

Paper Number 1110
Nomadcity: Connection for the Terminally Transient

Frank Prince

Hewlett-Packard
300 Apollo Drive
Chelmsford MA 01824
(508) 436-4945

WHAT IS NOMADICITY AND WHY SHOULD WE CARE?

Being a nomad is just what we might expect — moving around and sometimes getting stuck in an inconvenient place. Nomadicity is the electronic counterpart to being a nomad.

It used to be, before we had to deal with nomadicity, that "getting stuck in an inconvenient place" meant the place we were wasn't well suited to the things we wanted to do. If we went on a business trip we had to pack up all the materials for our presentation in a briefcase, toss in our calendar, a supply of pens and pencils and anything else we might need. It was all because when we got to the customer we wouldn't have our office to support us.

Now we have nomadicity. Things that are important to us are on computers. Our businesses use electronic communication: telephones, FAX, private cable systems. That means those technologies are part and parcel of our everyday work. What happens now when we go on a trip? We go through an electronic version of what we used to do. We work out a way to get access to the parts of our electronic environment we need wherever we're going. Most of the time, doing that is not particularly easy.

WHAT'S BEING DONE NOW: MOBILE EQUIPMENT AS A RESPONSE TO NOMADIC CONCERNS

There are two ways we might arrange access to our electronic "stuff" when we are on the move. We might have computing, application and communications support at the remote location. Doing that, we could get at our data or applications either through a network or through removable media. Any equipment that met our requirements would be good enough. In a way, we would be equipment independent. The second way we could solve our problem is take equipment with us. We could either pre-load the our data and applications or establish a communications path back to our home base. For this second option we would need mobile equipment.

For a lot of historical reasons, the mobile solution to nomadic computing is more common than the equipment independent one. That is not true for telephony. Functionally interchangeable public phones have long been available. Telephony has been largely equipment independent.

A Mobile Response Generates Its Own Problems

If we hadn't chosen mobility to solve nomadic problems there would still be mobile equipment. But it wouldn't have to be engineered to respond to nomadic needs. Once we make a decision to use mobile equipment to solve nomadic problems, we take on a whole set of second-order considerations. Each secondary consideration has to take nomadicity into account.

Task Complexity

The reason you pack up something and take it with you is that you feel you can't leave it (or what it does) behind. You are pretty sure you're going to need it. The relative new-ness of electronic functionality complicates decisions about what we need to take with us. Considerable research (not to mention trial and error) has gone into deciding what functions need to be built into portable computing and communications equipment. It has taken years to define the functionality set of personal information managers (PIMs). Now it is relatively well accepted. A general-purpose portable computer has very different requirements from a basic hand calculator even though each is portable computing equipment.

The starting point for any design is the user requirements it has to meet. For mobile equipment in service of nomadicity those tasks have to be drawn from nomadic requirements. If that means the computing and communications tasks that the user is unwilling to do without, packaging decisions for nomadic equipment must be user-centric.

Mobile Design Constraints

Once you know what you want the equipment to do you can decide how portable you want it to be. The degree of portability interacts directly with form factor requirements. Computer and communications equipment could be shipped in containers or mounted in vehicles and still be counted as mobile. Whatever design center is chosen, it will affect other form factor concerns. We tend now-a-days to think portable equipment will have to be carried by an individual human being — and that is the focus we will maintain in our discussion.

User Interface

The traditional complaint of the portable computer user is that input devices are too small. The keyboards and pointing devices are the most often cited examples. While the size of the interface component is certainly a factor, dealing with it is complicated by users functional expectations. Users who have become comfortable with a particular keyboard layout have a hard time moving back and forth between their more familiar environment and the one on their portable device.

Output devices — mostly displays, but to an increasing extent hard copy and audio — are also put under pressure in mobile environments. Physical size restrictions in portable devices are again a complicating factor. If it is big enough to see it may be too big to carry around. In a similar vein, if it is loud enough to hear it may be too loud to use in a public place. If it has the ability to plug in earphones it may be too awkward to use in some of the places people need it most.

The complexity of the interaction with the physical user interface is driven by the task to which the equipment is put. A "general purpose" computer will have a different set of requirements than a cellular telephone or a garage-door opener. The tendency of designers has been to tailor equipment user interfaces to reflect the operation of their non-mobile counterparts. The computer has a screen, keyboard and pointing device. The telephone has a keyboard. The garage-door opener has a button. These decisions all diminish the learning requirements placed on the user — but at the cost of requiring that the designer cram the interface into a portable form factor. A dearly learned lesson of interface design is nowhere more appropriate than in mobile equipment: users' previous experience is ignored only at the peril of product success.

Weight and Size

There are three "niches" in the size of portable equipment: pocket, bag, and mounted. By that we mean it will fit in a jacket pocket, a purse or briefcase, or it can be mounted on a vehicle. More recently a fourth slot has been added — "worn" — meaning it can be worn as one would a piece of jewelry or an article of clothing. Pioneering in this new category are wristwatch, pager and portable telephone manufacturers. In general we can appeal to personal experience about the weight of items in these categories. If it is worn it should only weigh an ounce or two. If it is in a pocket, even a pound is heavy. If it is in a bag seven pounds is too much. Ideal weights are much less than these maximums. In general — users want to carry around as little as possible.

The binding of function to packaging and the corresponding impact on the ability of designers to implement the combined requirements has led to secondary requirements for small, light, rugged components. This is an example of the "fan out" of nomadic requirements into apparently non-nomadic environments. In electronics there are good reasons to make components small. Small components can be faster or use less power or be more complex for a given size. These goals, however, are sometimes mutually incompatible. Knowing the component will be used in a portable device helps drive design tradeoffs.

Power

Power Budgets and Engineering Constraints

Any design requiring power has to take into account its operating environment and what it will be used for. Once those decisions are made a power budget can be established. All the components of the design can not, together, use more power than the total available from the power source.

Interaction with Utilization Requirements

Where, when, how long, and for what the equipment will be used all impact the power budget. Intense, continuous use for long periods of time where there is no external power source is the most difficult problem. It will usually be reflected in limited functionality. As restrictions on utilization are eased greater functionality can be built in — but sometimes at the expense of convenience. When an external power supply is available the user has to deal with carrying around and using a power converter and cord. The user goes from untethered to tethered.

Interaction with Weight and Size

The weight and size of equipment impact power budget. If equipment has to be small and light then engineers will have similar limits on the weight of the power supply. In power supplies, less weight usually means less power. Even when the application allows a power cable the components that provide conversion from "wall plug power" to a form the equipment can use take up space and weight.

Over all — small, light, powerful power supplies are an important goal of nomadic designs.

Communications Capabilities

There is some mobile equipment that stands completely alone. A personal information manager (PIM) might never get connected to another piece of equipment. When that is so, to lose the PIM is to lose all the information stored on it. The same is true if it breaks. In general — if mobile equipment is used for communication (a cell phone) or interaction (a bank card, a computer), or output (a portable television) or if it just needs to have its information backed up for safe keeping — it needs to have some means of communication with the outside world.

Interaction with Tethering

We have two broad classes of communications techniques based on the physical means of connection between the equipment and the rest of the world: tethered (wired) and untethered (wireless). Within each of these categories there is a range of different technologies that meet communications needs in different situations.

On the wired side, Ethernet in a 10BaseT configuration has become the most common low speed Local Area Network (LAN) at 10 million bits per second (Mb/s). When the need for more speed arises we have 100BaseT at 100 Mb/s, Asynchronous Transfer Mode (ATM) at up to 622 Mb/s, and various fiberoptic connections into the billion bit per second (Gb/s) range are available. Coaxial cable and combined fiberoptic/coaxial cable systems are providing connection over longer distances at speeds in the 10Mb/s range. Communications over the telephone network cover the range from 14,000 to 10Mb/s. In general, the faster you go, the farther you go the more costly the connection — both to the end user and to the service provider.

In many situations wireless communications are a useful alternative to wired communications. In the case of local area networks, the cost of wiring older buildings might make wireless solutions an attractive alternative. In a similar way, the cost of establishing a wired infrastructure over a particular geography might make a wireless solution more practical. If the choice is between running a hundred miles of line over a mountain or setting up a transmitter and receiver the wireless alternative may well

be a clear winner. Wireless systems may also be quicker to set up. If there are time constraints on deploying a system, wireless communications could be the answer.

Wireless Communications as a Response to Mobility

Notice that mobility of equipment or people has not entered into our discussion about the choice of wired or wireless communications. We wanted to establish that just as mobility is not always in response to nomadic concerns, wireless communications is not always motivated by mobility concerns. To take it a step further, wireless communication might not be driven by nomadic concerns either. But often it is.

One good reason to have mobile equipment is that the people who use it need to move around. One good reason to have wireless communication is that the best answer to mobile communication is often wireless communication. Once again we see that the requirements of the individual nomad drive nomadic requirements.

The Geographic Component to Wireless Communication

There are several predominate constraints for wireless system are design: transmission bandwidth, transmission distance, and cost. How many communicating parties there will be and the form of their communication — point-to-point, group broadcast, or broadcast — must also figure into the design. If there are logical groupings to the people and equipment communicating, the design has to ensure there is no interference between the groups. For mobile implementations all the considerations about power budget and size need to be taken into account.

When a design is finalized the theoretical distance over which it will operate is defined. Local conditions have impact on the actual range. An infrared system might be able to communicate point-to-point between two pieces of equipment over a distance of 10 meters. A cellular phone system might be able to communicate point-to-point over a distance of five miles between a particular phone and the cell tower. Multiple phones can communicate with the tower at the same time. The area covered by wireless system is a "cell".

Cells are, by design, disjoint. This is true for wireless systems like cellular phones — where the cell is the area served by one tower, and for wired systems like Ethernet LAN segments — where the cell is the area to which a broadcast message propagates.

WHAT NEEDS TO BE DONE: SYSTEMS ISSUES

Expectations

We talked about how mobile equipment can be a response to nomadic needs and how wireless communications can be a response to keeping mobile equipment connected. Keeping connected while moving has to deal with movement between disjoint cells.

In the last section we drew a picture of "cells-of-communication" — like circles drawn on a map with a radio tower in the middle. Cellular telephones are an example. As a phone moves it enters one circle (one cell) and leaves some other. If it is going to stay connected while it moves there needs to be some system in place for handing off the phone call from one cell to another — passing the baton like a series of relay runners. That is what the cellular phone system does. There are also protocols that connect the cellular system to the wired telephone system so you can call, and be called, by people at regular telephones.

There is similar set of issues for wired digital communications. If the equipment — perhaps a portable computer — can plug into a 10BaseT Ethernet port in your office you might expect it could plug into a 10BaseT Ethernet port in someone else's office and still work. Just like the cell-based wireless system,

for that to happen there would have to be a system in place for automatically disconnecting communication to one location and hooking it up to another.

Equipment Independence

We should point out that all this moving and connecting and disconnecting has to do with mobile equipment — not necessarily the people who are presumed to be using the equipment. Remember that the difference between a cell phone in your car and a cell phone in your pocket is that the cell phone in your car is just another place you might not be. Mobile equipment is just one response to nomadicity. If everywhere you cared to get a telephone call there was a telephone and some system was available to direct incoming calls to your location you would never have to carry a portable phone. If every time you wished to make a call there a phone were at hand, you would never have to have a phone of your own. A system of this type would free you from carrying around a phone of your own. You would be "equipment independent".

We can cast a similar example in computer terms. If everywhere you wanted to carry out an information processing task there was equipment at hand that would serve your needs, you would not need to carry any computing equipment with you.

Even mobile and wireless equipment sooner or later has to connect to some "equipment independent" infrastructure. For cellular phones that is the system of towers. One tower (given ownership and common technology) is just like every other. For portable computers with modems one phone jack is just like another.

What is important to remember is that no matter if we are equipment-dependent or equipment-independent, no matter if we are computing or communicating — sooner or later multiple interacting systems have to be marshaled on our behalf. No single system will do the job. In the following sections we will explore some of the things that have to get done if such systems are going to be a reality.

Now that I'm Here, What Can I Do?

Being stopped is just a special case of moving. The thing to focus on is transition from one environment to another. The things that make up our electronic environment are the pieces of equipment available to us and the communication paths available to them. We have "moved" if something in that environment changes status relative to the things we want to do. That could mean walking into a room where the available equipment is different than it is in the hallway. It could mean moving from one cellular telephone site coverage area to another. It could even mean a printer going off-line or a network connection going down.

If we want to continue working we have to deal with those changes.

Discovery

The first step in dealing with change is discovering what resources are available to us. That means not only identifying the communications paths and equipment available, but tracking changes in status over time.

A bootstrap process is necessary to get discovery off the ground. We could arrange it so we are never shut down — we keep our portable computer on "stand-by", never turning it off, so it comes "instant-on" when we open it up, or we keep our cell phone turned on and the system tracks us as we move. If we don't do one of those things, when we enter a new location or turn on a piece of equipment we have to have some protocol for how we discover what resources we have available.

Right now, as often as not, that convention is "ask the guy who works there." We bear the burden of reconnecting ourselves. "Can I get an outside line?" "Are you on the Intranet?" "Where's the phone jack?" "Does the network here support DHCP?" "Do you folks have a firewall?" "Does this computer have 123 installed?" "What's the name of that printer?"

Self-Organization

If there is no communication infrastructure in place the situation is worse. How do you bring ten laptops, two printers, a scanner, and a projection display into an empty room and make them all work together? Do you all have the same networking software? Do you have the cable, the connectors, compatible IR ports?

The ubiquitous nature of 10BaseT Ethernet and TCP/IP networking has given us some hope of solving that kind of problem. It is not entirely out of the question to expect that all the equipment that comes to the party will have those connections and protocols. Products are appearing that bridge between IR and Ethernet. While less standard, similar products are available for various other wireless transmission techniques.

And work is being done on self-organizing the services that make a TCP/IP network work. In such systems some machine will step up to handing out IP addresses, another will volunteer to be the name server. The systems will agree on a net mask.

Once we need to get off a single LAN segment, or if we want to dynamically configure multiple protocols, the situation becomes more difficult. If we want to coordinate across technologies — I enter the room, make myself known to the computer system and the telephone system now knows how to route calls to me — the problems are even harder.

Configuration

Getting things connected is not the same as getting them to talk. You can call Moscow, but you might not understand the person who answers the phone. Above the transport protocols are the protocols between devices and applications. What kind of printer is available? What operating system support is available for it? Can I get at the modem on the computer through the network?

Within single operating environments there is some work being done. Microsoft Windows 95 and Windows NT have protocols in place for downloading printer drivers dynamically from the machine hosting the printer. Telephony APIs for network access are being developed. But today's state of affairs is clear: discovery and configuration are the responsibility of the user.

So Where Is My Stuff?

Lets say that fortune is with us and we know what equipment and communications are available. The network is set up. The equipment is configured. We still have not done the first bit of work. There are other issues that need to be addressed.

We have four components to our computing and communications activities: communications, processing, storage and user interface. Any one of them could be any place — though conventionally they are all located on the same piece of equipment. But the nomad is moving around. Even if communications are in place there could be other problems. The communications paths available could be too slow or unreliable. What if he forgot something (the data is somewhere else)? What if his computing capability is not up to the task he has to do (the processing might need to be somewhere else)? What if the interface devices don't do the job he needs done (there is no color printer he can find)?

It is not the only situation, but in general you want to get the results delivered to the place you are.. Where the data is and where the processing takes place is less important. Putting the user interface with the user gives us a starting place for how we distribute computing and storage.

Separating the Interface from the Processing

If the processing is remote from the user there are a number of technologies available for "remoting" the user interface. Programs exist for controlling Microsoft Windows from remote machines. The Unix

X Window interface is inherently remotable. The Java language and Java Virtual Machine provide mechanisms for separating the user interface of an application from its processing. Web browsers and the HTTP provide another alternative.

Where ever the processing takes place ultimately the data has to get to it. How do you choose whether to move the processing to the data or the data to the processing?

If we consider only the simple case of "data" meaning whole files, we can illustrate the kind of calculation that needs to be made. It starts with how much time is available to do the task. There are three issues. How fast can the available processors do the processing? How long would it take to establish the processing environment on a particular processor and connect to the user interface at the user's location? How long would it take to move the data to the processor (and back if there is an update requirement)?

One answer is you can't get there from here. No combination of available processor, communications paths, and storage will meet the time constraints placed on the task. If that is so, either the constraints have to be relaxed or the task abandoned.

Moving Data to the Processor

If an available processor is already set up to do the task, the data can be moved, the processing carried out, the results returned to the user, and the data resynchronized: move the data to the processor.

Moving Processing to the Data

If a processor that can be configured to do the task is available at the site of the data, the time to configure it is not too long, and the results can be reported to the user: move the processing to the data.

Moving them both

It could be that the only way to get the job done is to configure a remote processor, move the data to it, do the job, report the results and move the data back. It might even work — if you could figure out how to do it from that conference room you just walked into in East Dubuque, Illinois.

People really have to solve these kinds of problems. What do you do about getting a five megabyte file to the machine you're going to use for the presentation when you are in the hotel, the file is in the office a thousand miles away, and you have to be on stage the next morning. Do you FTP it from the server over the phone line? Do you hook up via NetBEUI (Network BIOS Enhanced User Interface) or via Net BIOS (Basic Input / Output System) over PPP (Point-to-Point Protocol)? Which is faster? Which is cheaper?

Keeping It All Straight

Lets say that fortune is again with us and we have all the software and all the data in the right places, the connections established and the programs ready to run. We *still* have not done any real work.

To make things worse, our actual working environments are not made up of just one computer program using one file. We use several programs, some of which may use multiple files, or interact with other programs, or contact services over a network.

As an example, let's say we check into a hotel — what do we have to do to set up for business?

Reestablishing the Working Environment

For our example we will take as given that all the equipment and communications setup has gone smoothly. All the software is on our laptop. The files we normally use are either local or available over the network. The dial-up communications package worked. We are connected to the Internet. The job we want to do is integrate some material from an Internet source with spreadsheet data in a slide presentation and EMail the result to a co-worker. What is left to do?

First we have to reestablish our working environment. We have to start our web browser, spreadsheet program and slide preparation program. We have to start our EMail program.

Reestablishing Communications Paths

Now that everything is running, we have to reestablish communications with our data sources. For the spreadsheet and slide preparation programs that means opening data files. For the Internet browser it means navigating to the web site and loading the appropriate page. We have to connect our EMail program to our EMail server and log in.

At this point we are ready to begin work. We integrate the material into the presentation and send it as an attachment to our coworker.

Synchronization

It may be that we have to re-integrate the material we have changed on the road into a store of material that we keep in other places. When we are in the office our EMail program stores saved mail in one location. On the road it is stored on our laptop. We have to integrate EMail stored in two different places. Perhaps your group keeps all presentations on a networked disk drive. When we return to the office we need to copy the work we have done to the shared drive.

If we had been working with transaction oriented data our problems would have been complicated further. If we had been working on another set of activities, we would have had to go through our setup for that activity just as we had for the slide presentation creation activity.

Instead of spending precious time doing work, nomads spend undue amounts getting ready to work. This is less fun than most people want.

WHY DEALING WITH NOMADICITY IS SO HARD

It looks like there is a big opportunity for someone to come along and make this entire "getting ready" easier. In fact, a lot of organizations are working on doing just that. But there are both technological and organizational reasons why this is a very difficult problem.

We suggested earlier that system level nomadic problems are the most intractable. That stems in part from the existence of technological silos, bolstered by legacy industry interests. The telephone industry is not the computer industry. The digital communications industry is not the personal computer industry. The server business is different from the desktop business.

This is not to say there is no hope. If we limit ourselves to one operating environment, one application suite or one communications technology we can, through uniformity, achieve a kind of equipment independence in service of nomadic needs. Similarly, more complex environments, carefully tested, can be certified for operation in a variety of places. In either case, flexibility suffers. This is not the best solution. A central fact of nomadic life is change. Being flexible is important.

There are real technological problems as well. When developers work to deal with nomadic issues they often have no architectural framework into which they can fit their solutions. Similarly, especially if the problem is cross industry, there is no existing body of work from which they can draw for support.

Thus we see the cellular telephone industry solving the problem of moving telephones and user preferences with location registers while Internet Protocol Version 6 deals with it through home agents.

Automating Getting Ready

In a broad sense, the task of all these systems is to automate getting ready so you can be freed to actually do something. There are two parts to that. Some system has to keep track — with your guidance and on your behalf — of what you are doing and what you prefer when you enter a new location. Also, a system has to keep track of what things are available to you at that location. Working together these two systems can conspire to set up and maintain your environment at a particular location. If they can't do the job they can ask you what to do.

If these systems are carefully devised, they can be open-ended. They can offer support to and integration with a variety of implementing technologies. Not having these systems in place holds back the development of solutions to nomads' problems.

Keeping Track of What You're Doing

The system to keep track of what you're doing (relative to computing and communications) has to be an electronic system that spans communications and computing technologies. It has to be independent of any particular piece of hardware. It represents you in electronic space. There have been several names associated with this concept: surrogate, alias, secretary.

There is an important distinction between the services built on the system and the system itself. The services are the kind of activities you might classify as "secretarial". Taking and filtering calls — Voice and Electronic Mail, FAX, etc. Filing and organizing — making sure the things you need to take on the business trip are all loaded on your computer. Understanding and being able to communicate your preferences — knowing that even though the slides are in color its OK to FAX them this time.

The system itself is the framework, supporting the individual, into which these services fit. It has to provide mechanisms that support personal services and personal service development. It has to be built on rules for communication with other electronic components and with its owner. With such rules in place, developers can add function and interact with other personal services without having to build everything from scratch.

Keeping Track of Where You're Doing It

The second nomadic system knows the equipment and communication paths available at a particular physical location and their electronic characteristics. We call this system the site coordinator.

The basic function of the site coordinator is very specific. It knows about two things — equipment physical location and equipment electronic location. The physical location is defined in relative human terms. Equipment is "at hand" or "co-located" (in the same room) or "nearby" (not so far away it can not be easily reached). Electronic location is defined in terms of the media, protocol and address of the equipment. (10BaseT Ethernet using TCP/IP at address 1.2.3.4).

For site coordination to work there has to be a pre-arranged technique for contacting the site coordinator assigned to a given physical location. There are multiple candidates for such a technique. Infrared communications — being inherently location specific — is a good possibility. A laptop entering a new location could by default contact the location's site coordinator — sending a message on an infrared port. The site coordinator could listen on a location specific infrared receiver and communication could be initiated.

Like the secretary, the site coordinator depends on standards for communication, query and information update. Because it could be used to supplement the information reporting capabilities of older equipment (which might not be able to report, for example, that it is in a locked room and unavailable

for use except by persons with the key) the site coordinator database must have rules for open extensibility.

How Would It All work?

Let's construct an example: paying your bills on the airplane. Here is how it would work.

When you are working with your computer at home you use a personal accounting program and a bank electronic bill payment program. The accounting program uses a local file to hold your checkbook information. The bank program connects to the bank computer through a modem. You tell your electronic secretary that these two programs are the "bill paying" activity. It records the facilities the programs need and the files they use and puts an icon on your desktop — "bill paying".

Now you are going on a trip. You tell the secretary "take the bill paying activity in the laptop" (perhaps by dragging the "bill paying" icon to the "laptop" icon). The secretary clones itself on the laptop (if it is not already there) and sends all the parts of the bill paying activity to the laptop (if they are not already there).

On the airplane you decide you want to pay your bills. You open your laptop and turn on the power. The secretary broadcasts a query over the laptop's IR port to any site coordinator. The IR port in the seat back responds and the site coordinator for Row 6, seats A & B responds. The secretary says it needs a link to the telephone system at 14.4 Kb/s. The site coordinator responds with two options — if the laptop supports a Hayes 14.4 compatible modem over the IR port, then the IR link can be used. If a cellular telephone link is supported, a telephone jack is available in the seat back near the in-flight telephone. Knowing that the laptop supports the Hayes modem, the secretary chooses the IR port.

Having made contact with the telephone system and configured the laptop to expect a modem on the IR port the secretary updates the applications' initialization information and starts the applications. All the user did was click on the "bill paying" icon on the display.

But what if it had not worked? There are two possibilities. One, there was no site coordinator for the airplane. The user is certainly no worse off than without the system being in place. At least she gets a message saying that auto-connection is not possible and telling what was missing. Two: there is a site coordinator but the facilities are not available to carry out the activity. In this case the user gets an authoritative answer: no communication channel is available that meets the needs of the bill paying activity.

Who's Working On It?

Work on nomadicity has been going on for several years. There is an annual commercial conference on nomadicity. The Nomadicity Working Group of the Cross Industry Working Team (XIWT) has released a white paper on nomadicity and is drafting a reference architecture which includes the site coordinator and the alias (the name they assign to the secretary). The Internet Industry Standards Panel (IISP) has defined a set of nomadicity needs based on the XIWT work and has hosted a nomadicity round table to motivate standardization.

CONCLUSION

We believe nomadicity, especially the concepts of site coordinator and secretary, are *enabling technologies*. They provide a way to represent individuals and places in electronic terms, which in turn creates a framework to deal with people moving from place to place. That is important. When people move around, the world does not stop to wait for them. Their jobs go on, their private lives go on. Now that electronic environments — in banking, education, government, and personal communications — play so important a role in our lives we need a way to carry on those activities while we are on the move. Our lives are too big, they change too much, to be captured in a little box — be it personal information manager, cell phone, or portable computer.

Site coordinators and secretaries give us a way to cut down on getting ready and focus on doing. That is their central value. It is not that they accomplish it themselves. They provide the framework on which developers can build solutions to nomads' varied problems. With them, tools can be built to get the business that reorganizes up and running quicker. The businessperson on a trip can respond to outrageous fortune — and opportunity — more efficiently. The individual can be connected when he wants to be — and uninterrupted when he needs to be.

Several industry groups are in place working to advance the development of solutions to nomadic problems. If all goes well, those efforts will result in concrete definitions of site coordinators and secretaries. If there is agreement on those definitions we are a step closer to making life a little easier for all of us.