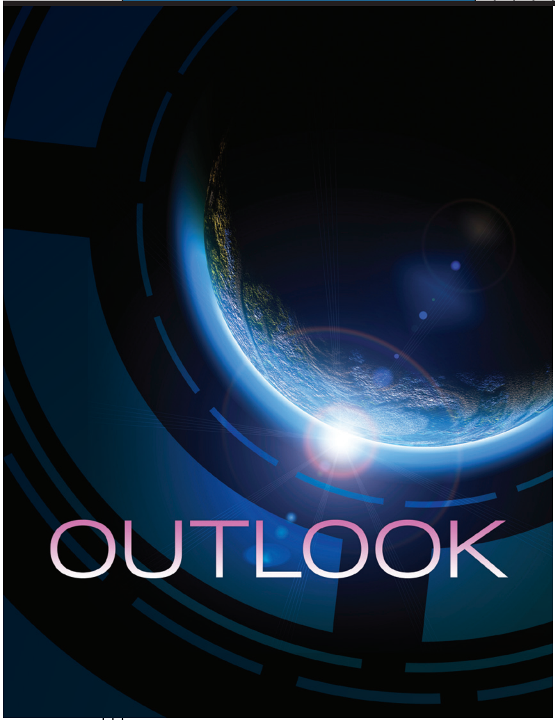

OPPORTUNITIES IN OPPORTUNISTIC COMPUTING

Marco Conti, *Italian National Research Council*
Mohan Kumar, *University of Texas at Arlington*

When two devices come into contact, albeit opportunistically, it provides a great opportunity to match services to resources, exchange information, cyberforage, execute tasks remotely, and forward messages.

In recent years, opportunistic networks have gained popularity in research and industry as a natural evolution from mobile ad hoc networks (MANETs). In opportunistic networks, nodes come into contact with each other opportunistically and communicate wirelessly. Opportunistic networks are human-centric because they opportunistically follow the way humans come into contact. Therefore, opportunistic networks are tightly coupled with social networks and can exploit human relationships to build more efficient and trustworthy protocols.

OPPORTUNISM

Technological advances are leading to a world replete with mobile and static sensors, user cell phones, and vehicles equipped with a variety of sensing and computing devices, thus paving the way for a multitude of opportunities for pairwise device contacts. Opportunistic computing exploits the opportunistic communication between pairs of devices (and applications executing on them) to share each other's content, resources, and services.

Opportunistic computing opens an exciting avenue for research and development, one hitherto not fully exploited, and at the same time expands the potential of opportunistic networks for real-life application problems.

Feasibility study

Yet we may still wonder about the feasibility of such a computing paradigm based on pairwise node contacts. Several motivating factors have led to the concept of opportunistic computing: key technological developments, wireless communications, and device architectures; burgeoning application areas; human-centered pervasive computing; and society-influenced social networks.

Mobile cell phones with integrated technology such as Wi-Fi, cameras, Bluetooth, and other, similar capabilities—along with embedded computing devices in moving vehicles and mobile and static sensory devices, including surveillance cameras and others—are available worldwide at reasonable costs. The widespread use of these devices creates a huge number of contact opportunities that are key to opportunistic communications. Significantly, the frequency and potential of opportunistic contacts are mind-boggling, thanks to the cell phone market's estimated 22 percent annual growth rate.

Published by the IEEE Computer Society

Worldwide market

Analysts estimate that 3.3 billion people worldwide use cell phones (www.reuters.com/article/technologyNews/idUSL2917209520071129)—a little more than half the world's population. This amounts to an estimated potential of one billion parallel opportunistic contacts worldwide at any given time, assuming two billion cell phones are turned on at that time. A conservative look at each cell phone's processor reveals a performance figure of 100 MIPS and communication at 200 Kbps. Exploiting these opportunistic contacts gives the potential to perform approximately one quadrillion processing tasks, and exchange 1 petabyte of data per second.

If we consider the 10 billion ARM processors (http://arm.com/products) in use today in embedded systems such as vehicles, appliances, and other devices, the estimates would be much higher. Indeed, a typical downtown or university has hundreds of—or O(100)—street cameras; O(1,000) user cell phone cameras; O(1,000) user devices, including laptops, PDAs, and cell phones; and O(100) desktops and information servers.

Given the plethora of wired and wireless communication technologies, such as Wi-Fi, Bluetooth, cellular, and WiMax, along with device capabilities, opportunistic contacts among pairs of devices are the norm rather than a rarity. The necessary infrastructure for opportunistic computing is thus all-pervasive. Opportunistic networks provide the concrete communications support, while the application scenarios provide the motivation to address opportunistic computing challenges. It is only a question of how and when we overcome the challenges to enhance existing applications and develop new ones. In effect, large-scale opportunistic computing, which can simply be defined as *delay-tolerant distributed computing* (DTDC), has tremendous potential.

BACKDROP

Up through the mid-1990s the computer occupied the center of the computing universe, and users were expected to make all possible adaptations to use it and its associated resources. In recent years, there has been a perceptible shift, however, as researchers and designers began focusing on the user. In his seminal paper, Mark Weiser[1] prophesied that the advent of pervasive computing would enhance the user experience. M. Satyanarayanan[2] made the case that while technological advances are inevitable, user attention is constant.

Indeed, pervasive technologies from smart spaces to iPhones have begun to recognize the user as the computing universe's center. Opportunistic computing takes pervasive and mobile computing further, to the as yet to be explored but highly potential realm of users' social behavior. Specifically in opportunistic computing, this behavior is utilized to send and forward messages, acquire and disseminate information, and acquire and distribute resources.

Hitherto, with few exceptions, there have been clear distinctions between servers and clients, producers and consumers of information, service providers and consumers, and resourceful and resource-poor entities. However, the Internet's encroachment into everyday life and the development of new technologies such as social networks and peer-to-peer (P2P) systems on the one hand, and anytime, anywhere wireless communications on the other, has made the gap between these sets narrower.

Further, the pervasive deployment of sensors and radio-frequency IDs (RFIDs) has enhanced the potential of user-generated information. Content is increasingly generated in a participatory fashion by the users themselves, following the user-generated content (UGC) model best exemplified by Web 2.0 services such as blogs, YouTube, and Flickr, along with grassroots journalism and similar movements. The users have become both content producers and consumers.

Given the plethora of wired and wireless communication technologies, opportunistic contacts among pairs of devices are the norm rather than a rarity.

The proliferation of powerful mobile devices and UGC models is strictly intertwined. Clearly, a user with a cell phone in hand that includes a camera, microphone, speaker, and perhaps an RFID reader has become a producer as well as a consumer of information and services.[3] The more powerful the user devices are, the likelier it is that they will generate and share content, leading to omnipresent content in devices and the diminishing role of centralized servers.

Utilizing the full potential of opportunistic contacts requires new networking and computing paradigms. While in the past few years significant research efforts have focused on exploiting opportunistic contacts to develop mainly message or content-forwarding applications,[4] opportunistic computing initiatives are still in their infancy.

Several mobile and pervasive computing projects have attempted to exploit all available resources in the environment in the presence of a significant degree of connectivity among the computing devices. Opportunistic computing's major challenge is to effectively utilize opportunistic contacts to make information available and accessible and to provide collaborative computing services to applications and users. The challenges include content distribution and management in an opportunistic P2P environment; management and sharing of scarce and seemingly disconnected resources; remote execution of tasks in a delay-tolerant environment; and cross-layer issues such as trust, authentication, and privacy.

THE ENABLERS

Several major challenges must be addressed to make opportunistic computing a reality, but significant research trends also push in this direction. Indeed, from the network perspective, there are already several examples of the opportunistic network paradigm's effectiveness, from special-purpose networks like the Sámi Network to general-purpose networking like the Haggle project (www.haggleproject.org).[4] At the same time, several application scenarios, from crisis management to pervasive healthcare, are emerging that naturally benefit from the opportunistic paradigm.

Opportunistic computing exploits humans' mobility and their gregarious nature to enable a transmission only if two users are sufficiently close.

Opportunistic networks

While human centricity lies at the core of pervasive computing's vision, legacy wired and wireless network architectures force human communications to follow network engineering paradigms. For example, exchanging a message between two participants at a conference entails at least two mail servers spread across the world—just to carry a few hundred bytes across a local space. Clearly, this overly engineering-centric communication paradigm is derived from the wired Internet. Wireless communications and mobile computing freed computing from the leash of tethered networking, but in these networks, mobility and the related disconnections continue to be an engineering challenge instead of an opportunity to communicate.

Mobility management in MANETs exemplifies the engineering-centric approach in the design of self-organizing networks: Mobility is a challenge to cope with, and routing-protocol design focuses on building stable end-to-end paths, as do mobile nodes. Opportunistic networks represent the first attempt to close the gap between human and network behavior by taking a user-centric approach to networking and exploiting user nodes' mobility as an opportunity—rather than a challenge—to improve data forwarding.[4]

Basically, this approach exploits humans' mobility and their gregarious nature to enable a transmission only if two users are sufficiently close. It might seem that the probability of a source coming into contact with a destination is rare, but the use of opportunistic, delayed paths comprising one or more opportunistic contacts between the source and the destination transforms this simple idea into a powerful one given the potential of opportunistic contacts. This communication model constitutes a theoretical basis for opportunistic networking.

In opportunistic networks such as MANETs, the communication is multihop, with intermediate nodes acting as routers that forward the messages addressed to other nodes. In this case, however, forwarding is not "on the fly" because intermediate nodes such as mobile relays store the messages when no forwarding opportunity exists—for example, there are no other useful nodes in the transmission range—and exploit any contact opportunity with other mobile devices to forward information.

For this reason, developers refer to the forwarding paradigm as "store, carry, and forward." In opportunistic networks, the nodes' mobility creates opportunities for communication, unlike MANETs, in which mobility is viewed as a disruption.

In the literature, developers often refer to opportunistic networks as delay-tolerant networks. The DTN architecture and protocols are currently under study in the Internet Research Task Force's Delay Tolerant Networking Research Group (www.dtnrg.org), which is concerned with the interconnection of heterogeneous networks. The DTN approach is based on building an overlay atop disconnected networks to combat network disconnections through persistent storage. The overlay provides functionalities similar to the Internet layer (www.ietf.org/rfc/rfc4838.txt) even if end-to-end connectivity may never be available.

Opportunistic networking is a more general concept as it does not assume any compatibility with the Internet architecture, nor any a priori knowledge regarding the network topology, areas of disconnections, or future link availability. In opportunistic networks, route computations differ from those in traditional Internet- or MANET-routing algorithms. Specifically, forwarding and routing protocols are merged because routes are actually built while messages are forwarded. Indeed, routes must be computed on the fly and hop by hop as each message progresses toward its destination. Nodes carrying messages to be forwarded opportunistically evaluate if any other node they contact could provide a good next hop toward the destination, then hand over the message if so.[4]

The Haggle project has developed and implemented a novel layerless architecture for opportunistic networks. Specifically, this architecture has been implemented on mobile phones with the Windows Mobile OS. Preliminary experimental results have also shown the potential of this paradigm when using simple devices like mobile phones. Similarly, the Metrosense project[3] uses Nokia mobile phones with the Symbian OS as sensing devices and opportunistic carriers of the sensed information inside a city.

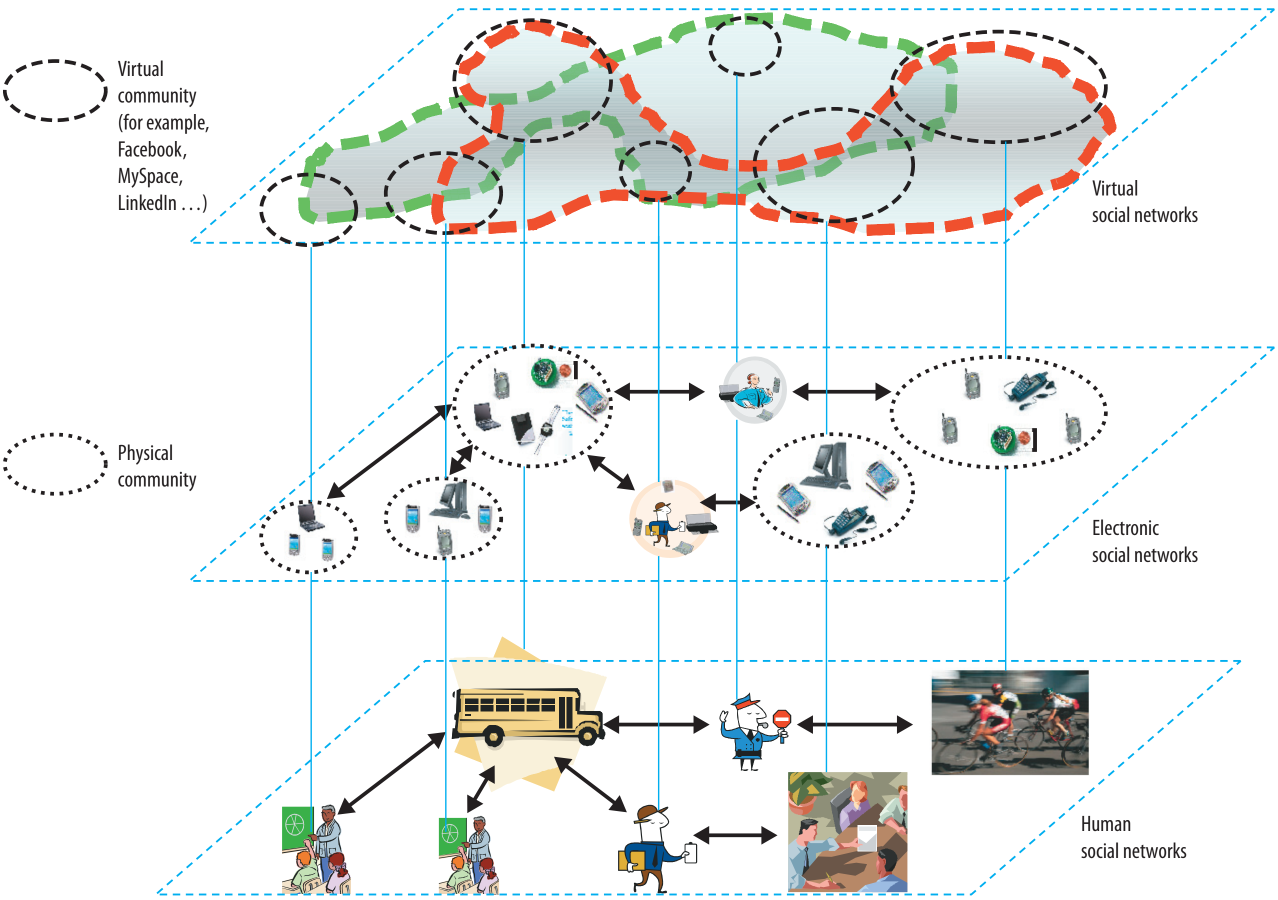

Figure 1. Human, electronic, and virtual social networks. Embedding the social relationships in the electronic world identifies at least two levels in an opportunistic environment: an electronic social network (where relationships depend on the physical properties) and a virtual social network that builds an overlay atop the electronic social network.

Social networks

Opportunistic networking tends to dissolve the traditional networking paradigms and integrate communication more closely with human behavior. Indeed, in the opportunistic networking field, a small but increasing number of attempts have been made to exploit social network features for driving the protocols' design. This looks to be a promising approach, as contacts between nodes are fundamentally tied with users' behavior and hence with social network structures.

However, as of today only some aspects of social networks have been exploited. Studying and modeling human mobility is a research area that has attracted increasing attention. Mobility models based on social behavior represent an important tool for testing the performance of opportunistic systems. Further, a clear understanding of the properties that characterize user movements (such as for any couple of nodes, their contact times, and their intercontact times) provides a cornerstone to design efficient protocols.[5]

A promising direction to fully exploit opportunistic networks involves building the networking solutions around the high-level communication patterns established by the users themselves, rather than applying a legacy engineering-centric approach to bring together devices in a common network plane (layer). An attempt to systematically exploit the underlying social network structure to develop effective social-inspired opportunistic network protocols is currently being carried out inside the Socialnets project (www.social-nets.eu). This project exploits social interactions and user habits to drive the design of protocols for a pervasive system. This is achieved through an interdisciplinary research effort aimed at merging a set of disciplines that are currently running in parallel and without integration—for example, statistical physics studies of complex networks, the social anthropology studies of human behavior, and the computer networking perspective.[6]

Figure 1 summarizes the basic ideas of the Socialnets project. By embedding the social relationships in the electronic world, we can identify at least two levels in an opportunistic environment: an electronic social network (in which relationships depend on the physical properties) and a virtual social network that builds an overlay atop the electronic social network.

Bubble Rap offers a promising forwarding protocol that tries to exploit the electronic social network idea to design effective opportunistic network protocols.[7] Specifically, it focuses on two aspects of a social network: the community and the centrality. Human society is structured in communities, and inside a community some are more popular than others: They have a high centrality. The basic idea of the forwarding algorithm is to use nodes with high centrality to deliver a message to the community that makes clear who the destination node belongs to.

Opportunistic computing can benefit from the ongoing and past research outcomes in pervasive and sensor systems, distributed and fault-tolerant computing, and mobile ad hoc networking.

While electronic social network relationships provide key information for designing opportunistic network protocols, the virtual social network provides a basis for the development of opportunistic computing services. For example, information and services can be replicated and distributed inside the community's electronic social network, taking into consideration its members' interests and locations.[8,9]

OPPORTUNISTIC COMPUTING

Essentially, opportunistic computing can be described as distributed computing with the caveats of intermittent connectivity and delay tolerance. Indeed, mobile and pervasive computing paradigms are also considered natural evolutions of traditional distributed computing. However, in mobile and pervasive computing systems, the disconnection or sleep device situations are treated as aberrations, while in opportunistic computing, opportunistic connectivity leads to accessing essential resources and information.

As Figure 2 shows, opportunistic computing exploits communication opportunities to provide computing services to meet the pervasive application requirements. Opportunistic networking research has benefited from past work in areas such as wireless mobile ad hoc networks and delay-tolerant networks, while pervasive, mobile, and social computing all motivate their respective applications.

Opportunistic computing exploits all available resources in an opportunistic environment to provide a platform for the execution of distributed computing tasks. The major challenge in opportunistic computing is to effectively utilize opportunistic contacts to make information available and accessible and to provide collaborative computing services to applications and users. To make opportunistic computing a reality, middleware services must mask disconnections and delays and manage heterogeneous computing resources, services, and data to provide a uniform view of the system to the applications.

Opportunistic computing can benefit from the ongoing and past research outcomes in pervasive and sensor systems, distributed and fault-tolerant computing, and mobile ad hoc networking. In particular, work in the areas of heterogeneity and interoperability,[10] proactivity and transparency, context-aware computing,[11] location-aware systems, sensor systems, failure handling techniques,[12] and others can be adapted to opportunistic computing systems. However, many challenges to opportunistic computing are unique from those in other systems.

Trusted collaboration

In a disconnected environment, mechanisms for establishing trust among peer nodes play a critical role. Trusted collaboration among the entities of social computing creates opportunities for distributed execution of computing tasks. However, the increasing trend toward decentralization has resulted in significant challenges because traditional security solutions often require centralized online trusted authorities or certificate repositories, which are not well suited for opportunistic networks in which connectivity as well as centralization requirements are both relaxed. Opportunistic networking requires a paradigm shift toward human-centric solutions to establish trust for interactions between peers.

As Figure 2 shows, social links between humans carrying the devices provide strong support for new concepts of trust and security to establish trustworthy relationships among devices and for incentivizing their collaboration at both the network and middleware levels.

Trust, security, and cooperation policies require strict interactions between the two layers, but we also envision several other cross-feed channels. For example, content-based forwarding strategies for routing and forwarding inside the electronic social network layer can exploit context information and content preferences provided by the upper layer.

The collective actions lead to the execution of high-level tasks, as they opportunistically exploit each other's resources. For example, a device that has collected a huge amount of data from the environment can utilize the data compression service offered by another user's device to optimize its memory resources.

More generally, when two nodes meet to perform a collaborative task, they are required to know each other's resources, data, and services. Therefore, when they come in contact opportunistically, they swap information, process the information, and take actions. For example, when two devices—say *a* and *b*—are within the wireless communication range of each other, they have an opportunity to exchange a description of their services. If *b* has a service relevant to *a*, and *b* is a trustable entity, *a* can send the service input parameters to *b*, which then runs the requested service on behalf of *a*. When the service is completed, *b* returns the output parameters to *a*, immediately if they are still in contact, or otherwise the next time they meet.

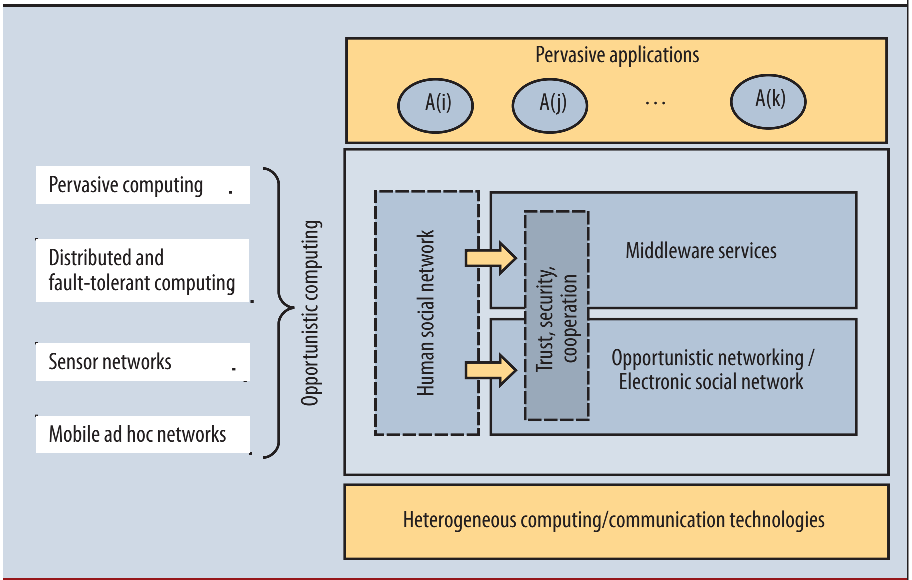

Figure 2. Positioning opportunistic computing, which exploits communication opportunities to provide computing services that meet pervasive application requirements.

Key challenges

Creating a distributed computing platform in a seemingly hostile and unstable networking environment requires overcoming the challenges the opportunistic environment poses.

Intermittent connectivity. In opportunistic networks the contact between pairs of devices provides the critical resource for collaboration. The connectivity problem is exaggerated by the lack of prior knowledge about the location, time, and communication bandwidth of such contacts. Hybrid routing protocols that employ context, a profile, or a history of mobile users and devices should be investigated for adoption into opportunistic networking environments.

It will be necessary to develop middleware mechanisms that mask delays and hide the complexity of the opportunistic paths from applications. The information acquired must be evaluated for caching, purging, and dissemination because resources such as memory and bandwidth are limited.

Delay tolerance. Successful implementations of DTN applications have demonstrated the usefulness of opportunistic networks. Delay tolerance is the key to distributed opportunistic computing.

First, protocols for the creation of delay-tolerant opportunistic communication paths should be developed. Second, delay-tolerant information acquisition and dissemination requires new cache consistency mechanisms to mask delays and the underlying network. Third, execution of remote services in opportunistic environments requires novel mechanisms for service discovery, service execution, and management.

Heterogeneity. Potentially, many kinds of devices will likely come in contact opportunistically—cell phones, handheld and notebook computers, sensors, cameras, and RFID-enabled objects. These devices can be supported by diverse communication capabilities and the radio frequencies at which they communicate. Contact interoperability among these pairs of heterogeneous devices is a major challenge.

RESEARCH ISSUES

Opportunistic networks, combined with social computing, herald the new paradigm of opportunistic computing for pervasive applications. Whereas pervasive computing seeks to enhance user quality of life through proactive application services, opportunistic computing also recognizes and exploits users' social behavior.

User devices, and indeed their BANs/PANs, possess complementary capabilities in terms of computing, communication, storage, energy, sensing, and related applications. This opens several lines of research for developing a set of middleware services that mask disconnections and heterogeneities and provide the applications with uniform access to data and services in a disconnected environment.

Middleware services

Middleware services provide mechanisms for managing information and access through a variety of applications, such as data and services placement, resource management (for example, storage, bandwidth, and energy), trust, security, and privacy for opportunistic computing, mobile agents, remote execution, and cyberforaging, among others. Trust and data privacy pose key issues. For example, reputation mechanisms should be in place to detect malicious users who might join a group to thwart collective actions or acquire sensitive information.

Users' social behavior should be embedded in the middleware mechanism to increase efficiency and security. Novel mechanisms for service sharing in a disconnected environment must be devised. Developing modular tools for message passing, information dissemination and acquisition, resource management, service discovery, service management, and other tasks poses a huge challenge to heterogeneous opportunistic environments.

Fault tolerance is critical to many distributed computing applications. While most current work assumes the existence of local networks, not much research has been done in the prevention, detection, and recovery aspects of faults in challenged environments.

Collaboration in opportunistic environments calls for new, robust strategies that facilitate collaboration in the absence of continuous connectivity. Mechanisms for replication and redundancy must take into consideration the limited resources in such constrained systems. Application tasks executing on one device will be required to interact with resources and services on other devices under time and connectivity constraints.

In an opportunistic computing environment, user devices and sensors carry or supply different kinds of information that is useful to other users and applications.

Information management

In an opportunistic computing environment, special attention must be dedicated to information management and provisioning because a vast range of information is embedded in the environment of pervasive computing and communication systems. However, the use of opportunistic techniques to provide situated information has not yet received much attention, even though there are developments of significant relevance in the distribution of content within P2P systems. Many developers have considered the extension of Internet-based content-sharing systems to mobile ad hoc networks by overlaying the P2P structure. However, very little development has taken place for P2P information provision in opportunistic networking.

The lack of distinction between information producers and consumers on the one hand and the utilization of opportunistic contacts to disseminate and acquire them on the other makes this task challenging. Aggressive broadcasting mechanisms, such as those based on epidemic dissemination protocols, have a tendency to load the network, abusing contact capacity and the content cache. From an information-centric perspective, use of opportunistic networks for information provision results in three fundamental issues: determining what to store, where to store it, and how to acquire relevant information.

Context awareness

Context awareness is a relevant key for searching the network. First, most of the content is relevant for people physically close to the source, who thus form a transient, local community with which to jointly interact. This requires establishing dynamic and temporary trust relationships between humans and machines. In addition, part of the generated content will be of interest to other users in virtual communities, which share common interests irrespective of their physical location. This means that a much wider range of objects can generate and store information while situated in an environment.

Context information and profiles of devices, individuals, and applications—together with cache optimization techniques—are needed for the effective management of the content cache. To share information within a social environment, researchers have proposed *social caches*. A social cache is a logical collective view of individual device caches that cache information objects useful to the members of its social group. Given that members are expected to meet more frequently, and information in the social cache can be effectively utilized by many members, social caching can significantly increase system performance.[8]

Services and data placement and replication

In an opportunistic computing environment, applications need different kinds of resources to execute services, and such resources may be available within the network. Similarly, user devices and sensors carry or supply different kinds of information that is useful to other users and applications. Users' cooperation is tightly coupled with their organization in social communities. Therefore, to increase system efficiency, it is critical to make services and data available in the environment closer to users who need them.

Replication of data or services increases their availability, while their migration may reduce the access delay. While placement of services is a well-investigated problem in traditional distributed systems, dynamic pervasive environments, such as those created by opportunistic contacts, pose new challenges. Andrea Passarella and colleagues[13] have investigated efficient and effective schemes for service provision in opportunistic environments. These investigations will lead to deployment of application-level services.

Resource management

The "product of connection time and available bandwidth," termed *contact capacity*, should be used effectively. The capacity is limited and varies in accordance with wireless communication conditions and the mobility of devices and their users. The contact capacity should be utilized effectively to ascertain reputations, establish trust, collaborate, and exchange information between the two meeting devices and their users.

The second most important resource is the memory or buffer space in the devices themselves. We call this the *content cache*. In opportunistic computing, devices carry one another's information in their content cache, which should be optimally maintained by purging unwanted data and keeping data useful to the applications on the device, such as peers it expects to meet in the near future. The content cache can be tuned to certain applications, contexts, or other criteria.

Energy is another key resource for an opportunistic environment, in which most devices are battery enabled. Energy management is a cross-layer issue with respect to the management of storage and bandwidth. Increased data transmission on the wireless interface results in more energy spent, while local data storage might incur significant energy costs for memory management.

Finally, the hardware and software resources on the devices must be exploited by providing seamless accessibility to applications executing on other devices. As most devices in opportunistic networks are mobile, they possess limited and varied hardware and software resources. Using resources distributed across the devices in a given space, such as a social network, is critical. Matching services to resources in opportunistic networks presents another challenge.

Trust, security, and privacy

Establishing trust and security for an interaction between a priori unknown peers in an opportunistic network is challenging. However, social network structures offer a basis to enhance trust and security provision by capitalizing on "communities" of devices that have commonality between them, either physically or logically.

The idea of using social network structures and properties for enhancing network security is not novel. Indeed, the literature contains several proposals based on using social networks to fight e-mail spam and defend against attacks. However, the use of social networks in completely decentralized networks is a completely new and challenging task because, in such an environment, legacy security solutions based on centralized server or online trusted authorities becomes infeasible. In this case, a natural direction to pursue exploits electronic social networks and the trust and security relationships naturally embedded in human interactions.

Economic model and social cooperation

We might argue that a solid economic model is fundamental to justifying implementation of an opportunistic computing paradigm. Why should one user make computing resources available to another? This is an even more critical question when the computing platforms are mobile devices that have very limited and critical resources, such as energy.

The development of an economic model to stimulate cooperation among peers has been extensively discussed in the framework of both P2P platforms and mobile ad hoc networking, where solutions based on incentives or reputation have been devised. Similar strategies can probably be applied here. However, we believe that exploiting the natural cooperation that exists in human social relations is the catalyst for opportunistic computing. In principle, rational users gain the most from an uncooperative behavior but, despite this, human society often exhibits cooperative behaviors. Characterizing and enforcing human cooperation is highly relevant for electronic social networks.

Mobile agents, remote execution, and cyberforaging

In an opportunistic computing environment, services are often only available on remote nodes outside direct communication of the requesting device. This requires developing mechanisms to support the remote execution of tasks and return the results to the node(s) requesting a service.

Exploiting the natural cooperation that exists in human social relations is the catalyst for opportunistic computing.

Mobile agent technology can be an effective tool to address this issue. Mobile agents may migrate from one node to another during contacts, carry input data and code, and exploit services and resources in the visited nodes. When a task execution is completed, these agents return to the source node together with their results. Similarly, mobile agents can be employed for information acquisition and dissemination.

APPLICATIONS

Opportunistic computing can be the basis for addressing challenges in many application areas. Three critical application areas can benefit from opportunistic computing services and hence constitute driving forces for research in this direction.

Crisis management

Legacy communications networks are not designed to withstand unplanned and unexpected disruptive events and are unsuitable to reliably support communication services for first responders. Opportunistic networking techniques can be adopted for interconnecting surviving parts of the telecommunication infrastructures, and services can be deployed for specific applications.

Infomobility services and intelligent transportation systems

Vehicular ad hoc networks (VANETs) exploit vehicle-to-vehicle communications, as well as the communication with roadside infrastructure, to implement cooperative systems and to increase traffic efficiency and safety. Other applications include tourist information and assistance such as parking availability notification and maps, and entertainment such as gaming and streaming video.

Pervasive healthcare

Opportunistic computing and network technologies can be used to create a pervasive system of intelligent devices comprising sensors and actuators that embrace patient surroundings at different levels. Transparently embedded body area networks and sensors can cooperatively gather, process, and transport information on our lifestyle and the social and environmental context around us without requiring any major change in the users' behavior.

Opportunistic networking techniques can be deployed as basic tools in distributed context-aware pervasive applications for performing real, noninvasive, continuous multiparametric monitoring of physical and physiological parameters.

Distributed computing on opportunistic networking platforms offers a new paradigm in computing, one with tremendous potential. Opportunistic computing will become a reality in the near future, given current growth in the proliferation of powerful and ubiquitous devices and the variety of applications.

A need exists, however, to develop effective solutions to the many new research challenges posed by this intriguing computing opportunity.

Going forward, we require architectures for reliable, secure, and delay-tolerant computing platforms built atop highly dynamic networks and characterized by transient and distributed interactions among devices. This will in turn require great flexibility, scalability, and a general ability to adapt what is typically embedded in human behavior.

Thus, exploiting the properties of human social links to build an electronic social device network offers a promising direction for developing efficient and effective pervasive computing systems that can adapt to highly dynamic and transient distributed systems. C

Acknowledgments

Mohan Kumar's research on opportunistic computing is supported under the US National Science Foundation's Grant CSR 0834493. Marco Conti's research is supported by the European Commission under the Haggle (027918) and Socialnets (217141) FET Projects.

REFERENCES

1. M. Weiser, "The Computer for the 21st Century," *Scientific Am.*, Sept. 1991, pp. 94-104.
2. M. Satyanarayanan, "Pervasive Computing: Vision and Challenges," *IEEE Personal Computing*, Aug. 2001, pp. 10-17.
3. A. Campbell et al., "The Rise of People-Centric Sensing," *IEEE Internet Computing*, July/Aug. 2008, pp. 12-21.
4. L. Pelusi, A. Passarella, and M. Conti, "Opportunistic Networking: Data Forwarding in Disconnected Mobile Ad Hoc Networks," *IEEE Comm. Magazine*, Nov. 2006, pp. 134-141.
5. C. Boldrini, M. Conti, and A. Passarella, "User-Centric Mobility Models for Opportunistic Networking," *Bio-Inspired Computing and Communication*, LNCS 5151, Springer, 2008, pp. 255-267.
6. S.M. Allen et al., "Social Networking for Pervasive Adaptation," *Proc. 2nd Int'l Conf. Self-Adaptive and Self-Organizing Systems Workshops* (SASOW 08), IEEE Press, 2008, pp. 49-54.
7. P. Hui, J. Crowcroft, and E. Yoneki, "BUBBLE Rap: Social-Based Forwarding in Delay Tolerant Networks," *Proc. 9th ACM Int'l Symp. Mobile Ad Hoc Networking and Computing* (MobiHoc 08), ACM Press, 2008; www.cl.cam.ac.uk/~ph315/publications/hoc86309-hui.pdf.
8. C. Boldrini, M. Conti, and A. Passarella, "Design and Performance Evaluation of ContentPlace, a Social Aware Data Dissemination System for Opportunistic Networks," *Computer Networks*, 2009; http://dx.doi.org/10.1016/j.connect.2009.09.001/.
9. P. Costa et al., "Socially-Aware Routing for Publish-Subscribe in Delay-Tolerant Mobile Ad Hoc Networks," *IEEE J. Selected Areas in Communications*, May 2008, pp. 748-760.
10. M. Kumar et al., "Pervasive Information Communities Organization (PICO): A Middleware Framework for Pervasive Computing," *IEEE Pervasive Computing*, July-Sept. 2003, pp. 72-79.
11. K. Henricksen and J. Indulska, "Developing Context-Aware Pervasive Computing Applications: Models and Approach," *Pervasive and Mobile Computing*, Feb. 2006, pp. 37-64.
12. Y. Yu and V.K. Prasanna, "Energy-Balanced Task Allocation for Collaborative Processing in Wireless Sensor Networks," *Mobile Networks and Applications (MONET)*, special issue on algorithmic solutions for wireless, mobile, ad hoc and sensor networks, Feb. 2005, pp. 115-131.
13. A. Passarella et al., *Minimum-Delay Service Provisioning in Opportunistic Networks*, IIT-CNR tech. report 11-2009; http://bruno1.iit.cnr.it/~andrea/docs/TR_11-2009.pdf.

Marco Conti *is a research director at IIT, an institute of the Italian National Research Council (CNR). His research focuses on the design, modeling, and performance evaluation of computer-network architectures and protocols. Contact him at marco.conti@iit.cnr.it.*

Mohan Kumar *is a professor in the Department of Computer Science and Engineering at the University of Texas at Arlington. His research interests include pervasive and mobile computing, opportunistic networking and computing, and sensor systems. Kumar received a PhD in computer science from the Indian Institute of Science. He is a senior member of the IEEE. Contact him at mkumar@uta.edu.*