

**MARIA PAPADOPOULI**  
DEPARTMENT OF COMPUTER SCIENCE  
UNIVERSITY OF CRETE

**RESEARCH AREAS**

My main research is in the area of wireless networks and mobile computing. Together with my collaborators, we have explored the following research domains:

- wireless networking: statistical analysis of measurements and modeling (since 2003)
- mobile peer-to-peer computing (since 1998)
- storage systems for multimedia servers (mostly during 1996-1999)

**PUBLICATIONS**

Some of our papers appeared in the ACM Sigmobile MobiHoc, ACM MobiCom, IEEE Infocom, IEEE Percom, IEEE Globecom, και IEEE Lanman proceedings. The ACM MobiCom, IEEE Percom, ACM Sigmobile MobiHoc, and IEEE Infocom are some of the most prestigious conferences in the research area of mobile computing, pervasive computing, mobile ad hoc networks, wireless networks, respectively. The acceptance ratio of these conferences in the year that our publication appeared is indicated in the following table:

Conference	Acceptance ratio (%)
ACM Sigmobile MobiHoc	16
ACM MobiCom	10
IEEE Infocom	18
IEEE Percom	15

Our journal publications appeared in the *Elsevier Performance Evaluation, Parallel Computing, MC<sup>2</sup>R Mobile Computing and Communications Review, ACM SIGMETRICS Performance Evaluation Review (PER)*.

The distribution of my publications in terms of research areas and type of publication appears in the following tables:

Research Area	Publications
wireless networking: statistical analysis & modeling	C4, C6, C7, C9-C11, C13-18, J1, J2
mobile peer-to-peer computing	C1, C3, C8, C12, C19-C26, J3
storage systems for multimedia servers	C27, J4, J5

Publications	Numbers
Books	1 (*)
Conference Proceedings	27
Journal	4

(\*) The monograph is currently in production by Springer (due to appear in November 2008).

**CITATIONS** (as of 1/11/2008)

My top five publications, in terms of number of citations in scholar.google.com, are the following: **C20** (166 citations), **C18** (67 citations), **C24** (63 citations), **C21** (50 citations), **C25** (36 citations).

My publications had the following distinctions: **C12** was ranked best paper in the review process and **C21** received the *best student paper/poster award in the Advances in Wired & Wireless Communications, IEEE Sarnoff Symposium 2001*.

This research statement will summarize our main research results in the areas of wireless networking and mobile peer-to-peer computing. Additional information about related research, professional, and educational activities, awards, and grants can be found in my curriculum vitae.

## 1. WIRELESS NETWORKING

We envision highly flexible environments of heterogeneous wireless networks of devices with different capabilities that have access to multiple channels or network interfaces and roam over different networks. In general, such environments are extremely complex and the interaction of different layers and technologies creates many situations that cannot be foreseen during the design and testing stages of technology development. This is especially true for wireless networks, which are used for many different purposes, and which are based on a shared medium that is inherently more vulnerable than its wired counterpart. It is critical to understand the *performance and workload of wireless networks* and develop wireless networks that are more robust, easier to manage and scale, and more able to efficiently utilize their scarce resources. While in several cases over-provisioning in wired networks is acceptable, it can become problematic in the wireless domain. It is therefore important to perform *comprehensive empirical studies in a wide range of production environments* to uncover deficiencies and identify possible optimizations and extensions.

To better understand the characteristics of wireless access and workload and uncover the weaknesses of wireless networks, we performed measurement-based studies to examine the spatial locality of the wireless information, user access and roaming patterns, network performance, and traffic workload characteristics. For this purpose, extensive real-traces were acquired from various large-scale campus-wide IEEE802.11-based infrastructure and were statistically analyzed.

The evolution of a technology is characterized by a sequence of steps that span from preliminary evaluation of this technology in small-scale testbeds to its deployment and adoption in potentially large-scale and real-life environments. It is common practice for a preliminary technology evaluation to explore its behavior under well-understood conditions and simple models. For more comprehensive performance analysis, it is necessary to use realistic and sophisticated models for the parameters of that technology. *Generating empirically-based models that capture the characteristics of the user traffic and access and providing a flexible framework for using them in performance analysis studies* were two other objectives of this research. We were among the first to statistically analyze and model the traffic and access demand using extensive empirical traces collected from large-scale IEEE802.11-based networks.

We distinguished the following important dimensions in wireless network modeling:

- (i) user traffic demand
- (ii) access patterns
- (iii) network topology
- (iv) channel conditions

This distinction enabled the superimposition of workload models on a given topology and focus on the right level of detail. We focused on access patterns and user traffic workload.

Our main contributions related to this research are as follows:

- A statistical analysis of the traffic workload of wireless infrastructures in various spatio-temporal scales, e.g., per Access Point (AP), per client, and infrastructure-wide
- A methodology for analyzing the performance of large-scale production wireless LANs in terms of delay, packet losses, and interrupted connections
- An application-based characterization of wireless demand in multiple spatio-temporal scales (per AP, per client, and infrastructure-wide)
- An analysis of spatial locality of wireless workload and impact of various caching mechanisms in wireless access
- A methodology and algorithms for forecasting the wireless traffic load at APs
- A methodology for statistical modeling of wireless traffic demand and roaming patterns
- An evaluation of the proposed wireless user demand models and highlighting the tradeoffs between scalability, accuracy, robustness, and tractability
- A data repository of standardized real-life and synthetic wireless traces, models, and monitoring tools (*FORTH Repository*).

The next paragraphs describe our research results in more detail.

## **1.1 STATISTICAL ANALYSIS OF WIRELESS NETWORK MEASUREMENTS**

### **Comparative analysis of wireless traffic load at APs [C15, B1]**

In **C15**, we analyzed the wireless traffic load of two major campus-wide wireless infrastructures and found a surprising degree of similarity in their characteristics. These results therefore provide strong evidence in support of the development of parsimonious workload models of campus wireless networks. Specifically, the analysis reveals the following interesting phenomena:

- There is a wide range of workloads and log normality is prevalent. In general, the traffic load in both wireless infrastructures is light, although there are long tails.
- No clear dependency with the type of building at which the AP is located exists, although some stochastic ordering is present in the tail of the distributions.
- An interesting dichotomy among APs is prominent in both of the infrastructures: APs dominated by uploaders and APs dominated by downloaders.
- The number of non-unicast wireless packets is substantial. Furthermore, the number of unicast received packets is strongly correlated in the log-log scale with the number of unicast sent packets. While the majority of APs send and receive packets of relatively small size, a significant number of APs show rather asymmetric packet sizes, i.e., APs with large sent and small receive packets, and APs with small sent and large receive packets.
- The distribution of the associations and roaming operations was found to be quite heavy-tailed and there is a correlation between the traffic load and number of associations in the log-log scale.

### **Application-based characterization of wireless traffic [C6]**

While statistics and measurements regarding the overall usage of individual APs are readily available, the limited knowledge of the wireless traffic demand, in terms of the type of application, hinders efficient network provisioning. In **C6**, we performed an extensive application-based characterization of traces collected from a large-scale wireless network, going beyond the port-number limitation, across three levels, namely, network, clients, and APs. Several interesting observations were made from this analysis: The most popular applications are web browsing and peer-to-peer accounting approximately for 81% of the total traffic. Most users are also dominated by these two applications. Network management and scanning activity are responsible for 17% of the total flows. While building-aggregated traffic application usage patterns appear similar, the application cross-section varies within APs of the same building. Most wireless clients appear to use the wireless network for one specific application that dominates their traffic share. File transfer flows, such as ftp and peer-to-peer, are heavier in the wired network than in the wireless one. The traffic share across applications is significantly affected when clients associate with new APs. This appears to be independent of the specific application type. Finally, there is an interesting dichotomy among APs, in terms of their dominant application type and downloading and uploading behavior.

#### **Performance of large-scale wireless networks [C13, B1]**

In **C13**, we explored the performance of a large-scale wireless network and compared it with a wired one. **C13** uses only passive monitoring techniques, namely, inference from TCP header traces. This enabled us to study a population of several thousand hosts in a real production environment, in which more than 31 million TCP connections were made. Our first contribution is methodological. Passive methods always have some degree of uncertainty, and we overcome this limitation by mostly relying on relative differences between wired and wireless traffic. Our analysis revealed that wireless clients experienced substantially higher packet delay variability than wired clients but their loss rates are surprisingly similar. We found that both the number of unnecessary TCP retransmissions and, even more substantially, the number of interrupted connections are higher for the wireless LAN than for the wired LAN. To the best of our knowledge, this is the first research effort to directly contrast wired and wireless traffic of a large production network.

#### **Caching paradigms in wireless access [C18, B1]**

Temporal and spatial phenomena and the impact of caching in a large-scale wireless infrastructure were examined in **C18**. Each client frequently requests objects that it has requested within the past hour, and occasionally requests objects that had been requested by other nearby users within the past hour. In general, same-AP caching is beneficial for APs with high hit ratios; such APs were found in the wireless infrastructure of our study. As in the case of wired networks, the single-client locality is a primary factor in wireless data. The overall ideal hit ratios of user cache, cache attached to an AP, and peer-to-peer caching (where peers are co-resident within an AP) paradigms in the wireless infrastructure of our study are 51%, 55%, and 23%, respectively. A cache at each AP would achieve an ideal hit ratio of 55% for the entire trace. On the other hand, a cache that serves the entire campus would achieve an ideal hit ratio equal to 71%. Thus, there is an opportunity to improve wireless access by more actively caching data in a user cache.

## 1.2 STATISTICAL-BASED MODELING OF WIRELESS TRAFFIC & ACCESS

A client initially disconnected from the Internet may associate with an AP in its wireless range. During its visit to that AP, this client may generate traffic by sending and/or receiving packets. Later, the client may reassociate with another AP and prolong its wireless Internet access, or disconnect from the wireless infrastructure. A transition is marked by two consecutive connections to distinct APs. A *session* of a client is a sequence of associations of that client to APs. A session is used to capture an episode of a *continuous* wireless access to the infrastructure; it begins when a currently disconnected client associates with the infrastructure and ends when the next disconnection message is received. During a session, the client generates flows by receiving and sending packets. Sessions capture the interaction between the clients and the network, while flows model the above-packet-level traffic activity masking the underlying network dependencies. We have modelled *the wireless traffic workload and access patterns* using *session-* and *flow-*based parameters.

### Access & roaming patterns [C14, C17,C18, B1]

The wireless access of a client is modelled as an alternation between sessions and disconnections. An access pattern of a client is characterized by an arrival process at certain APs and a sequence of transitions between APs. Various parameters can be used to characterize the mobility or roaming activity of a client, such as the *duration of sessions and visits, transitions between APs, number of inter-building transitions, duration of time spent and frequency of visits at a certain AP, duration of disconnection, predictability of the next AP associations, and arrival process at an AP*. We have statistically analyzed and modeled these parameters.

Predicting client arrivals at APs can improve the buffering, caching, load balancing, and prefetching at APs in order to mask the end-to-end delay, particularly in the case of regular clients. APs may not only predict client arrivals but also traffic demand. Based on these predictions, neighbouring APs can advise newly-arrived clients to avoid hotspots, suggest alternative APs, and better balance their load and channel utilization. **C17** analyzes rich set of empirical traces collected from large-scale campus-wide wireless infrastructures and proposes a new set of metrics to describe a session and characterize its mobility, transient nature, and spatial properties. It shows that the majority of the sessions last less than one hour. Wireless clients exhibited relatively low mobility, spending a large percentage of their wireless life at the same AP. In general, mobile sessions tend to have a small percentage of long visits and a large percentage of short visits at APs. A family of BiPareto distributions can model the visit and session durations.

**C18** shows that Markov-chain models can be used to characterize transitions of clients between APs and *accurately predict the next AP* with which a client will associate. These predictions can be further enhanced by incorporating networking and physical topological data as well as temporal information, such as time, day of the week, and visit duration. Unlike **C18** that models the roaming pattern of a client from the perspective of the local wireless device, **C14** focuses on the perspective of an AP and models the *client arrival at APs as Time-varying Poisson processes*. The main contributions of **C14** include a novel methodology for modeling the arrival processes of clients at APs and the use of a powerful visualization tool for finding detailed interior features and quantile plots with

simulation envelope for goodness-of-fit test. Time-varying Poisson processes can model well the arrival processes of clients at APs. We validated these results by modeling the visit arrivals at different time intervals and APs. Furthermore, we proposed a clustering of the APs based on their visit arrival and functionality of the area in which these APs are located.

### **Wireless user workload [C7, C9, J1, B1]**

In **C9**, we proposed a modelling methodology and models for the wireless user workload. The session arrival process provides the seeds of a cluster process, in which the arrivals of sessions imply the arrivals of correlated sets of flows. The following parameters are modelled:

- *session arrivals*
- *number of flows within a session*
- *flow inter-arrival times within a session*
- *flow sizes*

For each parameter to be modelled, the distribution that fits the best our empirical traces is selected. Several distributions were considered, such as the Pareto, Lognormal, Poisson, Exponential, BiPareto, and Generalized Extreme value. Maximum likelihood was used for the parameter fitting while the evaluation of the goodness of the various distributions was performed using formal and visual statistical analysis methods and tools, such as the quantile plots with simulation envelopes. The following distributions model the user demand well:

- a *Time-varying Poisson process* models well the session arrivals at various APs in the infrastructure
- the *BiPareto* models well the flow size and number of flows within a session
- the *Lognormal* is a great candidate for the flow inter-arrivals within a session

Extending **C9**, in **C7**, we *validated* these models using *various spatial scales* (e.g., AP-level, network-wide, groups of APs located at the same building) and *time periods* and found that *the above distributions persist*. At each spatio-temporal scale, the models for sessions and flows remain the same with only their parameter values differing. In **C7**, we showed how by selecting the appropriate spatio-temporal granularities of the models, the right balance between reusability and accuracy can be addressed. For example, when hourly periods and AP-scale are used, the models maintain sufficient spatial detail at the cost of a lower scalability and amenability. When a network-wide scale is used, simplicity is gained at the cost of a higher loss of detail. This evaluation of the models was performed using statistics-based metrics. Moreover, when we focus on hourly time periods of different traffic conditions—very heavy traffic and normal traffic workload—at APs, the above distributions still persist.

When the statistics-based metrics indicate a deviation of the models from the empirical data, systems-based criteria can be employed to evaluate the impact of this difference on the performance of a given system. In **J1** and (chapter 5 of) **B1**, we employed systems-based criteria, such as the hourly aggregate throughput, per-flow delay and throughput, and goodput, to evaluate the performance of a hotspot AP. For this evaluation, we

generated synthetic traces based on various models and spatio-temporal scales. Emulation- and simulation-based scenarios were performed using synthetic and real traces—that is, traces generated from a real-life wireless infrastructure—as input for the user workload. These traces were then “replayed” in emulation or simulation testbeds. An important result is that synthetic traces based on our models result in a performance *very close* to the one when the real traces are used as input. Furthermore, synthetic traces based on popular models—employed frequently in simulations—exhibit *large deviations* from the real traces.

### **Forecasting of wireless traffic demand [C11, C16, C4, J2]**

Admission control and load-balancing algorithms can benefit from short-term forecasting of traffic load at hotspots. An AP may predict its traffic load for the next time interval (e.g., next hour or five minute interval) and use their traffic load forecasts during admission control to not only better manage their traffic demand but also advice clients to associate with the appropriate APs to better utilize their local resources.

In **C11** and **C16**, we designed a number of new forecasting algorithms, such as the weighted moving-average, adaptive moving-average and flow-based algorithms. The flow-based algorithms exploits the high correlation in the log-log scale between the traffic load and number of flows in hourly and five-minute intervals. We also applied a normalized ARIMA multi-step ahead time-series forecasting algorithm. These algorithms were evaluated on real traffic traces acquired from the most heavily utilized APs from a large-scale wireless infrastructure. It appears that the time granularity and recent traffic history have dominant impact on the prediction accuracy. Specifically, the finer the time granularity and more recent the historical traffic data is, the larger their impact on the prediction error.

In **C4** and **J2**, we employed Singular Spectrum Analysis and found that the time-series of traffic load at a given wireless AP has a small intrinsic dimension. In particular, these time-series can be accurately modelled using a small number of leading (principal) components. The statistical analysis of leading components has demonstrated that even a few first components form the main part of the information. The residual components capture the small irregular variations, which do not fit in the basic part of the network traffic and can be interpreted as a stochastic noise. Based on these properties, we also studied contributions of the various components to the overall structure of the traffic load of an AP and its variation over time. Finally, we designed and evaluated the performance of a traffic predictor for the trend component, obtained by projecting the original time-series on the set of leading components.

## 2. PEER-TO-PEER COMPUTING FOR MOBILE NETWORKS

In this research, our attention shifts to wireless networks that are sparser and frequently disconnected from the Internet. In such networks, a device is *not always connected to the Internet, nor within wireless range of another device*. To enhance the formation of on-line communities of mobile users and attempt to create a new socio-technological paradigm, we proposed 7DS, a novel architecture and set of protocols that enables resource sharing among peers that are not necessarily connected to the Internet. 7DS encompasses three facets of cooperation: data sharing, message relaying, and bandwidth sharing. 7DS may relay, search for and disseminate information, and share bandwidth. It operates in a self-organizing manner, without the need for an infrastructure and serves as the underlying information and service discovery protocol. 7DS runs in the middleware and 7DS-enabled devices communicate with each other via wireless LANs. Unlike typical peer-to-peer approaches in wired networks, 7DS does not try to establish permanent caching or service discovery mechanisms due to the highly dynamic environment. Instead, 7DS hosts acquire the data from other peers within their wireless coverage using single-hop broadcast. Although IEEE802.11 APs and clients are rapidly deployed, there are still areas of limited or no wireless coverage. 7DS can “bridge” the access via wireless infrastructures and peers through caching and relaying. In that manner, it enables user devices to act as “short-term” service/access providers to other peer devices. We were among the first to study the resource sharing in such mobile networking environments.

Our main contributions related to this research are as follows:

- The design and implementation of 7DS, a novel system that enables information dissemination and sharing among mobile hosts
- The evaluation of the impact of the wireless range, host density, querying mechanism, power conservation, and cooperation on data dissemination via extensive simulations
- The generation of theoretical models for data dissemination that use random walks and diffusion-controlled processes
- The design, implementation, and evaluation of CLS, a location-sensing system that employs the peer-to-peer paradigm to enhance the position estimates
- The design, implementation, and evaluation of various mobile peer-to-peer computing and context-aware applications
- The design and evaluation of an on-demand channel switching for multi-channel wireless MAC protocols

The next paragraphs describe our research results in more detail.

### **Architecture of a mobile peer-to-peer computing system [C23, C24, C25, D1, B1]**

In **C23** and **C24**, we presented 7DS, a novel architecture and set of protocols that enable resource sharing among peers that are not necessarily connected to the Internet. **C25** focuses on bandwidth sharing and analyzes its performance.

**Performance analysis of information discovery and dissemination in mobile networks [C20, C21, C22, D1, B1]**

7DS can instantiate a different data access mechanism using either the server-to-client or the peer-to-peer paradigm. The thrust of this research is the information dissemination in mobile networks, which raises several questions: How fast does information spread in such networks? What is the impact of cooperation, data popularity, number of devices, and wireless range on information diffusion? How do the different mobile information access and caching paradigms compare?

Given the dearth of large-scale, non-controlled 7DS-like environments, we run extensive simulations to assess the performance of information dissemination and study the peer-to-peer and server-to-client paradigms. In addition, the impact of the transmission range, density of peers, popularity of information on the performance of information discovery and dissemination in relatively *short time periods* was studied. The percentage of hosts that acquire the data object as a function of time and their average delay were measured. **C20, C22, D1**, and Chapter 3 of **B1** present the performance analysis results. We found that the density of the cooperative hosts, their mobility, and the transmission power have the most pronounced impact on data dissemination. The synchronization of the periods that the network interface of peers is powered and the reduction in the frequency of querying can save energy. In the case of the fixed server-to-client paradigm, with a low density of hosts, the query frequency can be set as large as three minutes without impacting on the speed of data dissemination. Similar results hold in the case of the peer-to-peer paradigm.

The performance of data dissemination remains the same when the area is expanded but the density of the cooperative hosts and the transmission power are kept fixed. Also, for a fixed total wireless coverage, the larger the density of cooperative hosts, the better the performance. In the fixed server-to-client paradigm, this implies that for the same wireless coverage density, it is more efficient to have a larger number of cooperative hosts with lower transmission power than fewer with a higher transmission power.

In **C20**, the emphasis is on the *short-term behaviour* of the information dissemination. To analyze the *long-term* behaviour of the data dissemination, a simple epidemic model can be used to compute the expected delay for a message to be propagated to the population of an area. We also proposed a novel theoretical framework based on *diffusion-controlled processes, random walks and kinetics of diffusion-controlled chemical processes* to model the stationary server-to-client scheme (**Chapter 3** in **B1** and **D1**). For that, we used the *Rosenstock's simple trapping model* and found that the analytical results are consistent with the simulations.

Unlike the above studies that consider a mobile network with devices that are frequently disconnected from an infrastructure, **C21** considers a wide-area-network with devices mostly connected to that infrastructure. To access services within their local environment, these devices participate in a service discovery protocol which involves a master directory that registers all services available in the local environment. These directories typically are isolated from each other. Devices that move across service discovery domains have no access to information outside their current local domain. In **C21**, we

proposed an application-level protocol that enables data sharing among discovery domains. Each directory maintains a table of active links to other directories that share related information. A set of linked directories forms a data cluster that can be queried by devices for information. The data cluster is distributed, self-organizing, responsive to data mobility, and robust to failures. Using application-defined data schemas, clusters organize themselves into a hierarchy for efficient querying and network resource usage. Through analysis and simulation, we evaluated this mechanism under different workloads and showed that the protocol overhead for both maintaining a cluster and handling failures grows slowly with the number of gateways.

### **Design, implementation & evaluation of mobile p2p applications [C1, J3, B1]**

Chapter 2 in **B1** presents a number of *novel mobile computing applications* that use the 7DS as their underlying information discovery and sharing mechanism. One of them is the multimedia travelling journal application that enables users to build interactive multimedia journals that associate multimedia files with locations on maps. It runs on top of 7DS, and through 7DS, it allows local peers to share files associated with certain locations. The multimedia files and maps are stored in the cache of the local 7DS instance. A user can add or query for multimedia files related to a certain region on the map by clicking on the map. Moreover, the user can add, modify, or delete comments on a certain multimedia file, change its permission, and rate its content. A multimedia file can be public or private, and only public files are shared with other peers. The architecture of the system has been presented in **B1**, **C1**, and **J3**. A comparative performance analysis of the peer-to-peer- and centralized-based implementations of the multimedia travelling journal can be found in **C1**.

### **Collaborative location-sensing (CLS) [C3, C19, B1]**

In **C3** and **C19**, we applied the peer-to-peer paradigm to positioning for mobile computing devices. Our proposed system, CLS, positions wirelessly-enabled devices using the existing wireless communication infrastructure adaptively without the need of specialized hardware or training. To improve its accuracy, CLS enables hosts to cooperate and share positioning information and also allows the integration of external information, such as maps, popular routes, and user mobility patterns. CLS applies the peer-to-peer paradigm by enabling devices to gather positioning information from other neighbouring peers, estimate their distance from their peers based on signal-strength measurements, and position themselves accordingly. Periodically, CLS can refine its positioning estimates by incorporating newly received information from other devices. CLS adopts a grid-based representation of the physical space; each cell of the grid corresponds to a physical position in the physical space. The cell size reflects the spatial granularity/scale. Each cell of the grid is associated with a value that indicates the likelihood that the node is in that cell. These values are computed iteratively using one of the following approaches:

- A simple voting algorithm, through which a local CLS instance casts votes on cells of the grid. A vote on a cell indicates the likelihood that the local device is located in the corresponding area of that cell.
- A particle filter-based model.

CLS can incorporate additional information to improve its location estimates. Examples of such information are: position estimates from different network interfaces (e.g., Bluetooth, RF tags, IEEE802.11), contextual semantics (e.g., topological information about the environment, mobility patterns, hotspots of the area), and signal-strength-based signatures of the physical space, to improve the location estimation. Several CLS variants have been implemented and evaluated via extensive simulations and empirical measurements. For the empirical evaluation, we ran experiments using IEEE802.11 signal-strength measurements in various environments, such as FORTH and Crete Aquarium. CLS can utilize signal-strength maps of the physical space by superimposing statistical properties of the signal-strength values acquired during the training phase on their corresponding positions. Such maps can significantly improve its performance. Through empirical experiments, we showed how the different statistical properties of signal-strength measurements, the particle-filters model, the AP failures and additional peers and APs affect the performance of CLS.

#### **On-demand channel switching for multi-channel wireless MAC protocols [C12]**

In **C12**, we proposed a novel broadcast-based medium access control (MAC) protocol for *ad-hoc* wireless networks with *multiple channels* at each host. It enables devices to perform an on-demand, dynamic, channel selection based on the traffic conditions of the channels and communication pattern of each host. A host stays on a channel as long as its traffic share on that channel is above a certain threshold, below which it switches to another channel. It aims on reducing the “unnecessary received” traffic and channel switches, while still being able to send and receive legitimate traffic. It can enhance the bandwidth utilization, reduce packet delays, and increase energy savings, without requiring any specialized hardware.

#### **Framework for supporting context-aware services [C2, C5, C8]**

The desire to build a framework for supporting such context-aware services has been the driving force behind **C2**, **C5** and **C8**. Specifically, these papers proposed a context-aware platform for mobile devices in dynamic environments which uses Semantic Web technologies to model context information and advanced interactive map-based interfaces for accommodating pedestrians. To test and demonstrate the approach, a prototype has been developed and a number of further extensions have been studied. The design and implementation of this system have enabled us to identify and address the main challenges for realizing the objectives of next generation semantics-based pervasive information systems.