Prototyping Wearables: A Code-First Approach to the Design of Embedded Systems

Daniel Graham and Gang Zhou, Senior Member, IEEE

Abstract-As wearable devices become ubiquitous, there will be an increased demand for platforms that allow engineers and researchers to quickly prototype and evaluate new wearable devices. However, many of these platforms require that the hardware be configured before the code is written, thereby limiting the programmer to the limitations of the hardware. In this paper, we present a platform that allows researchers and engineers to quickly prototype new wearable devices using a code-first approach. This approach allows software developers to create new prototypes by first writing the code that the prototype is required to run. Once the code has been written, the hardware that is required to run the application can be generated by analyzing the code that the software developer has specified. This code-first approach is possible because of the system's architecture which is comprised of both a hardware and software component. The hardware component consists of a main board with four expansion ports, while the software platform is a modular middleware which consists of a collection of stateless libraries that abstract each hardware module. These modular abstractions allow us to synthesize the hardware configuration from the software definition. We evaluated our design using it to prototype three wearable devices: 1) an environmental exposure monitoring smartwatch; 2) an infrared indoor localization system; and 3) a step counter.

Index Terms-Firmware, hardware, wearable computing.

I. INTRODUCTION

I N 1988, the chief scientist at Xerox PARC, Weiser, coined the term ubiquitous computing. He envisioned a future of ubiquitous computing that he called: "calm computing." He believed that computers should create calm by being quiet and invisible servants that help us to be more efficient in a way that feels intuitive [39]. His vision has inspired several new computing devices, including wearable devices.

These wearable devices have changed the way we perceive personal computing devices. Devices such as the Galaxy Gear [3], the Pebble [2], and the Fitbit [1] have created new ways for us to track our health and check our mail. However, as researchers and engineers begin to explore the potential of wearable devices, they are faced with the challenge of developing and testing custom hardware prototypes.

As we begin to consider the question of prototyping devices by generating hardware from code, there are two fundamental contexts in which the question should be considered. The first context is automatically generating hardware prototypes

Digital Object Identifier 10.1109/JIOT.2016.2537148

by analyzing the code they are required to run, thereby allowing software developers with limited hardware experience to develop their own prototypes. The second context is that of hardware/software co-design, where the experienced hardware engineer is interested in optimizing a hardware design utilizing information from the code that the platform is required to run. Thanks to research done by several researchers, we know a great deal about problems related to hardware/software co-design [11], [12], [14], [16], [36], [37].

In this paper, we consider the first context and attempt to reduce the time that it takes to develop a hardware prototype by proposing a code-first approach to the design of embedded systems. A code-first approach allows a software developer to begin developing a hardware prototype by first writing the code that it will run. After the code has been written, it can be analyzed to determine the hardware configuration that is required. Once the configuration has been determined, a list of modules and their appropriate ports are displayed and the system can be configured by plugging in the appropriate modules. After the software developer has tested the code on the configured hardware prototype, he or she can then automatically generate the design files that are required to fabricate a custom board. This is possible because the platform is comprised of a collection of modular software and hardware components. The hardware components consist of a main board with several hardware modules, while the software component is a modular middleware which consists of stateless libraries, which abstract each hardware component. This modular abstraction creates a direct relationship between the software module and the hardware module, thereby allowing us to synthesize a hardware configuration from a software definition. A code-first approach to designing embedded systems can be achieved by creating a direct mapping between a stateless modular middleware and modular hardware.

The intuition behind the board's design is that there are a collection of interfaces (i.e., SPI, I^2C , and UART interfaces) that components normally use to interface with microprocessors. By abstracting these interface communication protocols and directly exposing the associated pins, it is possible to design a board that allows software developers to automatically generate the hardware configurations that are required to run their software.

Fig. 1 shows a picture of the main board, which we call the Eunit. The E-unit is comprised of four expansion ports. Each port is designed to be compatible with a specific interface and therefore accommodates a particular type of module. The first port is designed to be compatible with I²C interfaces and accommodates modules (such as an accelerometer and gyroscope).

2327-4662 © 2016 IEEE. Personal use is permitted, but republication/redistribution requires IEEE permission. See http://www.ieee.org/publications_standards/publications/rights/index.html for more information.

Manuscript received October 23, 2015; revised January 19, 2016; accepted February 18, 2016. Date of publication March 09, 2016; date of current version September 08, 2016.

The authors are with the Department of Computer Science, College of William & Mary, Williamsburg, VA 23185 USA (e-mail: dggraham@cs.wm.edu; gzhou@cs.wm.edu).



Program MIC Display Bluetooth library library library Kernel Kernel Kernel functions functions functions Hardware MCU Bluetooth MIC Display module module module

Fig. 1. Layout of the mainboard (E-unit).

Fig. 2. Overview of the middleware architecture.

The second port is designed to interface with UART-based modules (such as a Bluetooth module). The third port is designed to interface with analog sensing modules (such as a humidity sensor or finger pulse sensor). The fourth and final port is an SPI port that is designed to be compatible with a display module.

In the past, hardware platforms have helped to catalyze innovation [19], [32]. In 2005, researchers at the University of California Berkeley released the Telos mote along with the TinyOS operating system. This platform provided computer scientists with the tools to design and evaluate new protocols for wireless sensor networks and devices. Unlike the Telos platform, our proposed platform is reconfigurable. This means that components can be added or removed from the platform. Reconfigurable platforms for prototyping large-scale devices are becoming increasingly popular.

This paper makes four contributions as follows.

- 1) It presents a platform for prototyping wearable computing devices.
- It proposes a code-first approach to the design of embedded systems that allows programmers to synthesize hardware configurations from software definitions.
- 3) It presents a method for automatically generating custom schematics from a hardware configuration.
- 4) It presents an approach for automatically identifying port conflicts in hardware/software co-designs.

II. SOFTWARE DESIGN

This section is divided into two subsections. In the first subsection, we present our proposal for a code-first approach to developing embedded systems. In the second subsection, we discuss the design of the modular middleware architecture that is used in our reconfigurable platform.

A. Code-First Approach to Embedded System Design

A code-first approach allows software developers to write an application without worrying about configuring its resource dependencies. For example, Microsoft's entity framework allows software developers to abstract database models as classes [9]. Once the user compiles the program, the framework will automatically generate the database's structure from the class definitions. This increases the programmer's productivity since she does not need to manage the database by writing SQL queries to create, update, and delete tables. Instead, the entity framework ensures that the database is compatible with the programmer's implementation.

A code-first approach to embedded system design allows the framework to configure the hardware platform by analyzing the code for hardware dependencies. These dependencies are then used to generate the hardware configuration that is needed to run the application. This is possible because of the modular architecture that maps one software module directly to one hardware component. This one-to-one mapping allows hardware dependencies to be determined by analyzing the software dependencies. However, it is not sufficient to only analyze software dependencies, the microcontroller may have physical constraints that may prevent a valid software model from being executed. For example, the software developer may include two UART libraries in their program. Though this is a valid software model some microcontrollers such as the MPG430G2553 can only accommodate a single UART module without the use of resource sharing hardware. To ensure that our approach only produces valid hardware configurations we need to verify that the software models are compatible with our hardware platform and microcontroller. In the following sections, we discuss the design of the modular architecture and demonstrate how it can be used to construct the software models.

B. Modular Middleware Architecture for Embedded Systems

Developing embedded software is tedious. The software developer needs to know what control registers to set and what data registers to read. This means that the developer needs to have intimate knowledge of the chip's architecture and the board's layout, and therefore must spend hours reading datasheets. In an effort to reduce the burden on the software developer, several researchers have proposed a collection of hardware abstractions to help address this problem [17], [18]. We take a similar approach by proposing a modular middleware architecture that allows the software developer to quickly build software for our platform. Fig. 2 shows an overview of the modular middleware architecture. The middleware is comprised of a collection of libraries/software modules. Unlike previous systems that associate modules with generic functionality [13], our middleware architecture associates software modules directly with hardware modules. This allows the middleware to be tailored directly to the hardware platform's current configuration, resulting in a smaller code size.

Each software module is responsible for abstracting the hardware configurations of its corresponding hardware module. In particular, each software module is responsible for three tasks: 1) managing the appropriate hardware control and data registers; 2) performing the calculations associated with the module; and 3) managing the system's power. For example, the light-sensing software module is responsible for setting the control registers associated with controlling the MCU's analogto-digital converter (ADC) and CPU. Since the sensing module is allowed to set the CPU's control registers, it can disable the CPU to save power while it is waiting for the ADC to settle on a value. Once the ADC has settled on a value, the sensing module can wake up the CPU and perform the necessary calculations.

All software modules are required to be stateless. This means that the libraries do not assume that the MCU's control or data registers are in a particular state. This is an important requirement since memory limitations prevent us from using an operating system to provide resource management and protection. The absence of an operating system would be concerning if the platform was required to run multiple programs simultaneously. However, our platform is designed for specialized embedded applications where resources such as memory are too limited to support an operating system.

Extending the middleware is relatively easy, since it is a collection of decoupled stateless modules. The stateless nature of these modules allows a developer to extend the platform without impacting the other libraries. If the software developer chooses to use these libraries in addition to controlling the registers, he or she is able to do so without impacting the middleware since each module is stateless. Regrading the design of the libraries, there are two seminal questions which we believe need to be addressed. The first question is how do we partition functionality between software and hardware and how does this affect the performance of the design? This problem of dividing the systems functionality between hardware and software is known as the "partitioning problem." Several researchers have studied this problem and have proposed ways to achieve the fastest or lowest cost solution [14], [21], [28], [40].

This implies that there are several partition options depending on the constraints of our design. So this leads us naturally to the second question. How will a software developer with limited hardware experience choose and configure these different hardware implementations? This is where the library abstraction proves to be extremely useful. One advantage of creating a direct mapping between the hardware modules and the software libraries is that new hardware/software partitions can be selected by including different libraries. This means that a software developer can select the appropriate partition by simply including the library that meets their performance and cost constraints, without having any knowledge of the underlying hardware.

III. CONFIGURATION GENERATION PROCESS

In this section, we present an example that shows how a hardware configuration can be generated from a software

#include ''LCD.h''	1
#include ''tempSense.h''	2
#include ''helpers.h''	3
1 I	4
void main(void) {	5
unsigned long degrees;	6
char * reading;	7
board init();	8
LCD_init();	9
while (1) {	10
$LCD_gotoXY(0,2);$	1
LCD_writeString(' 'AT: ' ');	12
degrees = tempSense();	13
LCD_writeChar('(');	14
reading = convertADC (degrees,	1);
LCD_writeString(reading);	10
LCD_writeChar(0x7f);	17
LCD_writeString(''C'');	18
}	19
}	20

Fig. 3. Example program written using the modular middleware that displays the temperature on the LCD.

definition. In particular, we go through a detailed example of how to implement an indoor temperature sensor with an LCD using the proposed code-first approach. The example code in Fig. 3 shows a program that was written in C using our modular middleware architecture. The program begins by including three software modules/libraries. The first is an LCD module which abstracts the display component. The second module is the temperature sensing module. And the third module is the helper module which includes helper functions. These helper functions include the *board_init()* function that abstracts the setup of the Watchdog Timer and the *convertADC()* function that converts longs to formatted strings.

Lines 10–19 represent the program's running loop. In the loop, the program instructs the LCD screen to go to position (0, 2) and write the string "AT" Once this step has completed, the program calls the tempSense module which sets the appropriate control registers, reads the appropriate data registers, and returns the result, which is then converted to a string and displayed. The LCD software module supports unique characters with special codes, e.g., 0x7f represents the degree (°) character.

A. Introducing Mathematical Abstraction and Notation

There are currently solutions for partitioning software and hardware to create a performance optimized design using microcontrollers [14]. However, we were unable to find any research that addresses the problem of port conflicts that occur when implementing a collection of partitioned designs. A port conflict occurs when two partitions are included in the same design and require the same port. Fig. 4 shows an intuitive example of a port conflict between two hardware partitions. If a port conflict occurs, it is not possible to implement the design given current hardware/software partitions. Engineers at Texas Instruments have developed a tool called pinMux that helps resolve conflicts in complex designs [5]. However, the port requirements must be specified before the software is written. There are also cases in which the pinMux tool is unable to find a solution and therefore port conflicts must be manually



Fig. 4. Example of a port conflict on the UART port of the microcontroller reference design. BLE represents a Bluetooth module.

TABLE I Example of the Set P'

Property	Description
p_1	UART module
p_2	Sensing ADC module
p_3	3.3 V power requirement
p_4	Display module
p_5	On-board module
p_6	5 V power requirement

resolved by modifying the design to more efficiently use the ports. In this section, we present an approach for solving this port conflict problem by abstracting the problem as a constraint solving problem. The process is comprised of three steps:

- abstracting the collection of hardware/software co-design libraries as a set, which represents the system's configured state;
- 2) representing the microcontroller port limitation as a collection of constraining set functions; and
- applying the constraint set functions to the system set to automatically identify the software/hardware partition libraries that create port conflicts in the design.

B. Abstracting Libraries

The first step in the process of identifying port conflicts is abstracting the hardware/software co-design libraries as a set representing the system's state. These state abstractions are best explained by an example. Throughout our explanation, we will use the code in Fig. 3 to show how each abstraction applies to the specific example.

We can abstract each library as a tuple t which is comprised of a library identifier λ and a set P which contains n properties such that P is a subset of the defined set P' which contains the set of all properties p. Table I shows an example of the set P'

$$t = (\lambda, P = \{p_1, p_2, ..., p_n\}).$$
(1)

We define t_{λ} as representing the library identifier of the tuple and t_p as representing the property set in the tuple. Each value p_i describes a particular property of the library. These properties are the requirements that the associated hardware partition needs to interface with the software library. Table I shows a collection of example properties. For example, p_1 may mean that the library is a UART-based library, while property p_2 may mean that the component has an operating voltage of 3.3 V. Consider the "LCD.h" library from Fig. 3, it can be expressed in tuple form as: $(\lambda_3, \{p_3, p_4\})$. Table II shows a collection of example library identifiers and their associated properties.

Now that we have presented an abstraction for the software libraries and their associated properties, we can use this

TABLE II Example of the Set S'

Library identifier	Description	Properties	Include statement
λ_1	Bluetooth UART library	$\{p_1, p_3\}$	"Bluetooth.h"
λ_2	IR sensing	$\{p_2, p_3\}$	"IRSense.h"
λ_3	LCD display module	$\{p_3, p_4\}$	"LCD.h"
λ_4	Temperature sensing module	$\{p_3, p_5\}$	"tempSense.h"
λ_5	GPS UART library	$\{p_1, p_3\}$	"GPS.h"

approach to model the software system S as a collection of tuples that are a subset of the defined set of tuples S'. Table II shows an example of the set S'

$$S = \{ (\lambda_1, P_1), (\lambda_2, P_2) \dots (\lambda_n, P_n) \}.$$
 (2)

Consider the following example code shown in Fig. 3. Now that we have defined the sets P' and S', we can determine the system's tuple representation by examining the library dependencies of the program. The program may contain other libraries that are not associated with hardware modules. By checking the libraries in the program against a known set of libraries it is possible to extract the relevant libraries from the program. Consider the program in Fig. 3, it has three library dependencies: "LCD.h," "tempSense.h," and "helpers.h." We construct the tuple set by checking for the library in the tuple set S'. If the library is found we add it to the set S. For example, the "tempSense.h" and "LCD.h" libraries are both in the set S'so we add them to the subset S. However, the "helper.c" library is not in the set S' so we do not add it to the subset S. This results in the following set:

$$S = \{ (\lambda_3, \{p_3, p_4\}), (\lambda_4, \{p_3, p_5\}) \}.$$
(3)

C. Specifying the Constraints

Now that we have abstracted the hardware/software codesign libraries as a set representing the system's state, we need to represent the microcontroller's port constraints, so that they can be checked against the system's state. Recall that a port conflict occurs when there is a discrepancy between the system's state and the microcontroller's port constraints. For example, a program may include two libraries that have the UART property. This means that both libraries require the use of universal serial communication interface (USCI) module in the microcontroller. However, the msp430g2553 microcontroller does not have the required pins to support two UART devices, consequently even though the software definition of the program is correct, the hardware is unable to run it because the microcontroller does not have the required number of ports to support the design. In order to verify that a hardware configuration can be generated from the software model, we must show that the software model does not violate any of the hardware constraints. Formally, we can think of these constraints as a set of set functions

$$C = \{c_1(S), c_2(S), \dots, c_n(S)\}.$$
(4)

A constraint $c_i(S)$ is considered to be satisfied if it returns an empty set $c_i(S) = \{\}$. If a constraint is not satisfied, we say that it is violated. For a software model to be considered valid it must satisfy all of the constraints in the set C. Each hardware platform will have its own collection of constraints. For example, the proposed platform has a port constraint for UART modules p_1 . The platform can only accommodate one UART module. We call this type of constraint a uniqueness constraint. The following equation shows an example of an uniqueness constraint for property p_1 :

$$c_1(S) = \{t | t \in S \land \exists t' \in S : p_1 \in t_p \land p_1 \in t'_p \land t \neq t'\}.$$
(5)

Another type of constraint is a universal property constraint. An example of a universal property constraint is a power constraint, which requires all modules to meet the 3.3-V power requirement. The following equation shows an example of a universal property constraint for the property p_3 :

$$c_2(S) = \{t | t \in S \land p_3 \notin t_p\}.$$
(6)

We can define the violation set V as the set of the constraints that a system S violates

$$V(S,C) = \{ c | ((c \in C) \land c(S) \neq \{ \}) \}.$$
(7)

We say a system S is consistent with a constraint set C if the violation set is empty. An empty violation set means that the system set does not violate any of the hardware constraints. This formalization is important because as the complexity of the system grows we expect that these abstractions will form the basis for explicitly specifying constraints, though we do not expect to supplant formal systems such as TLA [24].

Now that we have defined the concept of a constraint, let us consider it within the context of the hardware platform in Fig. 1. Because of the microcontroller's architecture and the design of the reference platform, the platform has three uniqueness constraints for properties p_1, p_2, p_4 and one universal property constraint for property p_3 . Now let us consider how these constraints can be used to check the system set from our running example, shown in (3).

To ensure that the set represents a valid software model, we must ensure that all constraints are satisfied. First, let us consider the uniqueness constraint shown in (5). The uniqueness constraint requires that a property p_i is only found in a single tuple in the set S. In the set S, the property p_4 occurs only once so the constraint is satisfied. Since properties p_1 and p_2 are not in the set, their uniqueness constraints are also satisfied. Now that we have shown that all the uniqueness constraint are satisfied, we need to show that the universal property constraint in (6) is satisfied. The universal property constraint requires that all tuples in S have the property p_i . This constraint is satisfied for property p_3 since all the tuples in S have the property p_3 . Since all the constraints in C are satisfied, the software model is considered to be valid.

It may be difficult to grasp the purpose of the constraints in a scenario where they are not violated. So let us consider a system where the constraints are violated. To see how the constraints could be violated, let us extend the temperature display example to include a GPS module and a Bluetooth module. Fig. 5 shows the code for the new design. By following the procedure

#include ''LCD.h''	1
#include 'tempSense.h''	2
#include 'Bluetooth.h''	3
#include 'GPS.h''	4
#include 'helpers.h''	5
······································	6
void main(void) {	7
unsigned long degrees:	8
char* reading.	9
char* onsReading	10
hoard init():	1
LCD init();	12
$\mathbf{while}(1) \int$	13
$ICD \operatorname{goto} XY(0, 2)$:	1/
$LCD_gotoXI(0,2),$	14
degrees = tempSense();	1.
LCD write $Chen(f(2))$	11
$DCD_wineChar(()),$	1.
reading = convertADC (degrees, 1);	10
LCD_writeString(reading);	15
$LCD_writeChar(0x/t);$	20
LCD_writeString((,,C,));	21
$gpsReading = GPS_getValueString();$	22
bluetooth.send(reading);	23
bluetooth.send(gpsreading);	24
	25
}	26
}	27

Fig. 5. Modification of indoor temperature sensor to include GPS and Bluetooth. For this example, we assume the implementation of GPS module and it corresponding library.

outlined, we can define a new system set S_2 which presents this new GPS and Bluetooth capable system. The following equation shows definition of the new set:

$$S_{2} = \{ (\lambda_{3}, \{p_{3}, p_{4}\}), (\lambda_{4}, \{p_{3}, p_{5}\}), \\ (\lambda_{1}, \{p_{1}, p_{3}\}), (\lambda_{5}, \{p_{1}, p_{3}\}) \}.$$
(8)

Now we have a definition for the system set, we can once again apply the constraints. Notice that this time the uniqueness constraint is violated by both the GPS module and the Bluetooth module, since the GPS module and Bluetooth module both require the use of the UART ports and there are not enough pins on the microcontroller to accommodate. Notice also that the constraint does not simply return true or false but instead returns the violating tuples: $\{(\lambda_1, \{p_1, p_3\}), (\lambda_5, \{p_1, p_3\})\}$. These are then returned to the programmers as an error, notifying them of the libraries that have violated the limitation of the reference platform. Once these violation have been identified, the programmer may select new libraries that implement similar functionality but do not violate the port constraint of the microcontroller. For example, the programmer may choose a GPS implementation whose hardware/software partition uses an I^2C interface instead.

D. Automating Microcontroller Selection

Currently our system is designed to be a proof of concept, but commercial alternatives would allow engineers to select from a wide variety of microcontrollers. This means that these constraints could be used to inform processor selection. If a particular microcontroller does not meet the constraints of the application, these constraints could be used to search for a microcontroller that might be capable of running the

	T.	ABLE III	
EXAMPLE	OF	LIBRARY	MAPPINGS

Library identifier	Hardware id	Mapped module
λ_1	α_1	Bluetooth UART module
λ_2	α_2	IR sensing module
λ_3	α_3	LCD display module
λ_4	α_4	On-board temperature module
λ_5	α_5	Accelerometer gyroscope module

TABLE IV Example of Port Mapping

Hardware id	Port id	Port mapping
α_1	Υ_1	UART Port
α_2	Υ_2	Sensing ADC port
α_3	$ \Upsilon_3$	Display port
α_4	$\uparrow \Upsilon_4$	On-board no port
α_5	Υ_3	Sensing ADC port

application. In this case, the problem becomes an optimization problem where the system is attempting to find the best possible microcontroller whose constraints meet the requirements of a given application. Since companies such as Texas Instruments currently offer more than 40 different microcontrollers [4], approaches that provide automatic microcontroller selection would be extremely useful.

E. Generating the Hardware Configuration

Now that we have validated the software model, we can begin synthesizing the hardware configuration. The synthesis process is comprised of two steps. The first is to determine the list of hardware modules that are needed. The second step is to determine which ports each module plugs into. Our modular middleware architecture simplifies the first step by providing a mapping from the software libraries to the hardware modules. Table III shows the mapping of the libraries to the hardware modules. The design of the main board simplifies the second step by providing four unique expansion ports. Once, we determine what hardware modules are required, we can look up the associated port in Table IV.

Now let us consider the example in Fig. 3, we determined that the software model for this code was $S = \{(\lambda_3, \{p_3, p_4\}), (\lambda_4, \{p_3, p_5\})\}$. Formally, we can think of the synthesis process as a set function that converts the software model set S to a hardware model set H. We define a hardware model set H as a collection of tuples

$$H = \{ (\alpha_1, \Upsilon_1), (\alpha_2, \Upsilon_4), \dots, (\alpha_n, \Upsilon_m) \}$$
(9)

where α_i represents the hardware module id and Υ_i is the port identifier. We perform the first step of the synthesis process by identifying the corresponding hardware modules for libraries λ_3, λ_4 which are α_5, α_4 , respectively. Now that we know what hardware modules we need, we can look up the corresponding ports. This results in the tuple set $H = P\{(\alpha_3, \Upsilon_3), (\alpha_4, \Upsilon_4)\}$. From this configuration, the software developer knows to plug the LCD module into the display port and that the temperature sensor is already a part of the main board. The process of generating the hardware configuration from the software model can be easily automated since it is only a collection of lookups. Fig. 6(a) shows a picture of the resulting hardware configuration.



Fig. 6. (a) Temperature sensor with LCD display. (b) Overview of the steps in automating the synthesis process.

F. Automating the Process

In this section, we present the system architecture that implements the proposed formalization. Fig. 6(b) shows an overview of the system's architecture. We implement an alpha version of the system by developing an eclipse plug-in that is compatible with Code Composure Studio [20].

The system is comprised of four stages: a parsing stage, a constraint checking stage, a configuration generation stage, and a schematic generation stage. A key insight for implementing the system is realizing that the super sets S, H, and P can be represented as relational tables. By representing the sets as relational tables, it is possible to express and validate the system constraints as queries against these relations. This SQL-based architecture also allows the system to be updated as new libraries and constraints are added, since the database can be stored in a central location and queried remotely. However, before constructed. The system set is constructed by parsing all the files in the project directory and extracting the relevant include statements.

1) Parsing Step: During the parsing step, the system reads all the files in the project folder and extracts all of the libraries that are included in the system. This is done by examining the include statements in the project files. Once all the libraries have been extracted, the libraries that are not associated with the middleware must be removed before the system definition can be generated. This list is filtered by querying the tables associated with the S relation. The libraries that exist within the P relation are kept and those that do not, are discarded. The remaining libraries represent the system set S. This system set is then represented as an arraylist of library ids, and is encapsulated as a member of a system node class. 2) Constraint Checking Step: Now that the system set has been determined, the next step is validating this system set. A valid system set is one that does not violate any of the constraints associated with the system. Since each constraint can be expressed as a single SQL statement, it can be further abstracted as a visitor class which operates on the system node. When a visitor operates on the system node it returns a violation set containing all of the libraries that violate the constraint that is associated with the visitor. This violation set is then added to a global violation set that is associated with the system node class. Once all the visitors have visited the system node, and if the global violation set is empty, the system set is considered valid.

3) Configuration Generation Step: Now that the system set has been validated, the next step is generating the hard-ware configuration. The hardware configuration specifies which hardware modules connect to a particular port. Recall that there is a one-to-one relationship between hardware modules and software libraries. This direct relationship allows us to query the S and H relations to determine the appropriate port and hardware module.

4) Schematic Generation Step: Once the programmer has finished testing the hardware prototype by configuring the platform, she may want to fabricate a custom hardware prototype. However, in order to do this she will need a custom schematic that represents her prototype. In this section, we explain that it is possible to automatically generate a custom schematic using the proposed code-first approach. Schematic generation is possible because each hardware module maps directly to a schematic block and since each hardware module directly maps to a library, each library also maps to a schematic block.

Intuitively, this can be thought of as generating a schematic of the current hardware configuration while removing the unnecessary sections. A hardware schematic is normally represented using a .sch file. Fortunately, the .sch file is a xml-based file. Since the schematic for each hardware module is known, each module can be represented as a xml block. Because of this, each block can be added as a child in the .sch xml files. Once the schematic block is added to the schematic, it needs to be appropriately wired. Given that the tree-based hardware configuration already presents the appropriate wiring, the wiring for a particular schematic block can be determined by examining the port associated with the appropriate parent node. Since the connections from the ports to the modules are fixed, the appropriate wiring can simply be added to the schematic by looking at the wiring for the associated port.

IV. EVALUATION

This section is divided into three subsections. In the first subsection, we evaluate the design by building a step counter that communicates with an Android application. In the second subsection, we evaluate the flexibility of the design by: 1) prototyping a smartwatch that was designed to monitor the user's environmental exposure to temperature and light and 2) prototyping an indoor tracking module which tracks the user's location using a collection of landmark infrared transmitters. Tables V and VI show the libraries that were used in each prototype along with their associated properties.

TABLE V System and Associated Libraries

System	Library ID
Watch	"Bluetooth.h," "LCD.h," "tempSense.h," "lightSense.h"
IR system	"Bluetooth.h," "irSense.h"
Step counter	"AccelerometerRead.h," "Bluetooth.h"

TABLE VI System and Associated Libraries

Constraint type	Properties
Universal property constraint	p_1, p_2, p_4
Uniqueness constraint	p_3



Fig. 7. Different possible configurations of the platform: (a) environmental monitoring watch; (b) step counter; and (c) infrared system. (d) Screenshot of the accompanying Android application.

A. Step Counter

To test the system on a nontrivial example, we developed a prototype step counter and compared the results to the Fitbit. Fig. 7(b) shows a picture of the prototype step detector. The prototype was programmed to use a windowed peak detection algorithm [7]. We evaluated the prototype by having two participants wear the step counter for 2 days between the hours of 9 A.M. to 9 P.M. The results of each participant are shown in Fig. 8(a).

To accompany the wearable platform we build an Android application that is compatible with the platform. Fig. 7(d) shows a screenshot of the application. The application displays a real-time update of the readings that it receives from the platform. The readings are sent to the application from the module via Bluetooth. Once the Android application receives the values, it updates the appropriate section of the interface.

B. Environment Monitoring Smartwatch

We prototyped a smartwatch designed to monitor the user's ambient temperature and light exposure. The watch was constructed using four components: 1) the LCD display module; 2) 3.7-V battery pack; 3) the Bluetooth module; and 4) the mainboard. Once the components were placed on the



Fig. 8. (a) Results of the step detection experiment: results of the case study, in which two participants (Person 1 and Person 2) wore both the prototype and a Fitbit for 2 days. (b) Readings from the IR sensor: results of the localization experiment. When the signal is high, it indicates that the prototype received the signal, and when the signal is low, it indicates that the prototype has not received the signal. (c) and (d) Graphs of the light and temperature readings collected more than a 1-h period at 10 min intervals. During the first 50 min, the device was placed indoors, and during the final 45 min, the device was placed outside.

mainboard, we developed software that captured the light and temperature readings and displayed them on the LCD screen. Fig. 7(a) shows a prototype of the watch. Fig. 8(c) and (d) shows the results of the tests.

C. Infrared Indoor Localization Device

The third device that we prototyped was an indoor infrared localization device. This wearable device localizes an individual using unique infrared signatures from landmark infrared devices. Each landmark device is placed in a separate room and produces a unique infrared signal. As the user moves from room to room an infrared sensor located on the wearable device picks up the unique infrared signature of the landmark device. Since this infrared signature is unique to each room, this information can be used to localize the user.

We used our platform to quickly prototype this device using the mainboard and three modules: the infrared sensor, the Bluetooth module, and the 3.7-V battery. Fig. 7(c) shows a picture of the final prototype. This prototype is designed to fit in the user's shirt pocket, with the infrared sensor sitting slightly above the rim of the pocket so that it can pick up the infrared signatures from the landmark devices. The wearable device communicated with a smartphone using the Bluetooth module. However, if users would like to explore a lower power option they may elect to remove the Bluetooth module and replace the large 3.7-V battery with the smaller 3.3-V coin battery. Though the Bluetooth module is compatible with the 3.3-V voltage battery, the smaller 3.3-V coin does not have the capacity to sustain the Bluetooth module for long periods. So instead of using the Bluetooth module to transmit readings to a smartphone, the software developer can record the signature by writing the signature values to the chip's internal flash memory. These values can be retrieved later by connecting the platform to a PC and reading the chip's internal flash memory. In our evaluation, we used an infrared LED to present the landmark devices. We tested the device by walking into the room with the landmark device for 2 min and then walking out and waiting 2 min. The device was polled at 1 min intervals. Fig. 8(b) shows the results of the experiment.

Recall that constraints are fixed and are associated with the processor and do not change based on the application.

D. System's Limitations

One of the requirements of this approach is that the hardware/software partitions are sufficiently isolated. If the components are not sufficiently isolated they can affect the performance of the microcontroller and the other components. We tested this isolation requirement by connecting a motor controller circuit that was not sufficiently isolated to the E-unit. We also connected a Bluetooth module that was sufficiently isolated. We noticed failures in the Bluetooth module when both components were operated at the same time.

Another limitation of the system is the fixed power requirement. Since the ports on the reference design only provide 3.3 V, it is not possible to connect a module that has a 5-V requirement. Other researchers have proposed a custom port design that has two power supplies: a 3.3-V supply and 5 V [38]. Components that connect to a reference design simply select the appropriate power pin. Adopting this alternate design not only may help address these issues, but may also increase the complexity of the schematic generation process.

E. Considerations for a Commercial Solution

1) Interrupt Vectors: Interrupt vectors may prove problematic since the libraries are stateless and therefore cannot rely on the state of the microcontroller's interrupt vectors. This can be resolved using wait loops (spin locks) in place of interrupts. These spin locks limit the performance of the system where real-time performance is required. However, this can be mitigated using a real-time operating system with appropriately partitioned hardware/software libraries [6].

2) Operating Systems and Languages: In an ideal case, we would like to support a variety of operating systems and languages. Supporting languages such as Python and Java would make the approach more accessible to a wider variety of developers. However, the memory limitations of the microcontroller used in our prototype make it difficult to run an operating system and Java virtual machine, without which, it is impossible to run Java byte code.

3) Repository of Reference Designs: Facilitating a commercial implementation of the schematic generation process would require a repository of partitioned hardware/software codesign libraries that adhere to a specific format. Developing this repository would require significant engineering effort as both software libraries and hardware schematics would have to be developed for each additional module, before it could be leveraged by the synthesis process.

V. RELATED WORK

Research into rapid prototyping strategies can be divided into two major categories: fixed platform approaches and hardware/software co-design approaches. Fixed platforms use predefined designs, while hardware/software code design approaches use co-synthesis strategies to generate efficient hardware and software partitions.

The hardware/software co-design approaches that are the most similar to our design can be grouped into two subcategories: 1) interface-based designs and 2) platform-based designs. One of the earliest papers on interface-based design by Rowson *et al.* proposed a methodology for separating a component's behavior from how it communicates with other components in the system. Separating components in this way makes it easier to formally verify the component's behavior [33]. Since then, several researchers have explored this interfacebased design paradigm [10], [29], [30]. Though our platform uses an interface-based approach that is similar to previously proposed approaches, our approach extends the interface-based paradigm beyond the verification of a single component to the synthesis of an entire system.

The second hardware/software co-design approach is a platform-based design approach. A platform-based approach encourages the reuse of predesigned components through the use of automatic mapping tools [22], [34]. These automatic mapping tools use a layered approach to isolate and map an abstracted top layer description to a more detailed lower layer implementation. Consider the example of a field programmable gate array (FPGA), which uses a compiler to provide isolation and automatic mapping from the top layer VHDL abstraction to the lower layer implementation of logic blocks. In our approach, the libraries provide the abstraction for hardware/software partitions and the automatic mapping to lower-level hardware implementation is done by the automatic configuration generation process.

In addition to many hardware/software co-design approaches, several fixed platform-based approaches have also been proposed. Several companies and researchers have developed platforms that allow researchers and engineers to quickly prototype and test their new ideas. One of the earliest prototyping platforms was the Phidgets platform, which was developed in 2001 [15]. Afterward, in 2003, Plessl *et al.* advocated for the inclusion of FPGAs in sensor nodes [31]. They presented a sensor hardware architecture which coupled an FPGA with a CPU. By including the FPGA and allowing the CPU to configure it, they were able to dynamically configure the chips. Our approach does not use a reconfigurable chip. Instead, we focus on extending the capabilities of the platform by adding and removing external modules.

The earliest occurrence of an extensible platform that we found in the literature was the MetaCricket [26]. The MetaCricket consisted of a main board which connected to other devices and sensors. The board consisted of a main master controller and a supporting slave controller. The external sensors and expansion boards connected to the master controller while the slave controller controlled the board's internal components. Following the release of MetaCricket, researchers at Stanford University introduced the GoGoBoard in 2004 [35]. The GoGoboard was a low-cost programmable control and sensing board. The GoGoBoard was designed to be used as a learning resource in developing countries. With this goal in mind, the researchers focused on ensuring that the GoGoBoard could be assembled in developing countries. This decision influenced the board's design and the components that were selected. Unlike the GoGoBoard and the MetaCricket, our platform consists of both hardware and software components which make the process of prototyping easier.

In 2005, researchers at the University of California Berkeley proposed the Telos platform along with the TinyOS operating system [25]. The release of the reference platform and operating system sparked innovation in sensor network research. Following the release of the Telos platform, researchers at the University of California Berkeley released an updated platform called the TelosB. In 2009, researchers at University of California Berkeley [23] extended the capacity of the TelosB motes by creating two extension boards. The first extension board comprised of a triaxial accelerometer and a biaxial gyroscope. The second extension board provided electrocardiogram (ECG), and electrical impedance pneumography (EIP) functionality to the TelosB platform. These expansion boards demonstrated the flexibility of the TelosB motes. However, as new computing form factors emerge, such as body networks and wearable devices, researchers will need a platform that allows them to prototype devices for these new applications. Unlike the Telos and TelosB motes, our proposed platform can be extended using off-the-self components instead of custom extension boards. More recently, in 2014, researchers at the University of Florida have developed a reconfigurable RFID sensing tag [27]. The platform provides three pins that can be used to connect sensors to the platform. The platform also has an RFID antenna and Cortex M3 microcontroller.

Large companies have also attempted to develop reconfigurable platforms. For example, in 2008, Shimmer began developing their wearable computing development kit [8]. The Shimmer kit was a flexible health sensing kit which consisted of a collection of prebuilt expansion modules. Following the release of the Shimmer development kit, researchers at Microsoft proposed the Gadgeteer platform in 2011. The Gadgeteer platform is an extensible platform that allows researchers and industry professionals to quickly prototype hardware devices [38]. The Gadgeteer's main board is designed to use a collection of prebuilt modules which can be plugged into the main board. Unlike the Gadgeteer and Shimmer platforms, our proposed platform does not require custom modules. It works with off-the-shelf devices.

VI. CONCLUSION

In this paper, we have proposed an open source platform for prototyping wearable devices. The platform is comprised of a main board and four types of modules: 1) a Bluetooth module; 2) an LCD module; 3) a sensing module; and 4) a battery module. We designed the platform so that it can be built and assembled by researchers without having to depend on an expensive manufacturer. We also designed the device so that it can be easily programmed and debugged. We evaluated the platform using it to prototype three wearable devices: step counter, an environmental nb monitoring smartwatch, and an infrared-based localization system.

REFERENCES

- [1] "Fitbit Flex Product Page" [Online]. Available: https://www.fitbit.com/ flex, accessed on Jul. 8, 2014.
- [2] "Pebble Smartwatch Pebble Steal" [Online]. Available: https://getpebble. com/steel, accessed on Jul. 8, 2014.
- [3] "Samsung Galaxy Gear" [Online]. Available: http://www.samsung.com/ us/mobile/wearable-tech/SM-V7000ZKAXAR, accessed Jul. 8, 2014.
- [4] "Texas Instruments Development Kits and Software for Low-Power MCUs" [Online]. Available: http://www.ti.com/lsds/ti/microcontrollers_ 16-bit_32-bit/msp/tools_software.page, accessed on Oct. 14, 2015.
- [5] "Texas Instruments Pinmuxtool" [Online]. Available: http://www.ti.com/ tool/PINMUXTOOL, accessed on Oct. 10, 2015.
- [6] R. Barry, Using the FreeRTOS Real Time Kernel: A Practical Guide. Real Time Eng., 2010, FreeRTOS.org.
- [7] A. Brajdic and R. Harle, "Walk detection and step counting on unconstrained smartphones," in *Proc. ACM Int. Joint Conf. Pervasive Ubiq. Comput.*, 2013, pp. 225–234.
- [8] A. Burns *et al.*, "SHIMMER_{TM}—A wireless sensor platform for noninvasive biomedical research," *IEEE Sensors J.*, vol. 10, no. 9, pp. 1527–1534, Sep. 2010.
- [9] P. Častro, S. Melnik, and A. Adya, "ADO.NET entity framework: Raising the level of abstraction in data programming," in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2007, pp. 1070–1072.
- [10] P. Chou, R. Ortega, K. Hines, K. Patridge, and G. Borriello, "IPCHINOOK: An integrated IP-based design framework for distributed embedded systems," in *Proc. 36th Annu. ACM/IEEE Des. Autom. Conf.*, 1999, pp. 44–49.
- [11] G. De Michell and R. K. Gupta, "Hardware/software co-design," *Proc. IEEE*, vol. 85, no. 3, pp. 349–365, Mar. 1997.
- [12] S. Edwards, L. Lavagno, E. A. Lee, and A. Sangiovanni-Vincentelli, "Design of embedded systems: Formal models, validation, and synthesis," in *Readings in Hardware/Software Co-Design*. Norwell, MA, USA: Kluwer, 2001, vol. 86.
- [13] D. R. Engler et al., Exokernel: An Operating System Architecture for Application-Level Resource Management. New York, NY, USA: ACM, 1995, vol. 29.
- [14] R. Ernst, J. Henkel, and T. Benner, "Hardware–software cosynthesis for microcontrollers," in *Readings in Hardware/Software Co-Design*. Norwell, MA, USA: Kluwer, 2002, pp. 18–29.
- [15] S. Greenberg and C. Fitchett, "Phidgets: Easy development of physical interfaces through physical widgets," in *Proc. 14th Annu. ACM Symp. User Interface Softw. Technol.*, 2001, pp. 209–218.
- [16] R. K. Gupta and G. De Micheli, "Hardware-software cosynthesis for digital systems," *IEEE Des. Test Comput.*, vol. 10, no. 3, pp. 29–41, Sep. 1993.
- [17] V. Handziski, J. Polastre, J.-H. Hauer, and C. Sharp, "Flexible hardware abstraction of the TI MSP430 microcontroller in TinyOS," in *Proc. 2nd Int. Conf. Embedded Netw. Sensor Syst.*, 2004, pp. 277–278.
- [18] V. Handziski, J. Polastre, J.-H. Hauer, C. Sharp, A. Wolisz, and D. Culler, "Flexible hardware abstraction for wireless sensor networks," in *Proc.* 2nd Eur. Workshop Wireless Sensor Netw., 2005, pp. 145–157.
- [19] J. L. Hill and D. E. Culler, "Mica: A wireless platform for deeply embedded networks," *IEEE Micro*, vol. 22, no. 6, pp. 12–24, Nov./Dec. 2002.
- [20] Texas Instruments Incorporated, Code Composer Studio User's Guide, Texas Instrum. Lit. No. SPRU328B, 2000 [Online]. Available: http://www.ti.com/lit/ug/slau157am/slau157am.pdf
- [21] A. Kalavade and E. A. Lee, "A global criticality/local phase driven algorithm for the constrained hardware/software partitioning problem," in *Proc. 3rd Int. Workshop Hardware/Software Co-Des.*, 1994, pp. 42–48.
- [22] K. Keutzer et al., "System-level design: Orthogonalization of concerns and platform-based design," *IEEE Trans. Comput.-Aided Des. Integr. Circuits Syst.*, vol. 19, no. 12, pp. 1523–1543, Dec. 2000.
- [23] P. Kuryloski et al., DexterNet: An open platform for heterogeneous body sensor networks and its applications," in Proc. 6th Int. Workshop Wearable Implantable Body Sensor Netw. (BSN'09), 2009, pp. 92–97.
- [24] L. Lamport, Specifying Systems: The TLA+ Language and Tools for Hardware and Software Engineers. Reading, MA, USA: Addison-Wesley/Longman, 2002.
- [25] P. Levis et al., "TinyOS: An operating system for sensor networks," in Ambient Intelligence. New York, NY, USA: Springer, 2005, pp. 115–148.
- [26] F. Martin, B. Mikhak, and B. Silverman, "MetaCricket: A designer's kit for making computational devices," *IBM Syst. J.*, vol. 39, no. 3.4, pp. 795– 815, 2000.

- [27] M. S. Khan, M. S. Islam, and H. Deng, "Design of a reconfigurable RFID sensing tag as a generic sensing platform toward the future Internet of Things," *IEEE Internet Things J.*, vol. 1, no. 4, pp. 300–310, Aug. 2014.
- [28] R. Niemann and P. Marwedel, "An algorithm for hardware/software partitioning using mixed integer linear programming," *Des. Autom. Embedded Syst.*, vol. 2, no. 2, pp. 165–193, 1997.
- [29] R. Passerone, L. De Alfaro, T. A. Henzinger, and A. L. Sangiovanni-Vincentelli, "Convertibility verification and converter synthesis: Two faces of the same coin," in *Proc. IEEE/ACM Int. Conf. Comput.-Aided Des.*, 2002, pp. 132–139.
- [30] R. Passerone, J. A. Rowson, and A. Sangiovanni-Vincentelli, "Automatic synthesis of interfaces between incompatible protocols," in *Proc. 35th Annu. Des. Autom. Conf.*, 1998, pp. 8–13.
- [31] C. Plessl et al., "The case for reconfigurable hardware in wearable computing," Pers. Ubiq. Comput., vol. 7, no. 5, pp. 299–308, 2003.
- [32] J. Polastre, R. Szewczyk, and D. Culler, "Telos: Enabling ultra-low power wireless research," in *Proc. 4th Int. Symp. Inf. Process. Sensor Netw.* (*IPSN*'05), 2005, pp. 364–369.
- [33] J. A. Rowson and A. Sangiovanni-Vincentelli, "Interface-based design," in Proc. 34th Annu. Des. Autom. Conf., 1997, pp. 178–183.
- [34] A. Sangiovanni-Vincentelli, "Defining platform-based design," *EEDesign EETimes*, 2002 [Online]. Available: http://www.eetimes.com/document.asp?doc_id=1204965
- [35] A. Sipitakiat, P. Blikstein, and D. P. Cavallo, "Gogo board: Augmenting programmable bricks for economically challenged audiences," in *Proc.* 6th Int. Conf. Learn. Sci., 2004, pp. 481–488.
- [36] M. B. Srivastava and R. W. Brodersen, "Rapid-prototyping of hardware and software in a unified framework," in *Proc. IEEE Int. Conf. Comput. Aided Des. (ICCAD'91) Dig. Tech. Papers*, 1991, pp. 152–155.
- [37] K. Van Rompaey, I. Bolsens, H. De Man, and D. Verkest, "CoWare— A design environment for heterogenous hardware/software systems," in *Proc. Conf. Eur. Des. Autom.*, 1996, pp. 252–257.
- [38] N. Villar, J. Scott, S. Hodges, K. Hammil, and C. Miller, ".net gadgeteer: A platform for custom devices," in *Pervasive Computing*. New York, NY, USA: Springer, 2012, pp. 216–233.
- [39] M. Weiser and J. S. Brown, "The coming age of calm technology," in Beyond Calculation. New York, NY, USA: Springer, 1997, pp. 75–85.
- [40] W. H. Wolf, "Hardware–software co-design of embedded systems (and prolog)," *Proc. IEEE*, vol. 82, no. 7, pp. 967–989, Jul. 1994.



Daniel Graham received the B.S. degree and M.Eng. degree in systems engineering from the University of Virginia, Charlottesville, VA, USA, in 2010 and 2011, respectively, and is currently working toward the Ph.D. degree in computer science at the College of William & Mary, Williamsburg, VA, USA.

He is currently with the Department of Computer Science, College of William & Mary. His research interests include intelligent embedded systems and networks.



Gang Zhou (GSM'06–M'07–SM'13) received the Ph.D. degree from the University of Virginia, Charlottesville, VA, USA, in 2007.

He is currently an Associate Professor with the Department of Computer Science, College of William & Mary, Williamsburg, VA, USA. He has authored more than 70 academic papers in the areas of ubiquitous computing, mobile computing, sensor networks, and wireless networks. The total citations of his papers are more than 4500 according to Google Scholar, among which five of them have been trans-

ferred into patents and his MobiSys04 paper has been cited 800 times. Thirteen of his papers have each attracted more than 100 citations since 2004.

Dr. Zhou is a Senior Member of the ACM. He currently serves on the Journal Editorial Board of the IEEE INTERNET OF THINGS, as well as Elsevier *Computer Networks*. He was the recipient of an award for his outstanding service to the IEEE Instrumentation and Measurement Society in 2008. He was also the recipient of the Best Paper Award of the IEEE ICNP 2010 and the NSF CAREER Award in 2013.