The Personal Node (PN)

Gregory G. Finn and Joe Touch
USC/Information Sciences Institute
{finn, touch}@isi.edu
July 27, 1998

ISI/RR-98-461

Abstract

A Personal Node (PN) is a small, wallet-sized device that integrates people into the Internet. A PN incorporates wireless communication, limited user I/O, and local environmental telemetry to catalyze the coordination of other smart space and network devices for the user's benefit. By themselves smart spaces are not aware of the people in them and people are not be aware of what is in a smart space. The PN allows the smart space to interact continuously with a person, and a person to interact continuously with the space, mediating the interaction with the help of other devices throughout the system. A PN is an individual's networking focal point. As the user roams about, a PN persistently maintains user presence on the internetwork. This represents the final and missing link in smart spaces, bringing the user in as a system resource and participant.

1: Introduction

In the near future buildings, rooms and vehicles are expected to evolve into *smart spaces* that contain varieties of wireless *smart devices*. As individuals and vehicles roam they will swim in a virtual sea of such devices. Those spaces must become aware of the individuals in them and vice versa. A Personal Network Node (PN, or Personal Node, for short) is a small device that is continuously carried by an individual, **allowing the user to become a permanent part of the smart space**, and allowing the infrastructure ongoing access to the user for feedback and commands. The PN is the network corollary of the PC; it is a network node for the 'rest of us'.

A PN incorporates promiscuous wireless interfaces to allow it to act as a networking catalyst. PNs also include various telemetry sensors, such as location (GPS), temperature, and orientation. This allows a PN to be addressed by-interface, by-region, by-proximity or by-heading. The ability to form ad-hoc networks coupled with multiple addressing modalities offers considerable benefits for military, commercial and emergency services.

1.1: Need for a Sixth Sense

Wireless technology allows us to create smart spaces and devices; key to its success will be its ease of use. People do not have the capability of directly interacting with smart spaces. We are deaf, dumb and blind as far as smart spaces are concerned, and won't be aware when we are in a smart space or know what's in one. Conversely, smart spaces won't be aware of our presence. This mutual lack of awareness must be overcome before smart spaces can become well integrated into our daily lives.

In the distant future, smart spaces may include senses so that they can be aware of and interact directly with us. A more practical solution is for us to become aware of smart spaces and interact with them on their terms, and to extend them to interact with us continuously on ours; in effect, to carry that sense with us. We need to acquire a sixth sense that allows us to continuously inhabit, to 'see and speak' in this wireless spectrum.

One way to make us aware of our wireless surroundings as we roam is to carry a device that intermediates with smart spaces on our behalf. PNs see and speak for us in those areas of the wireless spectrum used by smart spaces.

This document outlines our vision of the PN, and how it uniquely enables smart spaces interaction via a continuous network presence for people. The remainder of this document is organized as follows:

- *Sec. 2: Defining a PN*
- *Sec. 3: Need for a PN*
- *Sec. 4: Research Issues*
- *Sec. 6: Related work*
- *Sec. 7: Implications*
- *Sec. 8: Summary*

1.2: A vision of smart spaces

The following example describes our vision of the future of smart spaces, and how it is affected by having people as a continuous part of the infrastructure.

Although this scenario superficially has much in common with similar visions (e.g. Active Badges), a fundamental distinction underlies the PN: this device serves as the user's "eyes and ears" to the Internet (i.e. the sixth sense). It is not intended to be a compute node, but rather the minimal set of I/O functions that a human needs to interact digitally.

> Consider a field operative, viewing a map with a monocular display. The monocular and its attached chordal keys provide primary I/O, while the PN provides location and orientation to direct the map to match the user's view.
>
> The user walks up to a truck, and the smart space within the truck signals to the monocular to hand-off the display. The display transfers to the heads-up display (HUD) or mounted monitor at the seat the user enters, as he sits. The user drives off, and the map zooms out as the vehicle accelerates, to provide a view more appropriate to the terrain as it moves by.
>
> The user's rucksack, being near his PN, is considered relevant by the truck's smart space, and is checked for rations. A supply truck passes by, and the user's truck asks the supply truck for a refill of water, food, and batteries, which are dumped by the roadside with a smart, encrypted beacon. The HUD points to the refill pack as the user's truck approaches, and the user's PN continues that function as a beacon-compass as the user walks over to pick it up. Because many people in the same group have also recently picked up their rations, the aggregate behavior triggers a shipment to be scheduled from the supply depot.

In the vision just described, the PN alerted the smart space of the user's presence, and allowed another user (via her PN and its proxy) to alert him an urgent query. This interaction did not require external terminals; simple interactions can be handled directly by the PN.

2: Defining a PN

A PN is a networking *vade mecum*[1] that is active whenever you are. It continuously maintains your presence in the network, and contains sufficient I/O capability to allow you to communicate until other resources are arranged.

1. *vade mecum* (latin) - lit. "go with me"

2.1: Basic Assumptions

For PNs to function successfully a number of conditions must be met. Most amount to assumptions regarding the way the Internet will evolve. We state these here:

- *Internetwork wireless services will be ubiquitous. Most buildings and metropolitan areas will provide such service.*
- *In-building networks will provide a basic level of Mobile IP [24] service.*
- *Exterior and interior wireless networking technologies will provide different levels of service. In-building wireless may offer greater bandwidth and lower latencies than exterior wireless.*
- *Smart-space computing will be increasingly 'user state' sensitive. Awareness of users and telemetry from them will become increasingly important.*

2.2: Design Principle and Content

The main design principle of the PN is, "**have only those capabilities that cannot be moved elsewhere**":

- *User I/O to bootstrap user-network interaction.*
- *All the I/O that's particular to the user's locale.*
- *Support for the above, i.e., wireless links, processing and memory for local operations and to bootstrap coordination elsewhere, and a battery that requires recharging "when the user does".*

A PN is a hand-held sized device, including:

- *A variety of wireless interfaces:*
  - Fast IR, for Mbps desktop roaming
  - Wireless LAN for 1 Mbps office roaming
  - Cellular for 100 Kbps MAN roaming
  - Satellite wireless for WAN roaming
- *A limited amount of user I/O:*
  - microphone, speaker
  - small LCD (PDA or smaller)
  - buttons (3-4, re-labellable; or a touchscreen)
- *As much telemetry I/O as possible:*
  - orientation: GPS, accelerometers, compass
  - environment: temperature, humidity, light / IR, sound, camera, EMI, pH
  - personal biometry: pulse, respiration (via sound processing), blood pressure, fingerprint
- *Support for the above:*
  - smart-card socket
  - CPU and memory (volatile and non-volatile)
  - battery for 48+ hours continuous "operation"

A PN is not a workstation or laptop. It does not include:

- *File storage:*
  - Fixed disk
  - Removable media
- *Traditional I/O:*
  - Keyboard and display

A PN mediates on behalf of its user with smart spaces that it encounters, making these smart spaces aware of the user and the user's environment. Conversely, a PN becomes aware of what is in the smart space environments. A PN also caches information of immediate or local relevance on behalf of the user.

To achieve these objectives, PNs must be carried by their users most of the time, much as pagers and cellular telephones are today. To be comfortably carried a PN must be small and lightweight. This precludes its having a standard keyboard or display. Unlike a wearable computer or laptop, a PN is not intended to replace a workstation [30]. It is assumed that a user's primary computing and storage resources are located elsewhere, and the PN acts as a bootstrap to catalyze the coordination of these other resources.

Compared to a PDA, a PN is a smaller device, intended for simpler network-to-user interactions, but with more advanced sensors and network access. A typical PDA is the size of a deck of cards, and limits its I/O to a fairly large display, stroke-based text input, and audio output (Figure 1).

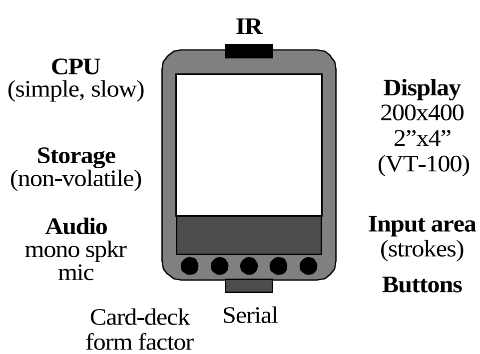

FIG. 1. PDA characteristics

A PN contains I/O peripherals that are intrinsic to an individual and that cannot effectively be located elsewhere. This includes user I/O, including microphone and speaker, control keys, and a limited display, and would eventually be the size of a large wristwatch (Figure 2). It also includes sensors that monitor you and your local environment. Sensors would include GPS to determine your location, electronic compass, accelerometers, photometer, barometer and biometry. Unlike PDAs, PNs are IP-addressable, operate continuously, autonomously and communicate with their surrounding wireless environment.

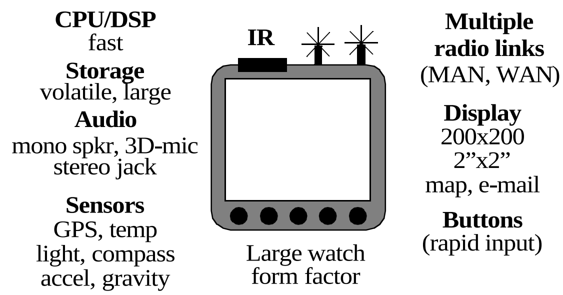

FIG. 2. PN characteristics

There are also a variety of research issues involved in the development of a PN. Primary among these is power conservation, which drives hardware design as well as protocols to support multi-capability signalling. The continuous connectivity of the PN presents challenges to the traditional network model of a node (host), as well [4].

2.3: How is a PN different?

The general concept of a PN is related to that of the ParcTab. The PN is different from either a ParcTab or PDAs in that it:

- *Supports continuous internetwork presence*
- *Is I/O rich: multiple wireless interfaces, audio, biometry and environmental instrumentation*

The PN extends the capabilities of smart spaces, allowing them to include participants outside the conventional range of interaction. It is the continuous inclusion of the user as a node in the global infrastructure that uniquely distinguishes the PN. These differences are summarized below in Table 1.

property	PDA	PN
on time	mins/mo	100% on high-duty days/mo
usage	primary when disconnected	always secondary (catalyst)
num conns	1	2-3 to facilitate handoff
output	display min. speaker	display stereo speaker
input	keystrokes (primary) buttons, mic	mic (primary) buttons
other I/O	as peripherals	integrated GPS, temp, light accel...

Table 1: Comparing PDAs and PNs

3: Need for a PN

As has already been discussed, a PN helps integrate people with the network infrastructure, enabling new uses for smart spaces. This section elaborates on how PNs help complete the network component design space, and examines the variety of capabilities it enables.

3.1: PN as a "Smart Spaces Device"

Since the earliest days of the ARPANet, networking has been based on the notion of two kinds of components - hosts and routers [4] [6]. This model has been modestly extended via Mobile IP and DHCP to support hosts that relocate periodically, such as laptops. IP telephony also seeks to introduce a new kind of node into the Internet, i.e. the IP telephone.

New types of devices are extending this model even further. The notion of a host as an aggregation of host- and router-like devices is the basis of desktop area networks (DAN [1], Viewstation [14]). In this case, the previous model of host as a single, terminal point on a network is insufficient. The network of workstations has become another example of this kind of virtualization and aggregation.

Even so, the most novel recent network devices are expressible in this model. Wearable computers are typically modeled as laptops, where they are not worn continuously, and move from network to network only sporadically. They are designed more as a sophisticated PC, one that replaces the conventional desktop workstation, but is intended as the same kind of stable network device.

By analogy to the telephone system, the PN can be shown to fill a gap in the design space of network devices (Figure 3). There have been earlier attempts to address this 'non-host, non-router' role [27]. In the early days of telephony, devices were scarce and expensive, so party lines (e.g., mainframes with batch sharing) were the norm. As telephones became less scarce, direct dial and per-home installations became the norm (PCs).

Users began to rely on access to telephones, due to their pervasiveness in permanent installations. They extended this demand to mobile use, initially by "move and reconfigure", e.g., mobile telephones requiring operator assist and pre-configuration.

Up to this point, message services and answering machines were necessary, because the desired party wasn't always available when called. Telephone use was significantly asynchronous, except during business hours.

The telephone system eventually evolved to support true mobile phones and pagers. Both pagers and mobile phones are both ubiquitous and present with the desired person, with a single, persistent number. Prior to this, a phone number was a business or home, or car or boat. Now phone numbers are, in effect, "people".

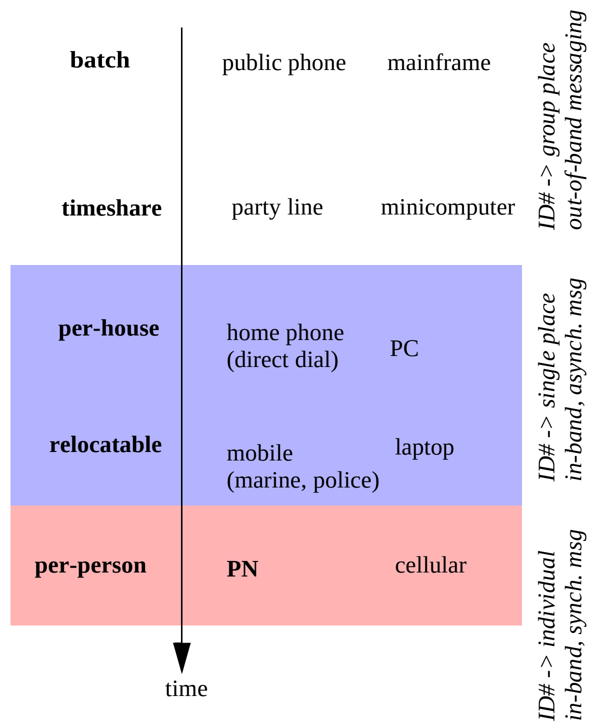

FIG. 3. Computer evolution emulates telephony

Network hosts have evolved from 'office' addresses to laptop addresses. The next step in this progression is the PN, which allows people to become nodes.

3.2: PN Application Domains

There is a phase change in behavior when accessibility is continuous, as has been seen in the use of pages or cellphones. The PN makes networking access continuous, makes smart spaces aware of the user and provides feedback on current user state. By design the PN is a catalyst enabling personal interaction with smart spaces.

We present a variety of examples of how a PN catalyzes personal interaction with smart spaces. These include:

- *Presence Sensitive Applications*
- *Smart Emergency Spaces and Services*
- *Autonomous Information Gathering*

Presence Sensitive Applications

The earlier field operative example demonstrated how sensitivity to user presence can be used by a smart space. There will be many such applications.

- *Security system smart doors could be automatically unlocked for authorized individuals.*
- *A smart docent notices a PN in its space, asks the PN its user's language preference and level of subject knowledge, then begins delivery of personally tailored audio program.*
- *Soldiers' PNs are interrogated for equipment and supply levels. Such information would be aggregated and fed back to logistics support.*
- *A smart transportation space would collect user destination information, e.g., to skip subway stops where no users are waiting or want to depart*
- *Upon entering a smart business space PNs could be supplied with the businesses web pointers. Current prices, not listed on the public web site, could be provided explicitly to local PNs.*
- *In training scenarios the status of individuals entering a smart space could be ascertained and their behaviors captured.*

Smart Emergency Spaces and Services

A PN is a wireless point-of-contact for the individual carrying it that contains instrument and computational assets. This combination enhances emergency services.

- *A PN can be loaded with site-specific emergency services information and software.*
- *A PN can verify the authenticity of emergency-services messages.*
- *When local communication services fail, a PN can draw upon its other wireless interfaces for support, creating an ad-hoc base-less network*
- *Because a PN knows its location, it can be addressed by location and proximity, locating an individual or limiting alarms to affected PNs*

Autonomous Information Gathering

As a user roams, their PN will pass through smart spaces. Information that may prove useful could be regionally broadcast to passing PNs. The reverse operation, in which the PN broadcasts information into a smart space, is also useful.

- *PNs located in vehicles or carried by soldiers could respond to queries regarding current level of supplies. This data could be aggregated and fed upwards to logistics personnel for disposition. If a PN indicates a shortage of critical supplies, the driver or soldier could be directed to the nearest depot location that stocked those supplies while also debiting depot inventory level.*
- *The ability to locate an employee and their vehicle allows delivery, repair and pickup services to operate more efficiently. This also allows resources to be staged as people move, extending the smart space to a smart logistics capability.*
- *In a commercial setting, PNs would anchor a broad new set of autonomous customer services. For example: A PN can automatically identify its user, allowing the business to use profiling information to direct the user to particular specials.*

4: Research Issues

PNs provide fundamental and new capabilities that have not previously existed. Their use raises a number of research issues concerning protocols, naming and addressing, coordination and configuration, privacy, scale and user interfaces. Their existence makes possible entirely new application areas that involve roaming, information capture and ad-hoc networking.

4.1: Protocols

The PN is predominantly a communications device, so the main research issues relate to protocols.

Integration with existing protocols

Many of the current Internet protocols are challenged by the way in which PNs change the model of a host. The PN requires continuous communication as it spans different link technologies, and may require periodic hibernation to conserve power. State-oriented network protocols, notably TCP, do not react well to such idle periods, and are not intended to support endpoint renaming. TCP also requires several round-trip times per exchange, increasing the latency of simple request/response protocols. Newer variants, including T/TCP, reduce this effect, but require further modification to support seamless transitioning [5].

The PN may also need to support proxy operation of other protocols. These might include delegated request/response, where the request is issued to a smart space resource which coalesces responses on behalf of the PN. Other remote protocols may include switchboarding, or remote control of a variable delegation point, such as to coordinate or redirect multimedia streams among other smart space devices.

Device-Control Protocols

Smart devices and PNs are network attached peripherals (NAPs). NAPs are relatively uncommon today and most utilize media-specific transport protocols. To avoid recreating the tower of babel problem that now exists with NAPs, smart devices should support Internet rather than media-specific transport protocols [13]. This requires the creation, testing and standardizing of physical device-control protocols and their APIs.

Because peripheral devices themselves fall into classes of similarity, such as disks and displays, it is reasonable to propose standardizing large portions of the lowest-layer interface that a class of smart device presents.

Internet accessibility increases risks to device privacy and integrity. Both privacy and integrity concerns should be met by adopting a methodology similar to that of Derived Virtual Devices (DVDs) [28].

Link agility

Roaming will move a PN into and out of contact with wireless networks. To maintain network presence the PN must monitor the 'liveness' of its links and associate with new ones as needed. Liveness monitoring and link association algorithms need to be developed and investigated. Low-power consumption is a requirement that must guide this work.

There are a number of triggering events to consider:

- *Failure of 802.11 low-power beaconing*
- *Hearing mobility agent from another subnet*
- *Distance from wireless hub and heading*

Tracking the rate of damaged messages combined with characterization of expected error rate versus distance [22] could also trigger association with a new network. Geographic information and application requirements may be of use when there is a choice of foreign agents and networks with which to associate.

Proxy protocols

The PN is not necessarily the best device to run request/response or stream protocols. Its limited bandwidth and power hinder it from first-class participation as a router in the smart space. Instead, it may be appropriate to off-load some protocols to proxies at other smart space devices.

These proxies could scatter requests and gather responses, effecting nested transactions. They could coordinate ordering protocols, or redirect continuous stream multimedia traffic, switchboard-style.

Protocol trade-offs

There are a number of bandwidth vs. latency vs. power trade-offs in a PN, some already presented. In general, wireless sending costs 10x the power of receiving, so the PN is a good platform for asymmetric protocols. Periodic retransmission of popular data, or server anticipation of PN requests based on traces, both can greatly reduce the power requirements and latency of access.

Furthermore, the bandwidth in and out of the PN is highly variable, and highly asymmetric as well. It may be feasible to enable the receiver of high-speed IR, but only the transmitter of the medium-speed office network. These trade-offs also affect the use of CPU and memory resources, because both can increase latency and power utilization.

4.2: Naming and addressing

The PN also challenges the traditional network notions of naming and addressing. The name of a PN should be intrinsically linked to that of the owner, not particularly to the device itself. GSM cell phones have this property, where the GSM encoding card identifies the telephone number, independent of phone. However, there are other challenges, some of which are only beginning to be addressed in the research community.

Geographic Addressing and Broadcasting

Providing each PN with knowledge of its location allows it to be addressed via two modalities, by-interface and by-location. Geographic broadcasting is a particularly attractive new capability. Various approaches to realizing geographic broadcasting need to be examined as do protections against its unauthorized use.

Geographic routing and addressing in the Internet has been approached by creating a virtual network from geographically aware routers located within the Internet in a manner analogous to the initial implementation of multicast routing [16][21]. Earlier work pointed out that geographic addresses could be used to route packets in a hierarchic network and could support host mobility [11].

Geographic addressing could be grafted into IP as an option. However, IPv6 reserves 1/8th of its 128-bit address space for geographic addressing [12]. The possibility of providing hosts both interface-oriented and location-oriented addresses needs to be investigated, as do questions of the affect of geographic addressing on routing and transport protocols.

Smart Space Discovery

As a PN enters a smart space it should become aware of that smart space and vice versa. Mechanisms are also needed to discover what smart devices are in a smart space, which ones the PN can access and which application-layer interfaces they support.

Generalizing Multicasting

A geographic address naturally lends itself to description of a region. For example, a set of geographic addresses defines the interior of a polygon. A radius around a hub's location also defines a region. Such regions can be used to create routing tables [11] or to define multicasting groups [16][21].

Under the current Internet multicast routing scheme, a portion of the IPv4 and IPv6 address spaces are reserved for multicast addresses. Multicast groups are explicitly created and associated with a multicast address. A host explicitly joins a group to become its member and routers alter their routing tables to service group members. Once joined, a host remains a member until it explicitly leaves the group or the group is destroyed. Host mobility has no effect on membership in this type of group.

A multicast group could also be defined geographically. Examples are those hosts within a geographic region, a room, a building, and so on. Ignoring for now precision of location, a host is either inside a group's region or outside it. Under this definition, membership in a geographical multicast group is implicitly determined. Host mobility does affect membership in this type of group.

The are great differences in how these two classes of multicast group are defined and in how membership is defined. Mobility makes ephemeral the membership criteria in geographically defined multicast groups. The set of such multicast groups that a mobile host is potentially a part of could change frequently. Existing Internet multicasting is ill suited to geographically defined multicast groups.

4.3: Coordination and configuration

The determination of resources in a smart space that a PN has entered and the run-time matching of those to application requirements is a producer/consumer problem. Solving this producer/consumer problem effectively is a major research task.

The question of what the appropriate OS will be for a PN also arises.

Configuring Smart Devices

In conventional system architectures, both device controllers and devices are resident in the chassis. The addition or removal of a system device is a relatively rare occurrence that can occasion rebooting of the system.

By way of contrast, in a smart-space environment the set of devices that a PN may come into contact with as the user roams will be large and dynamic. It will be unrealistic to expect a PN to be preconfigured with drivers specific to the smart devices that it encounters. This immediately imposes requirements upon PNs. These requirements are summarized as:

- *Need for Dynamic Reconfigurability*
- *Interface Matching*
- *Push/Pull Configuration Software*

4.4: Scale

Both PNs and smart spaces increase the scale of networks in a variety of ways.

The dynamic range of bandwidth increases, mostly due to a lower bandwidth required for WAN signals to the PN.

The latency range increases, also because WAN signals are liable to use satellite paths. These paths are already used in networks, but PNs would make them prevalent.

The number of devices in the network increases because addresses are required for smart space components, and they are required for PNs. The smart spaces might support address aggregation, but the global roaming capability of PNs may inhibit such a simplification.

4.5: Security

There are two levels of security required for PNs. The data content itself must be secure (authenticated or encrypted), and the event of communication may also require privacy (source confidentiality). The latter would otherwise permit tracking, and behavior patterns themselves may constitute a compromise.

The need for security is especially important for the PN because it is so directly associated with a single individual, and because it contains so much local state (GPS, microphone, etc.).

There is an external aspect of security, that of device theft, which is also important. The small size required for vade mecum operation encourages theft. Cell phone designers have adopted an automatic lock, which must be re-keyed with the user's PIN every time the phone is powered on. The PN might additionally require the PIN be re-entered on a schedule, every day or so. Much of the identity of the user, and much of her state, may also be encoded on a secure card, as in PCS cell phones.

Other devices may be used to simplify or augment the security provided by a PIN. These include biometrics (voiceprint, signature, fingerprint, other personal telemetry). Single-chip fingerprint imager chips are now emerging as commercial products along with the software to perform recognition.[2] Logging in would then consist merely of picking a PN up, pushing a button and swiping a

2. Examples are the Thomson-CSF FingerChip, Harris Semiconductor Corp. FingerLoc and the Veridicom FPS100.

finger across a scan pad. If a PN is used by a pool of people, the minutiae from prints of multiple individuals could be stored for recognition.

4.6: User interface

Size limitations prevent a PN from having a traditional keyboard or display. It is also unrealistic to expect individuals to use a conventional display and keyboard while roaming. Consequently, the user-interface for a PN must be nearly hands-free, depending primarily upon voice control with limited use of buttons.

Developing an effective nearly hands-free user interface for roaming will be a major research area. The need for speech recognition does imply that a PN needs access to significant computing resources.

When not roaming the user could pick up the PN and use it as a 3-D mouse, insofar as a PN contains accelerometers. Use could also be made of its small display. However, it is our contention that when users are not roaming they will likely be able to use nearby conventional display/keyboards.

Designing a user interface for a small display was part of the ParcTab effort [29]. We expect the ideas developed there would prove useful, but we view the small PN display as ancillary.

5: Feasibility

The PN is currently feasible, even given our demanding combination of capability, portability, and continuous operation. It is possible to bootstrap its development with an off-the-shelf, rapid prototype, in parallel with the coordinated application of well-known low-power, integrated packaging design.

5.1: Rapid prototyping

The principal components needed to prototype a PN are readily available. Setting aside the size and packaging issues, much of the needed research and development of the system architecture, protocols, and user interface could be done using a prototype built from off-the-shelf components. Existing PDAs and handheld PCs, together with wireless PCMCIA network interfaces can be used to emulate a PN. The PN must be more conservative in capability, to provide continuous battery-powered operation, however.

Once that prototyping effort is finished, packaging, size and power consumption issues could be addressed separately. The industry has already developed suitable low-power processors, memories and interface circuits. Examples are the SA-1100/133MHz microprocessor from Intel/Digital [9] that consumes 550 mwatts and non-volatile, zero-power SRAM.[3]

5.2: Power conservation

PDAs achieve their month-long recharge intervals by remaining normally off. They await an explicit activation by the user, perform a small amount of resulting computation, display the result and then turn off again.

On the other hand, PNs must maintain persistent network presence and so cannot be turned off. PNs must achieve a minimum recharge interval of 24 hours.

Variations on paging and polling techniques let a PN approach the normally-off power consumption characteristic of a PDA without sacrificing its network presence at the cost of increased latency. Each technique requires the cooperation of wireless hubs.

- *Incorporate a low-power pager in a PN. Hubs page a PN when they have traffic*
- *PNs poll their hub. This has scaling drawbacks but is it simple.*
- *Have the hub prompt polling by specific PNs. This technique is adopted in 802.11 low-power mode.*

Allowing a PN when roaming to remain inactive most of the time, will extend battery lifetime to a reasonable recharge interval. Overall, our current estimate of a PN is 3 watts of continuous operation power, or 8 oz. of zinc-air batteries/day. Paging-based wake-up allows discontinuous operation and will reduce demand, requiring only a small, matchbox-sized lithium or NiMH rechargeable battery.[4]

5.3: Communication

The PN requires the integration of a variety of link technologies, from high-speed desktop to low-speed wide-area links. Current variants of these include FIR (fast InfraRed) for the desktop, 802.11 wireless ethernet for the office, Metricom for the city, and text paging for the wide-area. Most metropolitan and wide-area communication technologies include their own power support, capable of 24 hours of standby operation and 1 hour of continuous operation. In the case of pagers, batteries last weeks, and support continuous monitoring of a very sporadic data stream.

3. Dallas Semiconductor DS1250 Power Cap SRAMs incorporate their own 10-year lithium cells.

4. Alkaline - MnO_2 batteries have an energy density of 75 WH/lb. Lithium - Li/MnO_2 batteries have 105 WH/lb and Zinc-Air 140 WH/lb. Source: http://www.duracellnpt.com.

These different link technologies have widely varying bandwidth, latency, and reliability capabilities. The application protocols need to adjust to the available link capability, operating in loosely- or tightly-coupled modes as needed.

5.4: Packaging

Packaging issues include power conservation, thermal diffusion, ruggedness, and integration for miniaturization. Current sensor technology already supports component-level and chip-level versions of many of the devices proposed for the PN. Power consumption and heat buildup can be reduced by conventional techniques (power devices only when in use, or only periodically).

The overall package needs to be wallet-sized (or smaller), and 1 lb or lighter. There appears to be a common *vade mecum* size, that of a cell phone or PDA, that is acceptable.

6: Related work

The PN is a variant of PDA technology and wireless 'presence' devices. It also extends the networking efforts of recent wireless and mobile protocols.

6.1: Handheld PDAs

Other more recent PDAs and handheld PCs have integrated small displays, touchscreens, and sometimes keyboards to provide rapid user access to limited local resources. Examples include the 3Com Pilot and a variety of WindowsCE palmtop PCs. The Philips Nino[5] incorporates a microphone to support speech-based commands, in addition handwriting or script recognition provided by each of these devices. Compared to the PN, these palmtop examples lack environment sensors, and have only limited wireless capability (usually only IR on-board). More importantly, they are designed to be used only intermittently, in an "on, check/enter data, off" cycle lasting minutes. Their battery life is often measured in hours of runtime; they achieve weeks of life by being 'off' most of the time.

6.2: Wireless 'presence' devices

The ParcTab [29] was an early example of a wireless personal node. It had a single infra-red interface and provided persistent presence for in-building roaming. The PN extends the range of the ParcTab, supporting wireless access beyond the building confines. The ParcTab provided PDA-like services directly, whereas the PN focuses on catalyzing of other devices on its behalf.

The Lovegety is a simple wireless personal node that demonstrates mobile information capture [17]. It uses peer-to-peer beaconing to allow singles to meet when in one another's proximity, by exchanging a few bits worth of preference information. The device is small enough to fit on a keychain, and runs for days to weeks on battery power. This simple device can be considered a low-bandwidth variant of a PN, enabling singles to detect each others' presence with a smart space created by a multi-party ad-hoc baseless network.

6.3: Wireless and mobile protocols

The problem of providing persistent internetwork presence while roaming is the subject of the Mobile IP development effort [24][18]. A mobile computing environment that uses multiple types of wireless networks is called a *wireless overlay*. Maintaining connectivity while running applications in this environment is extremely challenging [19][26] and is the subject of a number of research efforts, including BARWAN [2] and Odyssey [23]. Operating systems originally designed for workstations require extensive changes in a nomadic environment [3].

Imielisnski and Navas proposed embedding of geographic routing and addressing in the Internet by creating a virtual network from geographically aware routers [16][21]. Geographic addressing can also be directly used to route packets, support host mobility and provide regional broadcast [11].

7: Implications

The concept of a PN extends and challenges smart spaces and general issues in network research. By including users as nodes in the network, it extends the scope of the network, and the capability of applications within it. It also provides an opportunity to apply technology being developed for low-power, integrated sensors in a unique way to provide ad-hoc mobile smart sensor nets.

7.1: Smart spaces

The PN avoids the distinction between a user's on-line and off-line presence. The user is always on-line, accessible to signal for feedback, supporting immediate urgent-mode interaction. Because of its integrated, multi-level communications links it provides an opportunity to catalyze the aggregation of other network resources for the user's benefit.

5. http://www.nino.philips.com

7.2: Network research

Traditional networking considers users as temporary presences at permanent end-points known as hosts. The PN extends this notion, where people themselves become end-points on the network. People are more mobile, even than laptops, and so require hand-off without dead time, and a truly persistent identifier. Location of a user becomes a key network resource discovery issue.

The PN provides an opportunity to review more conventional host and gateway requirements, using a model that challenges their assumptions. Overall network architecture, naming, addressing, and resource discovery all may require re-examination.

In addition, transport protocols may require additional support for continuous relabeling of the endpoints, as users move between smart spaces. The PN itself may provide bridging capabilities between adjacent PNs when necessary. Finally, the traditional request/response protocols may require redesign, to support a proxy-mode operation, to off-load capabilities to smart space resources and conserve local power.

7.3: Application of related technology

There are a number of related technologies that are required for a PN to be developed. Small, low-power sensors already exist, but need to be integrated with a small amount of processing and storage into a handheld device. The PN focuses on placing as much I/O technology where the user is as possible, so there is virtually no limit to the challenge to integration technology here.

By placing the sensors where the user is, the PN provides a unique opportunity for ad-hoc deployment of sensor networks, in effect a mobile smart space centered, and concentrated exactly where the users are. Deployment at the appropriate place is de-facto achieved.

There is also a challenge to integrate a number of wireless communication technologies into a single, low-power, configurable device. This includes bandwidths from 1-1Mbps, latencies from µs to 100's of ms, and ranges from feet to tens of miles, using IR, CDMA, GSM, and even simple analog paging technologies.

8: Summary

Wireless technology will soon be used to create and leverage smart spaces comprised of peripheral devices and sensors that communicate with one another and the network. As humans we can't directly perceive the wireless spectrum and so we aren't aware when we are in a smart space and we won't know what's in it. Conversely, smart spaces won't be aware of our presence. As long as people do not have the capability of directly interacting with smart spaces as they roam, smart spaces will remain restricted in their scope of application and ease of use.

These issues are addressed by creating small personal wireless nodes (PNs) that are carried with individuals as they roam. The PN's goal is to integrate the human being into the Internet.

The PN allows the user to become a persistent part of the network, by providing:

- - *continuous communication with the user via minimal I/O*
- - *a variety of user-centric telemetry and biometry sensors*

By providing a minimal initial access, these capabilities can be used to bootstrap the user's access to more advanced services, and to support ad-hoc base-less networking when disconnected from the rest of the net.

Acknowledgments

The authors would like to thank Bill Manning, Gene Tsudik and Jon Postel for their helpful suggestions.

9: References

[1] Barham, P., Hayter, M., McAuley, D., Pratt, I. Devices on the Desk Area Network IEEE Journal on Selected Areas in Communications, May 1995.

[2] BARWAN Project http://www.cs.Berkeley.edu/~randy/Daedalus/BARWAN/BARWAN_over.html

[3] Bender, M., Davidson, A., Dong, C. et al. Unix for Nomads: Making Unix Support Mobile Computing First USENIX Symposium on Location Dependent Computing, 1993.

[4] Braden, R (ed.) RFC 1122: Requirements for Internet Hosts - Communication Layers Network Working Group, Oct. 1989.

[5] Braden, R., RFC 1644: T/TCP -- TCP Extensions for Transactions, Functional Specification USC/Information Sciences Institute, Jul. 1994.

[6] Braden, R., Postel, J., RFC 1009: Requirements for Internet Gateways ISI, Jun. 1987.

[7] Davis, C., Vixie, P., Goodwin, T., Dickinson, I. RFC1876: A Means for Expressing location Information in the Domain Name System University of Warwick, Jan. 1996.

[8] Deering, S., Hinden, R.
RFC1883: Internet Protocol, Version 6 (IPv6) Specification
Xerox PARC, Ipsilon Networks, Dec. 1995.

[9] DIGITAL Semiconductor SA-1100 Data Sheet
EC-R8XUA-TE
Digital Equipment Corp., Maynard, MA., Jan. 1998.

[10] Farrell, C., Schulze, M., Pleitner, S., Baldoni, D.
RFC1712: DNS Encoding of Geographical Location
Curtin University of Technology, Nov. 1994.

[11] Finn, G. G.
Routing and Addressing Problems in Large Metropolitan-scale Internetworks
ISI Research Report ISI/RR-87-180.
USC/Information Sciences Institute, Mar. 1987.

[12] Hinden, R., Deering, S., editors
RFC1884: IP Version 6 Addressing Architecture
Xerox PARC, Dec. 1995.

[13] Hotz, S., Van Meter, R., Finn, G. G.
Internet Protocols for Network-Attached Peripherals
Proc. Sixth NASA Goddard Conference on Mass Storage Systems and Technologies and 15th IEEE Symposium on Mass Storage Systems, Mar.1998.

[14] Houh, H., Adam, J., Ismert, M., Lindblad, C., Tennenhouse, D.
The VuNet Desk Area Network: Architecture, Implementation, and Experience
IEEE Journal on Selected Areas in Communications, May 1995.

[15] Imielinski, T., Navas, J. C.
GPS-Based Addressing and Routing
Technical Report (LCSR-TR-262)
Rutgers University Computer Science,
Mar. 1996 (updated Aug. 9, 1996).

[16] Imielinski, T., Navas, J.
RFC2009: GPS-Based Addressing and Routing
Computer Science, Rutgers University, Mar. 1996.

[17] Lovegety
http://www.freeyellow.com/members2/onmarketing

[18] Johnson, D. B.
Mobility Support in IPv6
Carnegie Mellon University,
Sun Microsystems, (work in progress).

[19] Katz, R., Brewer, E.
The Case for Wireless Overlay Networks
Ch. 23 in *Mobile Computing*,
Kluwer Academic Publishers, 1996.

[20] Mockapetris, P.
RFC1035: Domain Names - Implementation and Specification
USC/InformationSciences Institute, Nov. 1987.

[21] Navas, J. C., Imielinski, T.
GeoCast: Geographic Addressing and Routing
Proc. of the Third ACM/IEEE International Conference on Mobile Computing and Networking (MobiCom'97), Budapest, Hungary, Sep. 26-30 1997.

[22] Nguyen, G., Katz, R., Noble, B., Satyanarayanan, M.
A Trace-Based Approach for Modelling Wireless Channel Behavior
Proc. of the Winter Simulation Conference, 1996.

[23] Odyssey Project
http://www.cs.cmu.edu/afs/cs/project/coda/Web/docs-ody.html

[24] Perkins, C., Editor
RFC 2002: IP Mobility Support
IBM, Oct. 1996.

[25] Rekhter, Y., Li, T., editors
RFC1887: An Architecture for IPv6 Unicast Address Allocation
Cisco Systems, Dec. 1995.

[26] Satyanarayanan, M.
Fundamental Challenges in Mobile Computing
Fifteenth ACM Symposium on Principles of Distributed Computing, May 1996, Philadelphia, PA.

[27] Stewart, B., (chair)
"Charter of the Special Host Requirements (shr) working group",
Proc. of the 21st IETF, Jul. 1991.

[28] Van Meter, R., Hotz, S. Finn, G. G.
Derived Virtual Devices: A Secure Distributed File System Mechanism
Proc. Fifth NASA Goddard Conference on Mass Storage Systems and Technologies, Sep. 1996.

[29] Want, R., Schilit, B.N., Adams, N.I., Gold, R., Petersen, K., Goldberg, D., Ellis, J.R., Weiser, M.
The ParcTab Ubiquitous Computing Experiment
Ch. 2 in *Mobile Computing*
Kluwer Academic Publishers, 1996.

[30] Wearable Computing FAQ, v1.0, 8/28/97.
http://lcs.www.media.mit.edu/projects/wearables/FAQ/FAQ.txt.