

A Case for Leveraging 802.11p for Direct Phone-to-Phone Communications

Pilsoon Choi¹, Jason Gao¹, Nadesh Ramanathan², Mengda Mao², Shipeng Xu²,
Chirn-Chye Boon², Suhaib A. Fahmy², Li-Shiuan Peh¹

¹Massachusetts Institute of Technology, USA

²Nanyang Technological University, Singapore
pilsoon@mit.edu

ABSTRACT

WiFi cannot effectively handle the demands of device-to-device communication between phones, due to insufficient range and poor reliability. We make the case for using IEEE 802.11p DSRC instead, which has been adopted for vehicle-to-vehicle communications, providing lower latency and longer range. We demonstrate a prototype motivated by a novel fabrication process that deposits both III-V and CMOS devices on the same die. In our system prototype, the designed RF front-end is interfaced with a baseband processor on an FPGA, connected to Android phones. It consumes 0.02uJ/bit across 100m assuming free space. Application-level power control dramatically reduces power consumption by 47-56%.

1. INTRODUCTION

Direct device-to-device (D2D) communication between smartphones has been available for years via WiFi's ad-hoc mode, but as operating system modifications are required to set this up, D2D has not been widespread until recently, with the adoption of the WiFi Direct standard. Video sharing, file sharing, as well as multiplayer games have started to leverage WiFi Direct. Clearly, there are peer-to-peer applications that benefit from the faster response times of D2D communications; these applications gather user input and sensor data from nearby phones, perform computations in-situ, and output results and user interface updates with higher responsiveness. However, existing D2D communication only works for short-range, low mobility scenarios. WiFi is challenged in long-range or high-mobility scenarios [1]. WiFi Direct facilitates easier setup of device-to-device networks, but one device must serve as an access point (the group owner) and all other devices must communicate through it, thus not supporting highly mobile networks with rapidly changing topologies. This largely limits WiFi Direct applications to close-range, static deployments between a few phones. LTE Direct can be regarded as a promising new D2D technology, but as it leverages LTE infrastructure, it requires modifications to the LTE base stations which may hinder adoption.

Vehicle-to-vehicle (V2V) communication is, in essence, a form of D2D communication, and has been burgeoning with the adoption of the IEEE 802.11p DSRC standard around the world [2, 3]. Nu-




Figure 1: RF front-end modules (FEMs) on the Apple iPhone 4.

merous V2V applications in the transportation domain have been proposed or deployed, such as mobile multimedia, safety, road pricing, and others [4]. These applications leverage the high mobility, long range and fast response times of 802.11p for next-generation transportation applications. 802.11p's increased transmit power enables longer range communications, but the high power consumption of 802.11p radios has, until now, precluded their integration into non-vehicular mobile devices¹.

In this paper, we demonstrate the feasibility of realizing 802.11p on phones by bringing together materials, devices, circuits, and systems researchers. We see this development opening up D2D communications to a much larger class of applications, with mobile devices on pedestrians, passengers, and drivers now interconnected at low latency and high bandwidth, enabling highly interactive mobile applications.

Among several building blocks for a communications system, the RF front-end is one of the most critical, with III-V semiconductor devices (e.g. GaN, GaAs, InGaP) showing much better power density and efficiency than CMOS. Figure 1 (photo from [5]) shows multiple RF front-end modules (FEMs) for a variety of standards in an Apple iPhone 4; together, these occupy a large portion of real estate. In addition, each FEM includes multiple semiconductor dies within it, further increasing area footprint, power, as well as cost.

In our work, we leverage a unique process, the LEES (Low Energy Electronics Systems) process, where both CMOS and III-V semiconductor devices can be fabricated on a *single* die. This allows the use of the most suitable III-V devices grown on top of a conventional CMOS device, interfaced via metal layers. Such single-die integration offers the superior performance required by 802.11p specifications at the small form factor and within the tight

¹The recently-released Qualcomm Snapdragon Automotive Solutions support DSRC for short-range vehicular safety detection, but not 802.11p and its extended range with high transmit power.

power budget of a smartphone implementation. In Section 2, we will show how the process can shrink the 802.11p front-end module and how the FEM can plug into the existing communications subsystem circuitry on a phone.

LEES devices and the relevant semiconductor processes are now being developed in conjunction with a commercial CMOS foundry, targeting to release the first prototype devices and circuits at the end of 2014. In parallel, device modeling and p-cell layout for the LEES devices are also in progress to develop a PDK, which will be the first integrated CMOS and III-V design kit for circuit designers to create innovative circuits using a conventional design flow. To demonstrate chip functionality before the LEES process is ready, we first design and fabricate a reference front-end circuit for our system prototype using standard commercial 0.18 μ m CMOS and 0.25 μ m GaN technologies on separate dies. This front-end incorporates a novel circuit design to realize the high transmit power (28.8dBm, $4\times$ or $19\times$ that of WiFi) required by 802.11p, at low power.

We demonstrate compatibility with existing phones by emulating an 802.11p baseband on FPGA (using a modified 802.11a baseband) and interfacing the FPGA with the fabricated 802.11p transmitter. Our system prototype is a transmitter chain consisting of the designed front-end circuits in standard CMOS and GaN technologies, a baseband processor in an FPGA board interfaced to an Android smartphone through USB, all 802.11p compliant. Application-level adaptive control of the radio's transmit power through a gain control interface means the Android application can tune the radio's transmit power (and thus its power dissipation) to match actual desired D2D communication distance. This joint hardware-software power optimization enables substantial further power reduction, allowing the prototype to meet the aggressive smartphone power budget.

2. BACKGROUND

2.1 LEES Process and Design Flow

By enabling monolithic integration of III-V materials with CMOS, the new process presents new challenges to integrated circuit design that are fundamentally different. It thus prompted the setting up of a materials-circuit-system team to explore new application drivers that can best leverage the monolithic, vertically integrated process.

The LEES process is based on conventional front-end silicon CMOS processing by a commercial foundry, followed by III-V integration and processing in a research lab, before returning the processed wafer back to the commercial foundry for back-end silicon CMOS processing. Figure 2 illustrates an example structure of III-V monolithically integrated with CMOS/Si devices. In this paper, we focus on GaN HEMTs as these are particularly suitable for the high-power RF circuits necessary for 802.11p.

As the LEES process is based on a commercial CMOS foundry, CMOS circuit design can leverage the existing CMOS PDK provided by the foundry. The III-V portion of the die, however, will require a new PDK which includes III-V device models, layout p-cells, and interconnect models between III-V and CMOS devices. The PDK has been developed with a physics-based compact model of III-V devices [6] (i.e. GaN HEMT first) coupled with device layout, sizing and spacing rules defined by device and process researchers, enabling CMOS+III-V circuit simulation using conventional CAD tools and layout for both the CMOS and GaN portions of the die. This integrated CAD flow enables joint CMOS+III-V circuit design and eases migration of the LEES process to commercialization.

2.2 Phone Communications Circuits

A typical smartphone incorporates several two-way communications radios, including WiFi (IEEE 802.11a/b/g/n/ac), Bluetooth,




Figure 2: LEES process integration of III-V (GaN) and CMOS: (a) A silicon-on-insulator (SOI) wafer with fabricated Si devices; (b) Si CMOS/GaN-on-Si wafer realized by two-step bonding technology; (c) GaN window open and device isolation; (d) Schematic cross-section view of the monolithically integrated GaN HEMT devices with final metal interconnection of fabricated HEMTs and Si CMOS devices.

and the cellular radios. The cellular radios in mobile phones available today do not support direct device-to-device (D2D) communications, and only communicate with the cellular base stations that coordinate access to the medium. WiFi Direct is a recent standard that allows D2D communications between mobile phones, and thus enables networks with star topologies, but not mesh or full peer-to-peer topologies. Ad-hoc WiFi is a pre-existing standard that allows for direct D2D communication without needing to appoint one of the devices as a centralized controller or access point, but is not widely supported among the major mobile operating systems, and thus requires kernel modifications.

Each radio typically contains a PHY (physical layer) and MAC (medium access control) implemented in hardware, with upper MAC and higher networking layers implemented in software at the device driver, operating system and application level. Most of the building blocks of the communications subsystem within a phone are increasingly being integrated with current standard CMOS processes, except for the RF power amplifier. While there is significant ongoing circuit research targeting CMOS power amplifiers to enable higher level of integration of the entire communications subsystem, the intrinsic low power density and efficiency of current CMOS devices presents a tough challenge [7]. As shown in Figure 1, a power amplifier for each communication standard is still a separate chip fabricated using III-V technology which enables higher output power and efficiency, but worsens system form factor.

2.3 802.11p Compatibility with 802.11a

IEEE 802.11p DSRC is an emerging standard originally proposed for vehicle-to-vehicle (V2V) and vehicle-to-infrastructure (V2I) communication, enabling truly distributed mesh D2D networking like ad-hoc WiFi. Table 1 compares the 802.11p specification with 802.11a. 802.11p adopts the same OFDM modulation as 802.11a, but its time domain parameters are double those of 802.11a to mitigate highly mobile and severe fading vehicular environments. Thus, when we implement the digital baseband processor for 802.11p, we can use 802.11a hardware as is, but run it at half the clock frequency.

However, with the increased transmit power and robustness necessary for longer range V2V communications, the high power con-

	802.11p (DSRC)	802.11a (WiFi)
User mobility	Vehicular (outdoor)	Personal (indoor)
Operating frequencies	5.85-5.925GHz	5.15-5.825GHz
Channel bandwidth	10MHz	20MHz
Max. output power	760mW	40/200mW
Data rate	3-27Mbps	6-54Mbps
Modulation	BPSK-64QAM	BPSK-64QAM
OFDM symbol duration	8us	4us
Guard time	1.6us	0.8us
Preamble duration	32us	16us
Subcarrier spacing	156kHz	312kHz

Table 1: 802.11p vs. 802.11a specifications.




Figure 3: Proposed interfacing of the 802.11p LEES single-die solution with existing phones' WiFi chipset.

sumption of 802.11p radios has precluded their use in mobile phones until now. The LEES process can resolve this issue through its novel CMOS/III-V integration technology that can optimize high power density III-V devices' performance for specific applications and integrate them with CMOS on a single die. This new device and process technology, combined with a novel circuit design, along with adaptive gain control by the application software, make it possible to implement a low-power and small form-factor 802.11p-based D2D solution in a smartphone. Figure 3 depicts how the 802.11p system can be implemented with the existing WiFi chipset and application processor on a phone using LEES technology, such that only a 802.11p RF front-end chip needs to be added.

3. SYSTEM PROTOTYPE DESIGN

Our system prototype is a transmitter chain from phone to FPGA to the RF front-end, enabling 802.11p compliant signal transmission with application level gain control for power saving. A USB-Ethernet adapter is used to communicate between the Android phone and the FPGA. To interface the baseband processor on FPGA to our custom RF front-end, commercial DAC evaluation boards are used to feed analog I/Q signals into the RF transmitter. As mentioned in Section 2.2, all these digital and ADC/DAC components can be shared with existing WiFi communications circuitry, and a single RF front-end can readily support both 802.11a and 802.11p by slightly extending its maximum carrier frequency range from 5.875GHz to 5.925GHz. Thus, to demonstrate the feasibility of 802.11p implementation and its compatibility with existing WiFi solutions, we first design and fabricate a CMOS transmitter and GaN PA using commercial foundries, then leverage existing 802.11a IP to implement 802.11p baseband on an FPGA, before interfacing the FPGA to the Android kernel.

3.1 RF Front-End

Figure 4 shows the circuit schematics of the CMOS+GaN RF front-end consisting of a CMOS transmitter and a GaN PA. A simple CMOS receiver circuitry is also designed for calibration to enhance




Figure 4: CMOS+GaN RF front-end block diagram




Figure 5: Fabricated die microphotographs: (a) CMOS transmitter circuit (b) GaN PA circuit.

the RF performance. A CMOS driver amplifier is designed using cross-coupled neutralization capacitors to improve linearity, along with the center-tapped on-chip inductor for source degeneration. A center-tapped on-chip transformer is used to remove an external balun. The CMOS circuits are fabricated in 0.18um CMOS technology and occupy $1.4mm^2$ including a whole transmitter chain and a receiver for I/Q mismatch calibration excluding the pad area, as shown in Figure 5(a).

A GaN PA is designed and fabricated using a commercial 0.25um GaN-on-SiC process [8]. Since GaN HEMTs are depletion mode devices with a negative pinch-off voltage, the PA requires external negative biases. To satisfy the EVM requirement at 28.8dBm maximum output power, we propose a new in-phase power combining technique utilizing both Class AB and C biasing devices, achieving both linearity and efficiency across output power levels, while eliminating complex design issues in traditional power amplifiers. Figure 5(b) shows the fabricated GaN circuit occupying $1.28mm^2$ die area.

3.2 FPGA Subsystem

FPGAs provide an ideal platform for prototyping complex radio baseband implementations in real-time, offering high performance, low power, and portability, in comparison with other software radio platforms [9]. The FPGA platform performs two vital functions in our setup: baseband processing and providing the interface between the application software on the Android phone and the analog/RF circuitry via the DAC. The FPGA system is implemented on a Xilinx XC5VLX110T FPGA on the XUPV5 development board. We have implemented the complete transmitter chain as described in Figure 6. Packets are transmitted from the Android handset, via the RRR Abstraction layer and Ethernet physical interface, into the 802.11p Airblue baseband on the FPGA. The resultant baseband output is passed through a digital low pass filter and scaled before delivery to the DAC circuitry.

The Airblue baseband [10] was originally designed for the 802.11a standard. Since the two standards are largely similar, we run the entire baseband design at half the clock frequency (10MHz) to achieve compatibility with the 802.11p standard. It is worth noting that for




Figure 6: FPGA system diagram.

actual 802.11p deployment, more stringent output spectrum shaping is required than for 802.11a [11]. The Android handset can access two functions in the FPGA hardware: the packet generator and the gain control module. The packet generator is responsible for configuring parameters, buffering, and synchronizing, the baseband transmission. The gain control module allows the Android handset to directly control power settings on the RF front-end. The FPGA receives and decodes power control commands from handset, applying the appropriate settings at the front-end via a parallel pin interface. This enables power saving capability to be applied from the Android application software. We add a digital low pass filter to reduce the noise caused by the sampling effect within the 40MHz spectrum range.

The Asim Architect’s Workbench (AWB) [12] is the development environment for hybrid hardware-software design. FPGA support is provided in AWB via the Logic-based Environment for Application Programming (LEAP) framework [13] that provides the Remote Request-Response (RRR) framework, an abstracted communication layer.

3.3 Phone-FPGA Interface

The Android smartphone is interfaced to the FPGA through a USB-Ethernet adapter connected via Ethernet to the FPGA and via USB On-the-Go (OTG) to the Android device.

In order to have the Android device recognize and enumerate the USB-Ethernet adapter, we recompiled the Linux kernel for the device to include the USB-Ethernet drivers for the particular ASIX AX88178 and SMSC 7500 chipsets in the adapters. We then loaded this kernel onto the phones, replacing the default kernel. This allowed the Android device to become a USB host and recognize the USB slave Ethernet adapters attached to it via the USB OTG cable.

4. EVALUATION

4.1 Circuit Measurements

For the CMOS transmitter, after the calibration process, -52dB single-sideband (SSB) rejection is achieved as depicted in Figure 7(a). With the same calibration settings and using the IEEE 802.11p baseband I/Q signals, -36.5dB EVM is achieved at the output. The GaN PA achieves -30.5dB EVM and 22% drain efficiency across one-decade output power ranges with its maximum output power of 28.8dBm. This circuit characteristic is suitable for system level power saving across all output power levels at high efficiency, unlike a conventional PA whose efficiency exponentially decreases as output power drops. Figure 7(b) shows the output power spectrum using the 802.11p 64-QAM signals.

4.2 System Prototype Evaluation

Figure 8 shows the experimental setup of the system prototype, illustrating that a smartphone, an FPGA board and commercial DAC evaluation boards are interfaced to the designed CMOS and GaN




Figure 7: RF performance: (a) SSB rejection of CMOS transmitter, (b) GaN PA output spectrum.




Figure 8: Snapshot of the system prototype.

PCB boards. An 802.11p compliant digital baseband implemented in the FPGA along with the TI dual 12-bit DAC, DAC2902, sampling at 40MHz, feeds the analog I/Q baseband signals into the CMOS transmitter. An Android application on the smartphone controls packet generation/transmission and RF gain.

Since the transmit mode dominates power consumption, we design and implement an entire transmitter chain to validate the LEES feasibility as well as potential power reduction through application-level adaptive power control (ALAPC). In addition, the DC power of the PA is more than 90% of the whole transmitter power with a complex modulation scheme like OFDM in 802.11p, since it requires back-off due to its high PAPR signals and the power efficiency is dramatically reduced as output power decreases from the saturation point. Thus, power management of the PA is crucial to fit the 802.11p front-end within a smartphone’s stringent power budget.

In the following subsections, we demonstrate that ALAPC, combined with our GaN PA’s improved power efficiency across all output power levels, can achieve dramatic power reductions. We cannot yet deploy our prototype system due to its complex system configuration, FPGA and DAC boards, and multiple power supplies for the transmitter and PA boards. However, we can use traces from prior deployments of two mobile apps which originally used off-the-shelf D2D communications, to estimate the potential system power savings that can be achieved if we replace those COTS D2D radios with our proposed single-die 802.11p radios integrated within phones.

4.2.1 RoadRunner Evaluation

RoadRunner [14] is an in-vehicle Android app for road congestion control, and speaks turn-by-turn navigation instructions to the driver, like existing navigation systems, while enforcing road-space rationing by allocating tokens among vehicles in the background. Tokens permit a vehicle to drive on a specific road segment, and are distributed to vehicles from a server over the cellular network (LTE), or exchanged directly between vehicles over 802.11p DSRC.

Original deployment. The original deployment took place in Cambridge, MA, USA, consisting of 10 vehicles driving among




Figure 9: Average power consumption of RoadRunner V2V exchanges with and without adaptive power control.

multiple possible congestion-controlled routes. Three different scenarios were evaluated: RoadRunner using only the cellular network as a baseline; additionally using ad-hoc WiFi for V2V communications; and additionally using 802.11p DSRC for V2V communications. With 802.11p, each smartphone was tethered via USB to an off-the-shelf 802.11p DSRC radio [15]. Using 802.11p enabled network response time improvements of up to 80% versus the cellular network, and cellular network usage reductions of up to 84%. Ad-hoc WiFi's performance did not suffice: with ad-hoc WiFi, only 5 V2V communications sessions occurred at an average distance of 29.2 meters, resulting in only 6.8% of requests being offloaded to V2V from the cellular network, while with 802.11p, 47 V2V sessions occurred at an average distance of 175.7 meters, offloading 43% of requests. This original deployment thus motivates the use of 802.11p as a mobile D2D communication standard for phones, while the cumbersome setup tethering a COTS 802.11p radio to a phone motivates a single-die 802.11p chip.

Adaptive power control. We obtained the RoadRunner traces and assume that with our adaptive power control, each V2V communications session (a token exchange) would be transmitted at the minimum power required to reach the other vehicle. We compare this to the original deployment traces as a baseline, in which every V2V token exchange is conducted at full radio power. The traces include vehicle location, communications on all radio interfaces, and distances at which V2V token exchanges occurred during the deployment. For each V2V exchange, we look up the minimum power level to transmit a packet across that distance from our experimental measurements of the GaN PA, using 64-QAM coding. We normalize the sum of these estimates to a situation with no adaptive power control, shown in Figure 9.

With ALAPC and our new PA design (22.5% efficient for all power levels), the V2V exchanges use 47% less power (from 3.37 W down to 1.77 W), indicating that many V2V communications sessions did not need the full transmit power in the original deployment to reach the other vehicle. With ALAPC, but without our new PA design (so efficiency is exponentially decreasing), V2V token exchanges use 4.8% less power than the baseline (from 3.37 W down to 3.21 W), underscoring the importance of the improved PA efficiency of our circuits in realizing gains from ALAPC.

4.2.2 SignalGuru Evaluation

SignalGuru [16] is a vehicular traffic light detection iPhone app that shares data among multiple phones to collaboratively learn traffic signal transition patterns and provide GLOSA (Green Light Optimal Speed Advisory) to drivers. Each vehicle contains a windshield-mounted iPhone that observes traffic signal transitions via the phone's camera and broadcasts the observations over ad-hoc WiFi every 2 seconds.

Original deployment. The original SignalGuru deployment also occurred in Cambridge, MA, USA, along three consecutive inter-




Figure 10: Average power consumption of SignalGuru broadcasts with and without adaptive power control.

sections on Massachusetts Avenue. 5 vehicles followed a route for 3 hours, generating GPS location traces. To surmount the limited range of ad-hoc WiFi, a phone stationed near an intersection acted as a relay.

Adaptive power control. We obtained the SignalGuru traces, and in the simulation of our proposed 802.11p radio's performance, whenever a vehicle broadcasts a packet (every 2 seconds), we calculate the power level required to reach the nearest vehicle to it, from 19.8 to 28.8 dBm. We compare this to baseline static power control, in which every broadcast is transmitted at the maximum power level of 28.8 dBm.

With ALAPC and our new PA design (22.5% efficient for all power levels), SignalGuru broadcasts use 56.3% less power (from 3.37 W down to 1.47 W), shown in Figure 10. With ALAPC, but without our new PA design (efficiency exponentially decreasing), SignalGuru broadcasts use 24.5% less power than the baseline (from 3.37 W down to 2.54 W), highlighting again that the improved power efficiency of our single-die circuits is important to substantially lowering overall system power consumption.

4.2.3 Power Reduction Summary

To put our power reductions of 1.6 W (RoadRunner) and 1.9 W (SignalGuru) in context, we measured the dynamic range of a Samsung Galaxy S4 smartphone's power consumption to be between 1 W (screen on, idle) and 11 W (running a CPU-intensive benchmark) using a Monsoon Power Monitor [17]. This indicates a significant power reduction in the overall platform power budget can be realized with our new power amplifier.

4.3 Simulations on the new PDK

To predict the circuit functionality and layout area of the combined CMOS + GaN design using LEES process technology, the LEES PDK is used for the simulation and layout of our prototype circuits. Since the PDK includes an unmodified commercial CMOS PDK, we validate only the III-V portion. At one GaN device's three terminals, we achieve the same waveforms at 5.9GHz, which means RF parasitics as well as intrinsic AC/DC parameters are successfully reflected in the model. We also model large-signal nonlinearity using the physics-based compact model to match the PA's nonlinear characteristics at high output power, which is essential for predicting accurate performance of a III-V (or III-V + CMOS) PA with CMOS power control circuitry on a single die.

Figure 12 depicts a draft layout for the designed CMOS transmitter combined with a GaN PA using the p-cells in the PDK. Compared with the sum of two separate die areas at $1.4 + 1.28 = 2.68\text{mm}^2$, the combined CMOS + GaN layout using our PDK totals 1.98mm^2 , demonstrating the proposed process integration can further shrink the form factor (as pads are no longer needed). Along with the smaller single die area, the technology can integrate all FEMs and




Figure 11: Simulation comparison between a commercial and LEES GaN models.




Figure 12: Layout for the designed CMOS transmitter and GaN PA using the integrated PDK.

related components currently in a smartphone into a single die (or package) in the near future.

5. CONCLUSIONS

This work is the result of collaboration between materials and device researchers, circuits designers and mobile systems and software architects. Motivated by the novel GaN-CMOS monolithic process, we leveraged the GaN HEMT devices to realize the high-power power amplifier necessary for 802.11p specifications, and coupled that with a CMOS transmitter. The RF front-end circuits were tailored for adaptive power control, targeting good power efficiency across a wide range of transmit power. An 802.11p baseband processor was emulated on an FPGA (using an existing 802.11a baseband) to connect an Android phone to the RF front-end, creating a full system prototype to demonstrate the feasibility of incorporating the RF front-end into a phone. Our results show that the GaN-CMOS process can realize an 802.11p front-end within the stringent power and area budgets of a smartphone.

6. ACKNOWLEDGEMENTS

We thank Zhihong Liu for detailed discussions on the LEES process, Zhaomin Zhu for the development of the LEES PDK, in SMART, Ujwal Radhakrishna for the compact modeling of the GaN HEMTs, and Kermin Fleming for helping with porting of the AirBlue baseband, at MIT. This research was supported by the Singapore National Research Foundation through the Singapore-MIT Alliance for Research and Technology (SMART)'s Low Energy Electronic Systems (LEES) and Future Urban Mobility (FM) research programmes, and by the United States Department of Defense NDSEG fellowship.

7. REFERENCES

- [1] A. Balasubramanian, R. Mahajan, A. Venkataramani, B. N. Levine, and J. Zahorjan, "Interactive Wifi connectivity for moving vehicles," in *Proc. ACM SIGCOMM*, 2008.
- [2] "The connected vehicle test bed," www.its.dot.gov/factsheets/connected_vehicle_testbed_factsheet.htm, US Department of Transportation, 2013.
- [3] "CAR 2 CAR Communication Consortium," car-to-car.org.
- [4] P. Papadimitratos, A. La Fortelle, K. Evenssen, R. Brignolo, and S. Cosenza, "Vehicular communication systems: Enabling technologies, applications, and future outlook on intelligent transportation," *IEEE Communications Magazine*, vol. 47, no. 11, pp. 84–95, 2009.
- [5] "iFixit iPhone 4 teardown," ifixit.com/Teardown/iPhone+4/3130, iFixit, 2010.
- [6] U. Radhakrishna, L. Wei, D.-S. Lee, T. Palacios, and D. Antoniadis, "Physics-based GaN HEMT transport and charge model: Experimental verification and performance projection," in *Proc. IEEE IEDM*, 2012.
- [7] G. Liu, P. Haldi, T. K. Liu, and A. M. Niknejad, "Fully integrated CMOS power amplifier with efficiency enhancement at power back-off," *IEEE Journal of Solid-State Circuit*, vol. 3, no. 43, pp. 433–435, March 2008.
- [8] P. Choi, C. Boon, M. Mao, and H. Liu, "28.8 dBm, high efficiency, linear GaN power amplifier with in-phase power combining for IEEE 802.11p applications," *IEEE Microwave and Wireless Components Letters*, vol. 23, no. 8, pp. 433–435, August 2013.
- [9] J. Lotze, S. A. Fahmy, J. Noguera, B. Ozgul, L. Doyle, and R. Esser, "Development framework for implementing FPGA-based cognitive network nodes," in *IEEE Global Communications Conference – GLOBECOM*, 2009.
- [10] M. C. Ng, K. E. Fleming, M. Vutukuru, S. Gross, and H. Balakrishnan, "Airblue: A system for cross-layer wireless protocol development," in *Proc. ACM/IEEE Symposium on Architectures for Networking and Communications Systems*, 2010.
- [11] T. H. Pham, I. V. McLoughlin, and S. A. Fahmy, "Shaping spectral leakage for IEEE 802.11p vehicular communications," in *Proc. of the IEEE Vehicular Technology Conference – Spring*, 2014.
- [12] J. Emer, P. Ahuja, E. Borch, A. Klauser, C.-K. Luk, S. Manne, S. Mukherjee, H. Patil, S. Wallace, N. Binkert, R. Espasa, and T. Juan, "Asim: a performance model framework," *Computer*, vol. 35, no. 2, pp. 68–76, 2002.
- [13] A. Parashar, M. Adler, K. Fleming, M. Pellauer, and J. Emer, "LEAP: A virtual platform architecture for FPGAs," in *Proc. of Workshop on the Intersections of Computer Architecture and Reconfigurable Logic*, 2010.
- [14] J. Gao and L. Peh, "RoadRunner: Infrastructure-less vehicular congestion control," in *ITS World Congress*, in press 2014.
- [15] "MK2 WAVE-DSRC Radio," cohdaireless.com/product/mk2.html.
- [16] E. Koukoumidis, L.-S. Peh, and M. R. Martonos, "SignalGuru: leveraging mobile phones for collaborative traffic signal schedule advisory," in *Proc. ACM MobiSys*, 2011.
- [17] Monsoon Solutions, "Monsoon Power Monitor," <http://msoon.com/LabEquipment/PowerMonitor/>.