C2P2: A Peer-to-Peer Network for On-Demand Automobile Information Services

Shahram Ghandeharizadeh
Computer Science Department
University of Southern California
Los Angeles, CA 90089, USA
shahram@usc.edu

Bhaskar Krishnamachari
Computer Science and Electrical Engineering
University of Southern California
Los Angeles, CA 90089, USA
bkrishna@usc.edu

Abstract

This short paper outlines challenges of delivering continuous media and traffic information to mobile Car-to-Car Peer-to-Peer (C2P2) network of devices. We analyze network connectivity of a C2P2 cloud as a function of radio range of each device. A novel concept introduced by C2P2 is on-demand delivery of continuous media, audio and video clips, to moving vehicles.

1 Introduction

Automobile manufacturers along with numerous science fiction movies envision smart vehicles that navigate their occupants to a target destination using the fastest possible path while entertaining the passengers (e.g., those children who repeatedly ask "Are we there yet?"). One may find elements of this vision in today's automobiles. For example, the current high-end navigation systems can provide turn-by-turn instructions to direct a driver to a destination. Many manufacturers of minivans offer their customers with the option to purchase an entertainment system that might include a fold-down screen, video game console, wireless headphones, a DVD player, etc.

A peer-to-peer network technology can enhance today's automobile capabilities to provide additional information services. This technology will depend crucially upon the availability of low-cost devices equipped with a powerful processor, abundant amount of storage, and a wireless communication devices[1]. As we show in Section 2, when a fraction of cars are equipped with such devices, they can form an effective information network. Such a mobile peer-to-peer network is capable of providing timely, personalized, *on demand* traffic information to the driver. The information can be used by an on-board navigation system to compute the fastest path from origin to destination using current road status (say, no more than tens of seconds old). The network may also provide other multimedia (audio/video-on-demand) and traditional information services to the passengers (email, web browsing), possibly in conjunction with a cellular-based approach.

For traffic information, each device in the network is both a consumer and a producer of data. It produces data as follows. First, it is equipped with a GPS and a database of the surface streets[2], e.g., Los Angeles basin. A freeway (or a surface road) is partitioned into segments. Each segment has a well defined start and end points, speed profiles, etc. A car traversing a segment might measure its speed or the number of cars it observes in that segment. If it observes either a lower speed than expected average or a larger number of cars than usual in that stretch, it produces a *belief* that this segment is congested.

The information flow within the network may be determined by either a pull or a push-based mechanism. For example, with a push paradigm, the belief/report generated by each car is broadcasted to the other vehicles in its range which repeat this information to other cars in order to propagate it through the ad-hoc network to a vehicle that might consider passing through this segment. Data recency is an important consideration that may be managed by attaching lifetime fields to reports (say 1 minute), which force reports to be deleted when exceeded. If the same report is independently generated by a different device, the lifetime may be doubled at each receiving device. In a pull paradigm, the reports (tagged with geographical identifiers) would be returned only if specifically queried for by a node in the network.

Each car in the network may choose to make route planning decisions either once – at the beginning of its trip – or dynamically, as new information about traffic conditions

[1] Intel, for example, offers small devices for approximately $85 that consist of a 500 MHz processor with a wireless device that operates in the 5 GHz spectrum. The radio bandwidth of the device is on the order of Megabits per second (Mbps) with a radius of tens of meters [17].

[2] It might utilize the GPS of an existing navigation system.

becomes available [3]. It may treat the map as an edge-weighted graph, with road stretches as edges and intersections as vertices. The traffic information obtained from the network is used to mark each edge with a cost that corresponds to expected travel time through the corresponding road stretch. In the absence of specific reports, a default travel time may be associated with each road stretch. The navigation software on the car can then solve a shortest-path problem on the graph in order to determine the least delay route from the current location to the destination indicated by the driver. This effectively turns the transportation network into a controlled network, and one question to be investigated is if this might result in oscillatory effects. We believe that such effects can be mitigated by data gathering strategies that filter out momentary changes in traffic. Indeed, we argue that the widespread use of such on-demand traffic information systems would result in "congestion-control" on the transportation network, through effective load balancing.

We term this network of mobile devices car-to-car, peer-to-peer network, or C2P2 for short. C2P2 shares many similarities with existing peer-to-peer networks. One may categorize existing peer-to-peer networks into those based on either an ad-hoc or an addressable placement of data. With an addressable placement of data, e.g., CAN [13], Pastry [14], OceanStore [2] etc., a request for a referenced data item is routed to a specific destination node containing this data item. With an ad-hoc placement, e.g., Gnutella [5], a request floods the network and may search many nodes for the referenced data item. Addressable networks are further categorized into those that utilize either a centralized directory, e.g., Napster, or a de-centralized addressing space such as a partitioned space to identify a target node containing the referenced data item.

Similar to Napster, some of the current efforts envision utilizing centralized base stations for communication with mobile devices. A C2P2 device communicates its information with a centralized base station that makes decisions on the road status. Similarly, when a C2P2 device is queried for a route to a specific destination, it requests the centralized base station to provide it with status of specific road stretches for route planning. In the absence of centralized base stations (remote areas), the devices may communicate with one another and render localized beliefs on the status of different road stretches. In this scenario, the propagation of information is similar to how Gnutella processes queries: the devices may broadcast beliefs amongst themselves in the hope that a device might be interested in a broadcasted belief. A device inquiring about the current status of a specific stretch of a road is aware of its own current geographical location and the target road's geographical location [9, 8]. Using the surface roads as a guide, the system knows how to route the information to communicate with a device that recently traversed that segment of the road. This dynamic routing mechanism shares similarities with an addressable peer-to-peer routing, e.g., with CAN's routing mechanisms, given a starting CAN node and a destination zone, CAN routes a request with the objective to get close to the destination. C2P2's routing may take advantage of the structure of roads in combination with GPS readings of different routing devices to guide a request to its destination.

If we wish to extend the use of C2P2 networks for the streaming of continuous media from a base station for entertainment, there are additional challenges. First, the size of a referenced clip might exceed the storage capacity of a C2P2 device. Second, once the device starts the display of a stream, it should display the rest of it free from disruptions and delays (hiccups). A novel concept here is for C2P2 devices that are within the same geographical proximity and moving in the same direction to form clusters[3]. While a single device may not be able to store the entire repository, a cluster of devices might be able to store a significant fraction of this repository. The system may also utilize hypothetical clusters for prefetching data in order to deliver the right segment of data at the right time to a consuming C2P2 device. Intelligent caching and prefetching of data are also approaches that will minimize latency for applications such as e-mail, web browsing, chat-rooms etc.

In the following section, we provide some of our preliminary research results on how to determine connectivity between two C2P2 devices when their placement is constrained by linear structures such as roads. This analysis includes a study of what percentage of vehicles must be equipped with C2P2 devices to provide connectivity. Next, Section 3 elaborates on several future research directions that will shape our activities.

2 Ad-hoc Network Connectivity

A key difference between the C2P2 network and traditional peer-to-peer networks is that its topology changes with time because of the mobility of component nodes. Several routing algorithms are designed to move data from one node to another in such mobile ad-hoc network. These are categorized into those that employ geographic location information such as LAR [9] and GPSR [8], and those that do not use this information such as DSR [7], AODV [12], and ZRP [6]. However, all these algorithms assume fundamentally that the underlying network is indeed connected and a path exists between any two nodes.

Only recently has the question of the network density required to ensure connectivity been studied in any detail. Results by Gupta and Kumar [4], Xue and Kumar [16], and

[3]The system may even hypothesize about clusters expected to appear sometime in the future using the destination information of different C2P2 devices. These are termed hypothetical clusters.

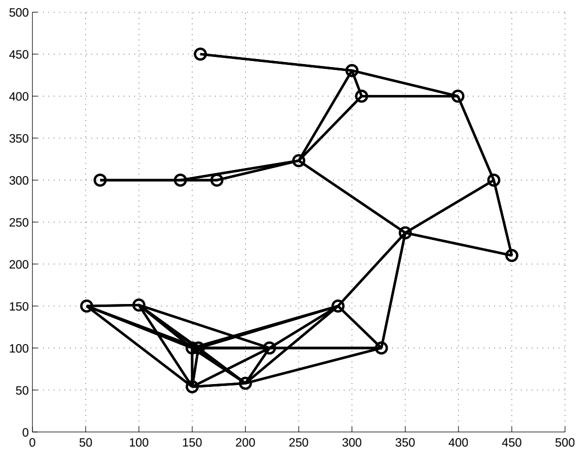

Figure 1. Screen capture of the simulator showing movement of various C2P2-equipped vehicles on a Manhattan-like grid of streets, along with the instantaneous network topology.

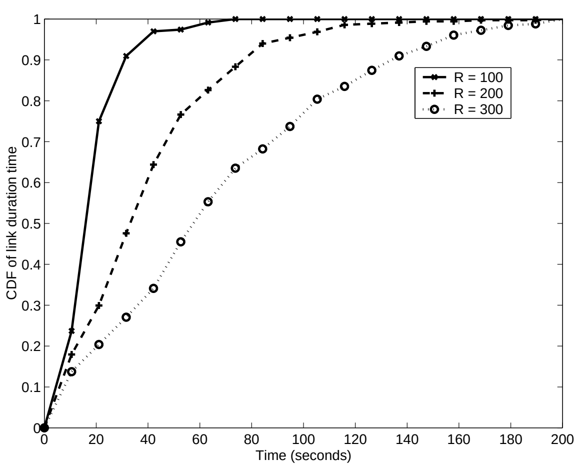

Figure 2. Cumulative probability density function for link duration times with different radio ranges.

Krishnamachari [10] etc., have shown that there is a "phase transition" with respect to the radio range. There is a critical value transmission range such that below this, nearly all randomly generated networks are disconnected, and above it all randomly generated networks are connected with a high probability. They have concluded that each node must be connected on average to $\log n$ nodes to ensure that the entire network is connected. With our target application, these studies have two limitations: First, they are focused on static, random ensembles, permitting nodes across the entire space. It is not clear how these results might hold for mobile vehicles that are constrained by surface roads. Second, they focus on the connectivity of ALL nodes in a network instead of application specific issues that might be less strict. For example, it may suffice to determine if a given car can obtain information from ANY C2P2 device near a region of interest through the network (regardless of its global connectivity).

We used simulation studies to investigate network connectivity and data availability in a C2P2 network. These studies assume a Manhattan-like model consisting of a regular rectangular city grid of roads. Each block is 50m by 50m. We consider a square region of 10x10 blocks. A given number of cars n move through the roads in this area. For each vehicle, a random velocity is chosen at each simulation time interval (once every .25 seconds) between 0 and 50 kmph (approximately 30 mph). In addition, each vehicle also pauses once during the entire simulation for a short while (a parameter typically set to 5 seconds) at a random time instance. The cars travel through a random set of intermediate waypoints in a randomly chosen grid. Between any two waypoints, the cars take (with equal probability) any of the several possible shortest paths. Each entire simulation is discretized into .25 second intervals, and is typically run for 400 seconds.

Our obtained results are shown in Figures 1-4. Figure 1 shows a snapshot of the simulator showing the vehicles in the grid forming a C2P2 network. The C2P2 network topology is highly dynamic, with relatively short link durations that increase with the radio range of individual devices, see Figure 2. Figure 3 shows the connectivity of C2P2 mobile networks. It also shows the critical threshold behavior associated with randomly located nodes: for a given value of n, there is a critical radio range beyond which the network is almost always connected. This threshold becomes smaller and the transition becomes sharper as we increase the number of nodes in an area (from 20 to 50).

Figure 4 shows both the connectivity and data availability in the C2P2 network. In the simulated scenario, one of the nodes in the network is requesting traffic information about a particular intersection, from ANY node in the network that is near (within 100m of) this location. This figure thus illustrates how even when the complete global connectivity is as low as 25%, more than 75% of the application-level needs of the C2P2 network may be fulfilled. This suggests that what may be important is not so much global connectivity of all nodes, but data availability within the network.

In general, we find that for radio ranges on the order of 100-200 meters, 20-50 nodes are needed per 10 km of roadway. For a typical traffic density of 1 car per 10 m, this translates to a penetration level of 2 - 5 % for effective performance of a C2P2 network.

3 Future Research Directions

Automobile industry defines telematics as an automobile network system that uses wireless communications, GPS tracking and call centers to give drivers better access to information and services. A cloud of C2P2 devices with cell phone base stations provides the infrastructure for telematics. Concepts similar to C2P2 have been described in sensor networks [1] and mobile information management [15]. The most relevant study is CarNet [11]. These studies re-enforce the vision of a qualitatively new class of mobile peer-to-peer networks. A novel concept introduced by C2P2 is on-demand delivery of continuous media, audio and video clips, to moving vehicles.

C2P2 systems present several challenges along two inter-related fronts: networking and data management. While there has been some initial work on routing protocols in mobile ad-hoc networks in recent years, a number of analytical issues remain unresolved. Scalability, bandwidth, and latency concerns are paramount and need to be quantified to answer questions such as: "How many traffic reports can an individual vehicle request?", "Under what conditions can the network guarantee a data latency of less than x seconds per hop?", etc.

There are also a number of significant open problems on how to manage information in such networks. It will be important to develop efficient and secure mechanisms for in-network storage, querying, and data processing. Particularly for the data-specific traffic information system, it may be necessary to develop data-centric routing mechanisms (the flow of named data from producers to consumers) in addition to address-centric routing (end-to-end communication of arbitrary data between addressable nodes).

Our efforts are also directed towards ensuring the interoperability of C2P2 networks with centralized infrastructures such as cellular and other static base stations. Since wireless communication is typically more expensive in terms of bandwidth and can suffer from harsh channel conditions, it may be necessary to employ the notion of "driving data to the ground as quickly as possible." It will be important to design C2P2 architectures that allow for a flexible compromise between centralized, decentralized,

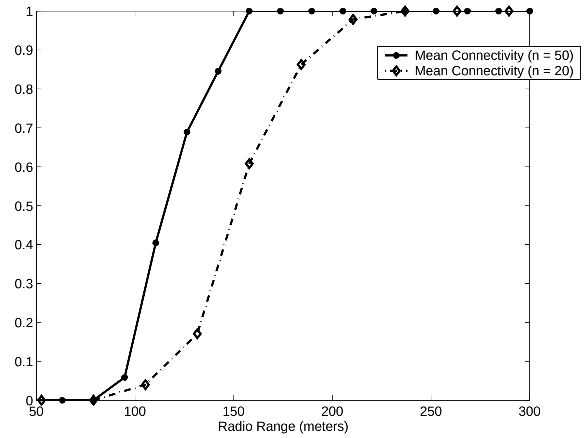

Figure 3. Probability that the entire network is connected with respect to radio range for $n = 20, n = 50$.

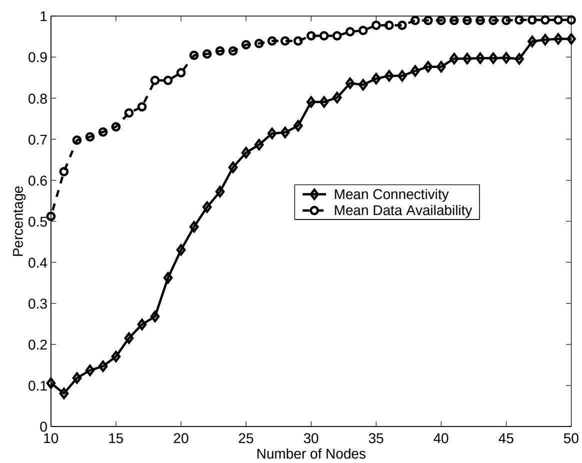

Figure 4. Mean connectivity and mean data availability with respect to the density (number of nodes/10kms) for a radio range of 150m

and hybrid approaches.

4 Acknowledgments

This research was supported in part by NSF grant IIS-0307908, National Library of Medicine LM07061-01, and an unrestricted cash gift from Microsoft Research.

References

[1] D. Estrin, R. Govindan, J. Heidemann, and S. Kumar. Next Century Challenges: Scalable Coordination in Sensor Networks. In *Proceedings of the Fifth Annual International Conference on Mobile Computing and Networks (MobiCOM'99)*, August 1999.

[2] J. Kubiatowicz et al. OceanStore: An Architecture for Global-scale Persistent Storage. In *Proceedings of ACM ASPLOS*. ACM, November 2000.

[3] S. Ghandeharizadeh, T. Helmi, S. Kapadia, and B. Krishnamachari. A Case for Mobility Based Admission Control Policy. In *Proceedings of the International Conference on Distributed Multimedia Systems*, September 2004.

[4] P. Gupta and P. Kumar. Critical Power for Asymptotic Connectivity in Wireless Networks, 1998.

[5] J. Guterman. Gnutella to the Rescue? Not so Fast, Napster Friends. *Gnutella Forum*, September, 2001.

[6] Zygmunt J. Haas and Marc R. Pearlman. The performance of query control schemes for the zone routing protocol. In *ACM SIGCOMM*, 1998.

[7] D. B. Johnson and D. A. Maltz. Dynamic Source Routing in Ad Hoc Wireless Networks. In Imielinski and Korth, editors, *Mobile Computing*, volume 353. Kluwer Academic Publishers, 1996.

[8] B. Karp and H. T. Kung. GPSR: Greedy Perimeter Stateless Routing for Wireless Networks. In *ACM Mobicom*, 2000.

[9] Y. Ko and N. H. Vaidya. Location-Aided Routing (LAR) in Mobile Ad Hoc Networks. In *Mobile Computing and Networking*, pages 66–75, 1998.

[10] B. Krishnamachari, S. Wicker, R. Bejar, and M. Pearlman. Critical Density Thresholds in Distributed Wireless Networks. To appear in *Communications, Information and Network Security*, Kluwer Publishers, 2002.

[11] R. Morris, J. Jannoti, F. Kaashoek, J. Li, and D. Decouto. CarNet: A Scalable Ad Hoc Wireless Network System. In *Proceedings of SIGOPS European Workshop*, September 2000.

[12] C. Perkins. Ad Hoc On Demand Distance Vector (AODV) Routing, <draft-ietf-manet-aodv-00.txt>, 1997.

[13] S. Ratnasamy, P. Francis, M. Handley, R. Karp, and S. Shenker. A Scalable Content Addressable Network. In *ACM SIGCOMM*, 2001.

[14] A. Rowstron and P. Druschel. Pastry: Scalable, decentralized object location, and routing for large-scale peer-to-peer systems. *Proceedings of the Eighteenth Conference on Distributed Systems Platforms*, 2218, 2001.

[15] O. Wolfson. Moving Objects Information Management: The Database Challenge . In *Proceedings of the Fifth Workshop on Next Generation Information Technologies and Systems (NGITS 2002)*, June 2002.

[16] F. Xue and P. R. Kumar. The Number of Neighbors Needed for Connectivity of Wireless Networks. In submission, 2002.

[17] R. Yavatkar. USC Distinguished Lecture Series: Internet Exchange Architecture. Private communications, October 2002.