at the security service provider site. Only after that barrier has been passed access is granted to the homenet which might use a second

layer of authentication if needed or desired. This barrier prevents any exploration of the home network by outsiders, whereas directly connected devices can be hacked until successful.

The protection of integrity – "the prevention of unauthorised modification of information" [15] – is more application than network dependent. In all proposals of Figure 1, a hacker can theoretically eavesdrop on traffic to or from the homenet and modify or spoof data. This can be prevented by e.g. sending encrypted data between the two communication endpoints but not by using a certain network infrastructure.

Availability including reliability deals with the "prevention of unauthorised withholding of information or resources" [15]. In practice, protection against e.g. denial-of-service (DOS) attacks and power failures belongs to this domain. The proposals of Figure 1a and 1b are directly exposed to the Internet and therefore also potentially exposed to direct DOS attacks. If a hacker has set his or her mind to flooding the home network, the home owner's only chance is to contact his/her ISP for help. In the case of a security service provider managing the access to the home network, the risk of a DOS attack aimed at a private home is small because no IP-packets must reach a home network unless authorised. Instead, a DOS attack aimed at the security service provider might cut off *all* connected home networks during the attack.

The end user's decision for or against certain architecture depends on his/her outcome of a personal risk analysis that lists the assets of a household and their value for the household. Household assets are not only replaceable items like a PC, an appliance or software but also values such as reputation, privacy, and dependency.

Questions to be asked are: What is the worst case scenario for a compromised home? What devices and services are accessible to a hacker if all security measures fail? What is the worst damage caused by a compromised system?

If the only connected device is a web camera in the refrigerator, the experienced damage of lost privacy might be low compared to the damage caused by a compromised alarm system, which mistakenly unlocks a door for burglars to come in.

How badly is a particular service needed? How much may security cost? How much time can be spent in an emergency situation or to keep up-to-date with software and security solutions?

It might be more cost-effective to delegate such management to a service provider with only minimal time and effort needed by household members. The more a household depends on a service, the more likely it will pay to assure that the needed service is available at all times.

2.4 Example solutions

We will now take a look at two examples that illustrate the different kinds of network infrastructures.

Ericsson [16] offers an example of an SP-controlled Internet-access (c.f. Figure 1c and 1d). An SP provides the home owner with a thin gateway that connects the Ethernet-based home network with this SP (optionally using the Internet with VPN-technology). The SP is responsible for initialising these home gateways, called *e-box*, and for the maintenance of the e-service software executing on the e-box. This architecture is chosen with the motivation to provide a convenient and secure way of remotely accessing home devices and having Internet-access from inside the house without the end user having to have a deep knowledge of networks and Internet security.

Another provider of a residential gateway infrastructure is Echelon Corporation with their i.LON 1000 Internet Server [17]. They assume no additional infrastructure, i.e. the gateway is directly connected to the Internet. Their gateway is intended to be used for managing and operating LonWorks devices that are connected to the gateway and managed through the power line with the LonTalk protocol [7]. They use an architecture as in Figure 1b, with their Internet Server as the connection point to the outside. The home network that this server can serve is not Ethernet but the power line. This architecture does not consider versatile electronic devices like PCs on the home network, only LonWorks devices are taken into account.

3. E-Service software and platform

In this section, we take a closer look at e-service software and its *execution platform*. Execution platform denotes the environment in which e-service software runs and communicates with other e-service applications. As mentioned in the introduction, e-services can be any kind of software that offers a service to its end users by using the Internet.

We start with a fictitious example application from a health care scenario to give an idea what an e-service can be. A home patient wears a pulse checker that reports the pulse rate once per minute during night time. The pulse checker sends its data to an application on a device on the home network, e.g. the gateway. This application collects and evaluates the received data. If it detects anomalies, it will alarm the house and – depending on the severity – also a health care provider. If everything runs normal, the application sends collected data to the health care provider once per day for backup purposes and for evaluation by health care personnel.

This example shows that many different platforms and network nodes can be involved in an e-service (c.f. Figure 2). A network-enabled pulse checker runs a pulse application. This application sends its acquired data regularly to another pulse application on the home gateway. The dotted line between the pulse checker hardware and the gateway denotes a radio network between them. The pulse application on the gateway checks the received data. If

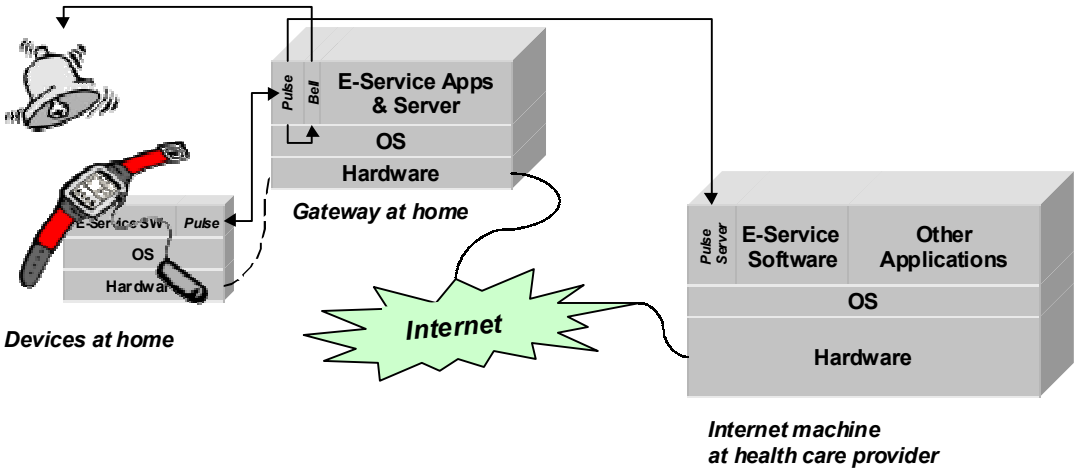

Figure 2: Pulse checker service

An example e-service. Lines with arrows denote message communication between involved applications. Other lines denote physical connections between platforms.

necessary an alarm is generated within the house (denoted by the bell in Figure 2) and passed on to a pulse server application at the health care provider site. Note that the bell is not a networked device. Instead, the e-service bell on the gateway "knows" how to ring a bell by switching it on and off. The bell, unlike the pulse checker, is thus a device *without* an IP address. In practice, the bell might be connected with the gateway through the X-10 or LonTalk [7] protocol for power line transmission, and the bell application on the gateway uses that protocol for communication with the bell. Once per day, the pulse application on the gateway sends accumulated data to the pulse server application at the health care provider site – through the Internet and possibly a security provider. This pulse server application performs evaluations with these data received from multiple serviced homes and stores them for later use.

Note that the gateway at home does not only function as a platform for running services but also as a platform for distributing e-service software within the house. Each time the pulse checker hardware is switched on, it checks with the e-service software server on the gateway if there is new software for this specific pulse checker device. If so, the software is dynamically downloaded and deployed on the pulse checker using e.g. JINI technology [18].

Why is the gateway and not another home network node used for running applications and serving home devices? For once, the gateway is always turned on in order to allow Internet access to and from the house. This cannot be said of a home PC. Secondly, the gateway is the only device in the house that connects both to the homenet *and* the access net; it is thus accessible by both the home devices and machines on the access net (i.e. potentially the Internet). In addition, the gateway's only function is to allow and support e-services in the house. It is therefore feasible to let it manage and maintain the services as well. Furthermore, if e-service software on the homenet is to be installable, configurable and updateable by an SP and not only by the end users, it is reasonable to put all software on this one server on the homenet that is accessible from outside.

Why not leave all the e-service software at an SP site? If software resides at an SP site only, end users always have to connect to that SP to run their software. This might be an excellent – simple and easy-to-maintain – solution for Intranets or Enterprise networks that have an SP node on their local network. For in-house TV or Video it is not suitable, though. In-house related traffic would be forced to leave the house, needlessly congesting the access network.

As we have already seen in the pulse checker example, e-service software for one service can execute on a number of platforms and is highly distributed: software for one e-service might involve modules that run on a home device, the home gateway and an SP site. This poses a number of requirements on the software and its platform.

1. E-service software shall be able to execute on all kinds of operating systems and platforms.
2. It shall be small and undemanding in RAM and CPU requirements to enable it to run even on embedded systems like the pulse checker in our example.
3. The platform shall allow for automated updates from remote but also for updates on demand of authorised users.
4. The platform shall allow for version negotiation. A node that has not been updated yet must be able to communicate with already updated nodes and vice versa.

5. The platform should be resistant to attempts to crash it or to crash an e-service on that platform.
6. The platform shall guarantee that no e-service is starved by the execution of another service, i.e. it must provide a quality of service guarantee.
7. The platform shall guarantee that a module belonging to e-service *A* can communicate *privately* with another module of e-service *A*, even if that module runs on another network node. E-service *B* shall not be able to eavesdrop on or to fake communication to e-service *A*.

An architecture that fulfils some of the requirements is SUN's Java Embedded Server (JES) [19], which implements the OSGi specification [20] for e-service software. JES is a small server, quite similar to a web server. It allows the connection of http clients, i.e. Web browsers. These clients can download servlet-generated html-pages. In addition to its web server properties, JES is designed to manage software that runs in the JES service space i.e. it allows for version management and interdependencies of the servlets and their libraries called *bundles* that are accessible from outside. Even though JES needs the Java Runtime Environment (JRE), it still fulfils requirements 1 and 2. Specially designed for version management and updates it nicely meets requirements 3 and 4. The fulfilment of requirement 5 depends on the stability of JES, JRE and the operating system that runs the e-service software. Given the requirement of platform independence it might be difficult to secure all three parties (JES, JRE, OS). Requirements 6 and 7 are not addressed by JES.

4. Conclusion and Outlook

In this article, we have described four kinds of network architectures that will connect a home to the Internet. The architectures range from simple and unsecured, to advanced with security checks at more than one network node. We have shown an example of a software infrastructure that enables the private home for e-services that are not bound to a home PC but can involve all kinds of networked devices.

Research and development is needed in the area of residential gateways and security issues, e.g. on how to guarantee that e-services do not – maliciously or incidentally – interfere with each other. Prevention of interference builds on resource management, which remains a difficult task. Operating system vendors fear its impairment on performance while platform-independent middleware like JES cannot efficiently do it.

In the end, there is the question if end users are security-aware enough to understand the implications of connecting their home to the Internet. Will they be willing to accept a proposed network infrastructure for e-services? Will they be ready to open their house to the Internet? Will there be service providers that develop e-services for them and help them manage security?

E-service systems still have to prove that they are secure and stable, but this does not stop them from evolving.

5. Acknowledgement

This paper has been written as part of the STEM (Software Technology and Methodology) Project Oriented Studies (POS) in fall 2000 [21] within the ECSEL Research School of Linköpings Universitet. The goal of this project was to control home devices such as lamps or radios and to show images from a Web camera on a Web client and/or PDA using Ericsson's software infrastructure [16]. In addition, security issues related to the area of e-home software were explored.

6. References

[1] TARASEWICH P., WARKENTIN M. *Issues in Wirless E-Commerce*. SIGecom Exchanges, Newsletter 1.1, Aug. 2000. http://www.acm.org/sigs/sigecom/exchanges/

[2] ELECTROLUX. Screenfridge. http://www.electrolux.com/screenfridge/

[3] MERLONI. *Margherita2000.com*. http://www.margherita2000.com/index/

[4] ALLNETDEVICES. *News and features for developers and vendors of Net devices, content and applications*. Set-top boxes. http://www.allnetdevices.com/devicespecs/set_top_box/

[5] MIT MEDIALAB. *Counter Intelligence*. http://www.media.mit.edu/ci/

[6] TIA, CEBUS TSC. *EIA-600 CEBus SET*. 1992. http://www.tiaonline.org/standards/

[7] ECHELON CORP. *LonWorks protocol.* http://www.lonworks.echolon.com

[8] HOME AUTOMATION & NETWORKING ASSOCIATION (HANA). *Standards*. http://www.homeautomation.org/standards.html

[9] DEERING S., HINDEN R. *Internet Protocol, Version 6 (IPv6) Specification*. IETF – RFC 2460. 1998. http://www.ietf.org/rfc

[10] REKHTER Y. ET AL. *Address Allocation for Private Internets*. IETF – RFC 1918. 1996. http://www.ietf.org/rfc

[11] EVANG K., FRANCIS P.. *The IP Network Address Translator (NAT)*. IETF – RFC 1631. 1994. http://www.ietf.org/rfc

[12] SRISURESH P., HOLDREGE M. *IP Network Address Translator (NAT) Terminology and Considerations.* IETF – RFC 2663. 1999. http://www.ietf.org/rfc

[13] SEMERIA C. *Internet Firewalls and Security – A Technology Overview.* 2000. http://www.3com.com/nsc/500619.html

[14] GLEESON B. ET AL. *A Framework for IP Based Virtual Private Networks*. IETF – RFC 2764. 2000. http://www.ietf.org.rfc

[15] GOLLMANN, D. *Computer Security*. Wiley&Sons. Chichester. 1999.

[16] ERICSSON. *E-box*. http://www.ericsson.com/wireless/products/ebox/

[17] ECHELON CORP. *i.LON 1000 Internet Server*. http://www.echelon.com/Products/ilon/default.htm

[18] SUN MICROSYSTEMS, INC. *JINI*. http://www.sun.com/jini

[19] SUN MICROSYSTEMS, INC. *Java Embedded Server.* http://www.sun.com/software/embeddedserver/index.html

[20] OSGI. *OSGi Service Gateway Specification*. http://www.osgi.org/about/spec1.html

[21] STEM POS, GROUP 1, 2000. *HomeCore*. http://www.ida.liu.se/~stempos1/