

Toward Self-Organized Mobile Ad Hoc Networks: The Terminodes Project

Jean-Pierre Hubaux, Thomas Gross, Jean-Yves Le Boudec, and Martin Vetterli
Swiss Federal Institute of Technology, Lausanne and Zürich

ABSTRACT

This article provides a technical overview of mobile ad hoc networks and describes their long-term potential. It covers current research, and describes major technical challenges, including networking, real-time services, and software. It shows that by their very nature, mobile ad hoc networks can bring a paradigm shift in the way networks are organized and operated, and can even lead to a fundamental change in the relationships between information technology and societal organization. As an illustration of these concepts, the article also contains an overall description of our long-term research project, called *terminodes*.

INTRODUCTION

With the success of IP networks, the involvement and level of responsibility of end users have dramatically increased; users are now entitled to develop and modify the applications they run on their end systems. This freedom has fueled creativity and has led to the deployment of a large number of new services and applications. Combined with the progress of technology, it has already changed the way we live and work, and has significantly affected the relationships between network operators and service providers [1].

At the dawn of this new century, an even more dramatic change is about to happen: the relentless progress of electronic integration and wireless communication is going to pave the way to infrastructureless, self-organized networks. Self-organization means that the network runs solely by operation of end users; an early low-technology example of a self-organized network is the citizen-band (CB) voice analog network.

Of course, such networks will neither be a replacement nor an alternative to current and future infrastructure-based networks. Instead, they will complement these infrastructures in cases where cost, constraints, or environment

require self-organized solutions. For this reason, in most cases they will have to be interconnected with other more conventional networks such as the Internet and cellular networks.

The next section will provide an overview of mobile ad hoc networks. We will then discuss the paradigm shift this new approach is likely to bring, and present the broad range of research opportunities generated by this shift. As an illustration of these concepts, we will describe a long-term research project we began recently, called *terminodes*. Finally, we will discuss the overall potential of this new generation of networks.

MOBILE AD HOC NETWORKS

AN OVERVIEW

The basic idea of mobile ad hoc networks is not recent: as early as the '70s, research projects led to significant results in terms of radio resource allocation and network organization. At that time, these networks were usually called *packet radio networks* [2].

In these early generations, each node was considered an individual, handheld device. The main thrust was on military applications (providing person-to-person communications on the battlefield). With the progress of technology, new kinds of networks have emerged, in which the nodes are projected to be millimeter-sized sensors [3]. At the same time, several new potential application fields have been identified, such as relief operations in natural or manmade disasters, car-based networks (for vehicle cooperation, notably), and the provision of wireless connectivity in remote areas.

Because of their very nature, mobile ad hoc networks have a certain number of peculiarities. First of all, as already mentioned, they can act independent of any provider. Moreover, they have to be highly cooperative: the tasks are distributed over the nodes, and any operation is the result of the collaboration of a group of them. In addition, the nodes rely on batteries for their energy; therefore, energy savings are an impor-

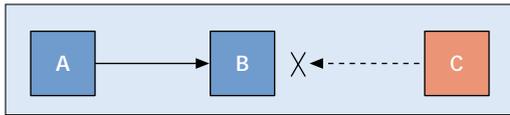


Figure 1. *The hidden terminal problem. Assume A is transmitting to B, and C is out of the range of A. If C starts transmitting, a collision occurs in B; A is said to be hidden from C.*

tant system design criterion; furthermore, nodes have to be power-aware: the set of functions offered by a node depends on its available power. The topology can be highly dynamic, hampering the stability of the links and routes. The links themselves are wireless; their dependability and capacity must be carefully scrutinized. Finally, security is difficult to implement because of the vulnerability of the links and the limited physical protection of each node.

These characteristics require specific solutions, which we are going to detail hereafter. Since most of the research has been focused on networking issues, we will start with them.

TECHNICAL ISSUES

Routing is a central function in any network. Routing in ad hoc networks poses two specific challenges [4]. First, traditional solutions (e.g., in the Internet or cellular phone networks) aim to quickly propagate changes in topology or reachability; however, they are designed under the assumption that the network is relatively stable. In contrast, in a mobile ad hoc network, the topology of the network is constantly changing. Second, traditional routing solutions rely on some form of distributed routing databases, maintained by the operators in either the network nodes or specialized management nodes. In mobile ad hoc networks, nodes cannot be assumed to have persistent data storage, and they cannot always be trusted. The Internet Engineering Task Force (IETF) mobile ad hoc networking group (MANET [5]) studies solutions that mainly address the first challenge. Such solutions are based on optimization of the traditional ingredients of distance vector, link state, and source routing. Addressing the second challenge may require a departure from such solutions. *Geographic methods* may be useful for this purpose: nodes are informed of their own geographic position, and those of other nodes; routing is based on this information. Geographic methods reduce the routing function of intermediate systems to the bare minimum, and avoid distributed databases. They can also be used as a complement to more conventional methods [6].

A network must **manage the mobility** of its terminals, and therefore be able to locate any of them. In particular, if a terminal wants to communicate with another, it will make use of the address of the latter; the network will have to locate it in some way. The simple solution of broadcasting a paging message to the whole network does not scale. In cellular networks, the location of the mobile stations is stored in centralized servers (e.g., the home location register of the Global System for Mobile Communica-

tions, GSM). The self-organization of an ad hoc network precludes the existence of such servers. Therefore, the information about location must be distributed among the nodes themselves. A solution to this problem is provided in [7]. Some researchers have also proposed to *predict* the future location of the nodes based on mobility patterns [8].

Mobile ad hoc nodes use **IP addresses**. For small mobile ad hoc networks, addresses are allocated in the traditional way, with an IP prefix identifying the mobile ad hoc network. For large-scale networks, the topology-based address allocation currently used in the Internet may not be optimal. In contrast, a node address should be interpreted as a stable node identifier, which carries no specific topological information.

The **transport layer** also requires specific attention in an ad hoc network. The experience gained from conventional wireless networks has shown that the performance of TCP may be severely degraded. TCP interprets losses as a signal of congestion and strongly reduces its sending rate, whereas wireless links may temporarily exhibit high loss rates due to transmission errors not related to congestion. A solution is to avoid the use of TCP, by putting data integrity in the application layer and using a rate-controlled forward-error-correcting application layer protocol.

The **radio interface** can be engineered in different ways, based on the requirements of a specific system. Several issues need to be taken into account. First, the signal strength decreases as the square of the distance, if not worse. Second, some of the traditional multi-access protocols used for wireline LANs cannot be used; for example, collision detection is not appropriate because a node is usually unable to listen while it is transmitting. Third, two terminals may unknowingly interfere at a third one. This situation is usually called the *hidden terminal problem* (Fig. 1). A possible solution consists in making use of carrier sense multiple access with collision avoidance (CSMA/CA), or defining master and slave roles, such as in Bluetooth, therefore entitling the master to control the allocation of the transmission resource in its power range.

Security is an issue of critical importance for most networks, and mobile ad hoc networks are no exception. Several security features can be required, such as availability of service despite denial-of-service attacks, confidentiality, integrity, authentication, and nonrepudiation. Guaranteeing these features is a major challenge. Solutions based on watchdogs and the use of tamper-proof hardware are being researched. In addition, in order to fulfill the requirement of self-organization, key management has to be totally distributed [9].

Power management is almost always a difficult issue in wireless networks. In the case of ad hoc networks, there are essentially two concerns. First, power has to be fine-tuned in order to maximize the throughput of the network: the higher the power, the larger the transmission range of the node, but also the higher the interference from other signals. As shown in [10], the best trade-off is obtained when there is an aver-

For large-scale networks, the topology-based address allocation currently used in the Internet may not be optimal. In contrast, a node address should be interpreted as a stable node identifier, which carries no specific topological information.

The attraction of self-organization is evident, be it from an intellectual challenge or a practical, operations point of view. But of course, the key question is if it can be done at all, and if so, up to what scale.

age exactly one packet in transit over each hop. Second, since the nodes are usually battery-operated, it is important to minimize their consumption. A typical solution consists in turning the devices to a sleep or idle mode whenever they have nothing to accomplish. A further optimization consists in making use of energy-conserving routing (and multicast).

Finally, it has to be mentioned that a number of **standards and products** that allow the development of small scale ad hoc networks are emerging. Wireless local area products (IEEE 802.11, Hiperlan) are now widespread and provide in-building wireless access; however, they are usually deployed as access links only, with packet relay performed by traditional bridges or routers. Bluetooth is a low-cost technology for short-range communications; its market is targeted toward PCs, phones, appliances, watches, and so on. It allows multiple nodes to connect to each other in a multihop arrangement.

THE PARADIGM SHIFT AND SOME OPEN RESEARCH QUESTIONS

As seen in the previous section, the development of self-organized ad hoc networks requires research in many technical areas. A wealth of contributions have been made on most of the topics we have described. However, the *overall potential* of this new family of networks has not been fully investigated yet.

This section will take a broader look, and sketch how the combination of all these techniques is going to lead to a paradigm shift for communications.

In our vision, the key word is self-organization, which cuts through all layers of the system, from the physical to the network and application layer. The attraction of self-organization is evident, be it from an intellectual challenge or a practical operations point of view. But of course, the key question is whether it can be done at all, and if so, up to what scale.

Let us start with an **information theoretic** view of the problem. Finding the capacity region (i.e., the set of achievable rates) of such a multi-terminal system is an open problem of multi-user information theory. A recent paper [11], based on a simple interference model (all nodes in an ad hoc network interfere in an omnidirectional fashion with a power decay law), derives a very interesting result. If there are N nodes in a bounded region (e.g., a circle) attempting arbitrary point-to-point communications, the total throughput capacity of an ad hoc wireless network grows at \sqrt{N} , which implies that the throughput per node decreases at $1/\sqrt{N}$. Thus, the total efficiency benefits from the number of nodes, but the performance per node decreases with the number of nodes. The article also provides a construct to achieve this rate. It is to be pointed out that:

- There is no collaboration between the users.
- The users are static.
- The antennas are omnidirectional.

Nonetheless, the result indicates that the promise of an ad hoc network with millions of nodes might be bogged down by interference, a

result that should come as no surprise to communications engineers. And solving the problem of interference might make the system less attractive for economic reasons. Progress in intelligent antennas and multi-user detection can significantly improve the situation.

Moving up in the layers, consider **networking issues**. As discussed in the previous section, self-organized routing is challenging, even if location is known. Questions of scaling to very large sizes are open in the methods currently proposed. However, avoiding centralized information could make the system more robust, easier to set up and operate. When location is unknown, the complexity grows even further. Questions of resource allocation and flow control, including original datagram services suited for highly redundant but changing topologies, are topics of investigation.

Next, consider services built on top of a networking layer. Among the most challenging are **real-time services** (e.g., interactive speech communication). Given the uncertainty in the layers below, special techniques to create real-time robust services are needed. At the same time, we expect that the density of nodes will help create redundancy that can be useful. We will describe later our approach to deal with this problem.

Let us look at one more layer, the **information system layer** that is built on top of the communications infrastructure. Constructing a decentralized information system that cannot rely on any centralized server and is robust to nodes being randomly up or down is one of the challenges. In addition, persistence of key information despite the unreliability of both communication and node availability is another challenge. Also, building up databases in a self-organized fashion is an interesting topic in itself. Such self-organized information systems have appeared recently in order to cope with server failures. For a recent example, see Gnutella, a distributed version of music download on the Internet.

THE TERMINODES PROJECT

In order to illustrate the research opportunities described in the previous section, we will now present a long-term research project (2000–2010), defined by the authors and their colleagues. This project aims to study and prototype large-scale self-organized mobile ad hoc networks. In this framework, the nodes are called *terminodes* (terminal+ node). The project distinguishes itself from other research projects in this area in several ways. First, it encompasses all layers and explores interlayer interactions, from the fundamentals of the physical layer up to software architecture and applications; second, it is free from short-term compatibility constraints such as interoperation with IP networks (which will be studied in the project, however); last but not least, it tries to capture the business and societal potential generated by the paradigm shift previously described. More specifically, this project can be perceived at three levels of abstraction:

- At the most concrete level, terminodes are a **technical challenge**: our purpose here is to make innovative contributions in the area of self-organized mobile ad hoc networks.

- At the intermediate level, terminodes are an **intellectual fantasy**; in other words, they are a way to stimulate creativity in order to identify new research challenges. An example already generated in this way is to try to define a formal model for fair exchange.
- At the highest level of abstraction, terminodes are a **societal/political vision**, in which the ultimate aim of self-organized communications is considered. Indeed, we believe that long-term research should contain a dimension that goes beyond purely technical questions. We will come back to this in the final section.

In the rest of this section we will describe some of the first results already obtained in the areas of networking, real-time services and software. Additional information is available on the Web site of the project (<http://www.terminodes.org>).

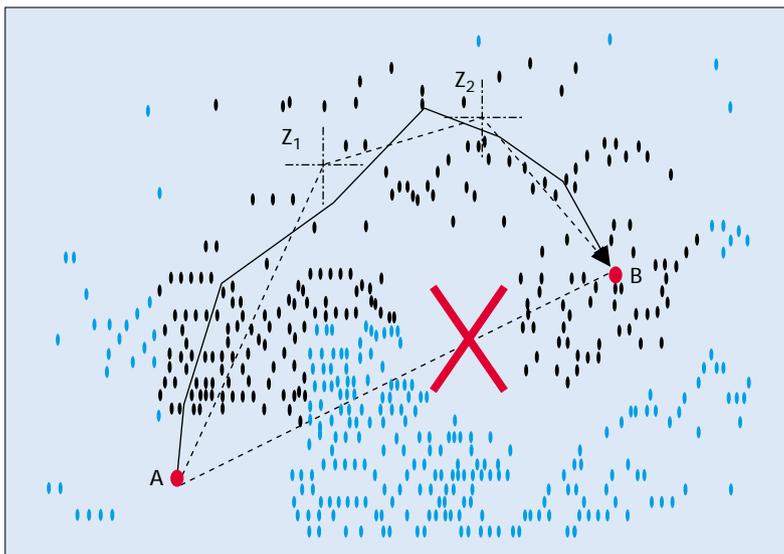
NETWORKING ISSUES

In this subsection we first explain how a packet can travel from a terminode to another. Then we explain how positioning can be achieved, even in the absence of global positioning system (GPS) receivers. Finally, we see how we can motivate users to cooperate with each other.

Each terminode has a permanent, unique node identifier, called the *end system unique identifier* (EUI), burnt in the hardware and mappable to an IPv6 address (we will request a reservation of a portion of the IPv6 addressing space for terminodes). In addition, it has a time-varying *location-dependent address* (LDA). For outdoor terminodes with a GPS receiver, the LDA is the location given by GPS. Otherwise, a relative positioning system is used, as described later. Terminode routing uses a combination of EUI and LDA in order to address the requirements of scalability, reduction of intermediate system functions, and robustness against malicious users. For unicast, it is based on *geodesic packet forwarding*: at each hop, the packet is forwarded to the neighbor closest to the direction in which the destination is located. However, this may not always work, since the aggregate transmission range of the terminode network is not necessarily convex. In order to circumvent holes, *anchored geodesic packet forwarding* is used. A source node, say A, must first acquire some approximate value of the destination's (say B's) LDA. Anchors are used to get the packet across the wide area, close to B's location (Fig. 2).

Once the packet comes close to B's location, another routing method is used (*terminode local routing*), which is a variant of MANET routing. Terminode local routing uses the destination's EUI, not its location. This is based on the assumption that the location of a terminode can be considered a valid indication only if the packet is far from the destination. In close vicinity to B, the location information is likely to be inaccurate, and terminode local routing uses the classical methods of link state, distance vector, and source route to reach B. Anchors are discovered by means of a variety of exploration protocols, currently under study. Multicast is left for further study.

Mobility management is based on the *virtual home region* (VHR) concept: we define a unique



■ Figure 2. Anchored geodesic packet forwarding is used across the wide area. A is sending data to B. The straight line to B does not work due to a gap in the terminode network. A uses two intermediate locations, Z_1 and Z_2 , called anchors. They are discovered by means of an appropriate method based on "small world" graphs. The packet is forwarded approximately along the dotted line AZ_1Z_2B . The actual path (plain line) uses intermediate terminodes, and differs from the dotted line.

hash function H , known to all terminodes, by means of which we can associate a point in the space with each terminode: $C = H(\text{EUI})$. The VHR is then defined as the disk centered in C of radius R , where R is a system parameter that depends on the local terminode density. All the terminodes located in the VHR are in charge of storing the current LDA of the terminode identified by the EUI.

A related area of concern is the awareness by each terminode of its own geographic location. As we have seen, this information is required by the mechanism of packet forwarding. A straightforward solution consists in integrating a GPS receiver in each terminode. However, this solution would be inapplicable in several cases, notably for cost or packaging reasons, or when the received signal is too weak (e.g., indoor operation). In order to cover these cases, we have devised a distributed algorithm based on measurement of the propagation time of the signal between pairs of terminodes; it provides each terminode with an estimate of its position expressed in a relative coordinate system.

Because of self-organization, service availability is a major requirement. There are two issues. First, end users must be given incentive to cooperate (and notably to relay packets for the benefit of other users). Second, they must be encouraged to not overload the network. In the mobile ad hoc networks proposed and implemented so far, these two aspects were irrelevant because of the small size of the network and the emergency situations in which they were expected to be deployed. We have devised a mechanism based on a virtual currency called a *nuglet*: terminodes remunerate each other for the services they provide to each other (Fig. 3).

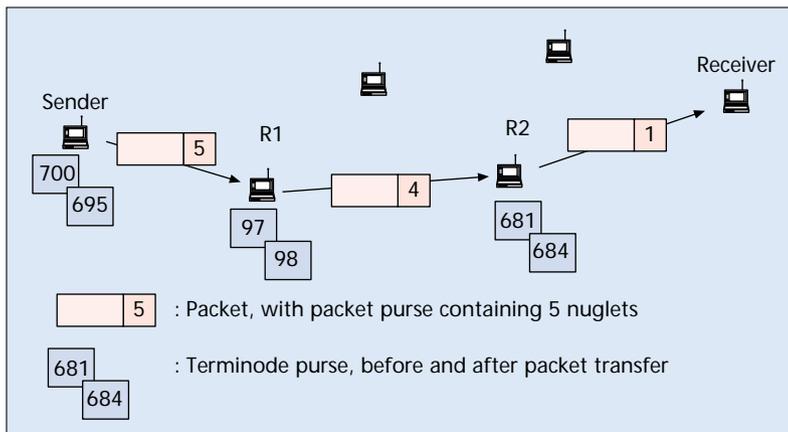


Figure 3. Virtual currency: nuglets. The sender puts 5 nuglets in the packet purse. The first relaying node (R1) charges 1 nuglet, which is taken from the packet purse. The second relaying node, R2, charges 3 nuglets. This higher greediness can be motivated, for example, by the fact that R2 faces a higher energy expenditure to relay the packet (typically if the distance to the next hop is higher).

REAL-TIME SERVICES OVER AD HOC NETWORKS

As hinted at earlier, providing real-time services such as voice or video over ad hoc networks represents a very challenging task. Besides the highly unreliable channels, stringent delay constraints do not permit using retransmission and/or large delay error correcting codes (the latter would be useless anyway if a link goes down completely, a rather likely event).

So what can be done to provide service in such an unpredictable environment? The key is to use the redundancy that would probably be present in such large and dense networks. Rather than trying to make one connection survive at all costs, the network should take advantage of the number of possible connections that can be established. Over this set of connections, various techniques are now possible to obtain higher reliability (e.g., using error correction codes across various parallel connections).

We will discuss in more detail a method that combines signal processing techniques for source coding with channel awareness for channel coding: multiple description coding. We are not arguing that this is the only or best solution, but rather take this as an example of cross-layer interactions. Recall that in a traditional communications system, the separation principle tells us to first compress the source as much as possible (i.e., to represent it with as few bits as possible while maintaining a certain quality level) and then to design a channel coding method to transport these bits to the destination. In multiple description coding, one is less optimistic: one actually expects some packets carrying information to never make it to the destination, but there is no prior knowledge of which packets are more or less likely to survive. Thus, one does not “put all one’s eggs in the same basket,” but chooses to send several different descriptions of the source in different packets.

To give an intuitive feeling of the change of approach, Fig. 4 shows on the left a regular sub-band decomposition (as used in standard sub-

band image coding), while the right side shows a decomposition as used in multiple description coding. In the second case, one can see the redundancy left between subbands, which allows good reconstruction even if not all subbands are received. Of course, there is a price to pay: if all packets do arrive, the performance is not as good as in regular compression; however, if only partial information is received, one can recover an adequate approximation.

SOFTWARE ASPECTS

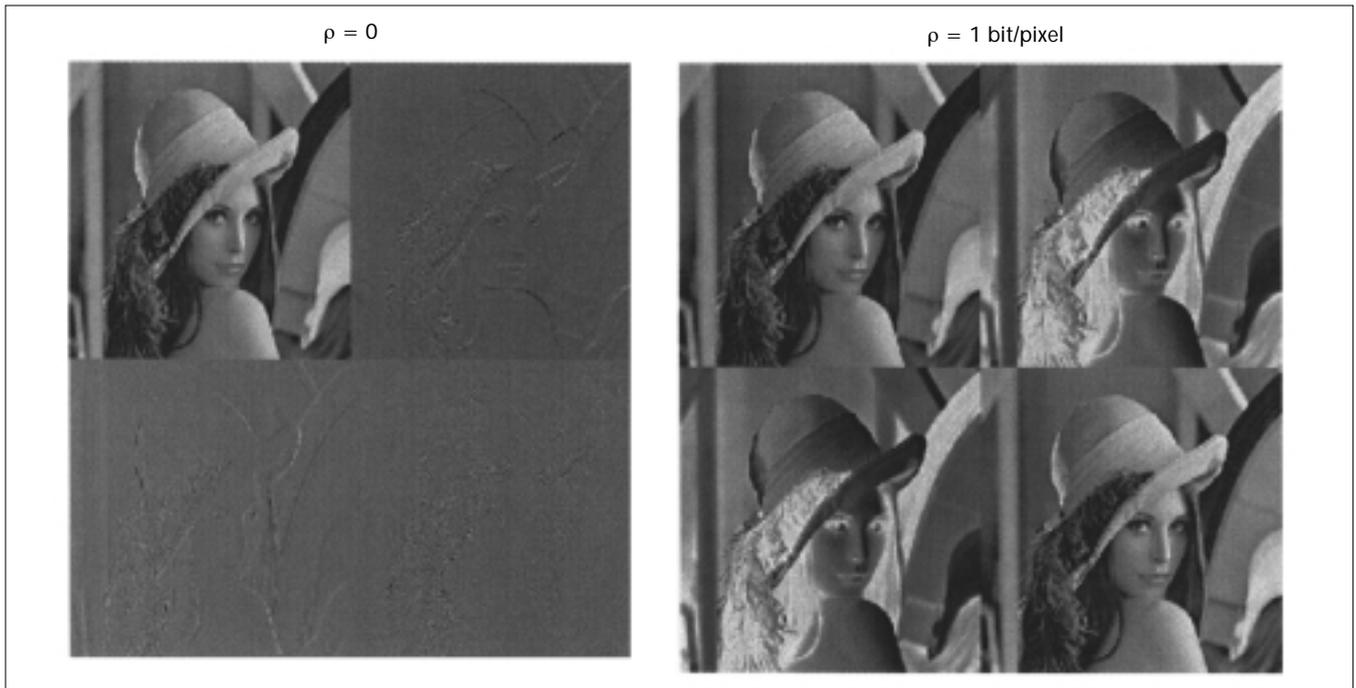
Software plays a critical role in mobile ad hoc networks, because software ultimately decides what kind of system is presented to users. The terminode project provides a framework to understand the software issues and gain practical experience. There are (at least) three issues:

- The (software) implementation of the base software, that is, the base “operating” system including the software for the various protocols, such as the implementation of the routing algorithms, accounting system, and security system.
- The (application) software that makes a collection of terminodes useful for a client, that is, the software that provides application-specific functionality. An example could be the software application that allows a client to use a terminode to perform a transaction.
- The software architecture of the overall system that structures how the various software components fit together. A crucial issue is that the software architecture support flexibility (both the base software and application software are likely to change frequently).

The characteristics of the base operating system (which controls the operation of a terminode device) will depend to a large degree on the properties of a terminode. If the terminode is a full-fledged computer, the operating system is provided by the hardware platform. If the terminode is a handheld communication device (e.g., to be used after some event has rendered the fixed communication infrastructure nonoperational), the operating system may be nothing more than a simple controller.

The development of application software for terminodes is still in its infancy. Thus, the other parts of the system must support flexibility. It is unlikely that any terminode will have *all* desired application software loaded. And if mobile wide-area ad hoc networks become a reality, with a large user base, new applications will be developed. Both constraints call for flexibility of the software architecture of a terminode system.

Flexibility, however, is not the only concern. Since each terminode carries (real or virtual) currency, it is also important that the system is not easily attacked and that an application can only access the services to which it is entitled. For example, if an application is granted the right to send X Mb/s, we must ensure that the application stays within this limit. Or, if the application has the right to use a segment of the frequency spectrum, the system must prevent the application from going outside this band.



■ Figure 4. Regular subband coding (left) and multiple description subband coding (right).

The solution we adopt is based on the concept of a *contract*. Each software component (e.g., an application module or a middleware module) specifies its resource demands and constraints. This contract is then processed by a sophisticated *loader*. If, at the time the component is brought into the system, the loader cannot establish that the component will obey its contract, appropriate dynamic checks have to be enabled to observe the runtime behavior.

DISCUSSION AND CONCLUSION

As we have seen, self-organized mobile ad hoc networks have the potential to bring a dramatic paradigm shift in communications. We illustrate this by Table 1, which gives a global view of the evolution of networking over several decades. Each row corresponds to a generation of networks:

- Telecommunication networks such as the public switched telephone network
 - The Internet
 - Self-organized mobile ad hoc networks
- For each generation, the table shows who (the end user, operator, or vendor) is primarily in charge of each of the most important issues: the infrastructure (including billing), security, and applications. It illustrates the strong tendency toward an increase of responsibility of end users. For example, the column related to security shows that in telecommunication networks the telcos are in charge of security; in the Internet, the Internet service providers have an important role to play, but users can organize themselves independently (e.g., with solutions based on public key cryptography such as Pretty Good Privacy). Finally, in self-organized mobile ad hoc networks, security is mastered by end users, who will make use of appropriate products provided by vendors.

The Internet has shown that when users are

Network	Infrastructure	Security	Applications
Telecom networks	Telcos	Telcos	Telcos (IN)
Internet	ISPs + telcos	ISPs + users (PGP)	Users
Self-org. ad hoc NW	Users + vendors	Users + vendors	Users

■ Table 1. The overall evolution of networking.

empowered with end systems, they can program at will; human creativity brings the amazing wealth of applications we witness today — e-mail, the Web, e-commerce — with a dramatic impact on society (globalization of the economy is but one example). Self-organized mobile ad hoc networks, if successfully developed, convey the promise of going a step further, entrusting users with the operation of the network.

Another fundamental question is related to our relationship with communications. Personal communications have been the dominant paradigm so far, but mobile ad hoc networks open new possibilities, such as the communication between objects (e.g., by means of sensor networks). The progress of technology will soon make it possible to produce *billions* of these devices at low cost. Here also, the usage of this potential requires appropriate investigation.

Very clearly, over the 20th century the impressive development of communications technology has been fueled by the usual rules of capitalism: identification of markets, and provision of products and services. We have probably reached a plateau now: mankind does not need so much to be saturated and distracted by additional “smart” appliances or new features on fixed and mobile devices. Rather, the real needs will be the “global” ones: bring development and democracy to

Personal communications have been the dominant paradigm so far, but mobile ad hoc networks open new possibilities, such as the communication between objects.

the Third World, and respond to the ecological challenges. We believe that self-organized mobile ad hoc networks can be a vehicle to accomplish these worthy tasks. How and when are the fascinating questions we intend to consider.

ACKNOWLEDGMENTS

The authors would like to thank the many contributors to the definition of the Terminodes project, notably Karl Aberer, Gustavo Alonso, Gérard Ben Arous, Ljubica Blazevic, Jacques Bovay, Levente Buttyan, Srdan Capkun, Jean-Dominique Decotignie, Silvia Giordano, Rachid Guerraoui, Maher Hamdi, Friedemann Mattern, Yves Pigneur, Bixio Rimoldi, André Schiper, Sergio Servetto, Anja Skrivervik, Sabine Süstrunk, Emre Telatar, Lothar Thiele, Patrick Thiran, Rudiger Urbanke, Serge Vaudenay, and Alain Wegmann. They also wish to thank Matthias Grossglauser and Taieb Znati for their useful comments on early versions of the article. Finally, they would like to extend their gratitude to Jim Massey for having suggested the nuglets neologism, and to Pierluigi Dragotti for providing Fig. 4.

REFERENCES

- [1] D. Messerschmitt and J. P. Hubaux, "Opportunities for E-Commerce in Networking," *IEEE Commun. Mag.*, Sept. 1999.
- [2] J. Jubin and J. D. Tornow, "The DARPA Packet Radio Project," *Proc. IEEE*, 1987.
- [3] J. M. Kahn, R. H. Katz, and K. S. J. Pister, "Next Century Challenges: Mobile Networking for Smart Dust," *MOBICOM '99*, Seattle, WA, 1999.
- [4] D. A. Maltz *et al.*, "The Effects of On-Demand Behavior in Routing Protocols for Multi-Hop Wireless Ad Hoc Networks," *IEEE JSAC*, Aug. 1999.
- [5] Mobile Ad-hoc Networks (manet) WG, Mobile Ad-hoc Networks (manet) Charter, Wg charter, IETF, 1999; <http://www.ietf.org/html.charters/manet-charter.html>
- [6] Y. Ko and N. H. Vaidya, "Location-Aided Routing (LAR) Mobile Ad Hoc Networks," *MOBICOM '98*, Dallas, TX, 1998.
- [7] J. Li *et al.*, "A Scalable Location Service for Geographic Ad Hoc Routing," *Mobicom 2000*, Boston, MA, 2000.
- [8] A. B. McDonald and T. F. Znati, "A Mobility-Based Framework for Adaptive Clustering in Wireless Ad Hoc Networks," *IEEE JSAC*, Aug. 1999.
- [9] L. Zhou and Z. Haas, "Securing Ad Hoc Networks," *IEEE Network*, vol. 13, no. 6, Nov./Dec. 1999, pp. 24–30.
- [10] L. Kleinrock, "On Some Principles of Nomadic Computing and Multi-Access Communications," *IEEE Commun. Mag.*, July 2000.
- [11] P. Gupta and P. R. Kumar, "The Capacity of Wireless Networks," *IEEE Trans. Info. Theory*, Mar. 2000.

ADDITIONAL READING

- [1] P. L. Dragotti, S. Servetto, and M. Vetterli, "Multiple Description Subband Coding," submitted for publication, Fall 2000.

BIOGRAPHIES

JEAN-PIERRE HUBAUX (jean-pierre.hubaux@epfl.ch) joined the Swiss Federal Institute of Technology, Lausanne (EPFL) as an associate professor in 1990; he was promoted full professor in 1996. He is co-founder and co-director of the Institute for Computer Communications and Applications. His current research is focused on mobile networking with special interest in self-organized mobile ad hoc networks and therefore in terminodes. At the beginning of his activity at EPFL, he defined the first curriculum in communication systems. In October 1999 he became the first chair of the newly created communication systems department. Before joining EPFL he spent 10 years in France with Alcatel, where he was involved in R&D activities, mostly in the area of switching systems architecture and software.

THOMAS GROSS is a professor of computer science at ETHZ, Switzerland. He received a Ph.D. in electrical engineering from Stanford University and was with Carnegie Mellon University, 1984–2000. He is interested in tools, techniques, and abstractions for software construction (especially compilers), and has worked on many aspects of the design and implementation of software systems. Further information (including downloadable versions of many papers) can be found at <http://www.lst.inf.ethz.ch>.

JEAN-YVES LE BOUDEC (leboudec@epfl.ch) is full professor at EPFL, Switzerland. He graduated from Ecole Normale Supérieure de Saint-Cloud, Paris, where he obtained an Aggregation in Mathematics in 1980. He received his doctorate in 1984 from the University of Rennes, France. He worked at Bell Northern Research, Ottawa, Canada, as a member of scientific staff in the Network and Product Traffic Design Department, then at IBM Zurich Research Laboratory where he was manager of the Customer Premises Network Department. He joined EPFL in 1994. His interests are in the architecture and performance of communication networks.

MARTIN VETTERLI got his engineering degree from ETH in Zurich, his M.S. from Stanford, and his doctorate from the Swiss Federal Institute of Technology in Lausanne (EPFL). He was on the faculty of Columbia University, New York, and the University of California at Berkeley before joining the Communication Systems Department of EPFL. He works on signal processing and communications; in particular, wavelet theory and applications, image and video compression, joint source-channel coding, and self-organized communication systems. His work has won him several prizes (best paper awards from EURASIP in 1984 and the IEEE Signal Processing Society in 1991 and 1996); he is the co-author, with J. Kovacevic, of the book *Wavelets and Subband Coding* (Prentice Hall, 1995).