

Software development for a novel WSN platform

P. Völgyesi, J. Sallai, Á. Lédeczi

Vanderbilt University
Nashville, TN, USA

akos.ledeczi@vanderbilt.edu

P. Dutta

University of Michigan
Ann Arbor, MI, USA

prabal.dutta@michigan.edu

M. Maróti

University of Szeged
Szeged, Hungary

mmaroti@math.u-szeged.hu

ABSTRACT

This work-in-progress paper introduces a new hardware platform for wireless sensor networks, summarizes the new challenges it creates for software development and describes a toolchain being developed to meet those challenges. The hardware platform is based on a low-power FPGA as opposed to a traditional microcontroller. The FPGA configuration includes a soft core microcontroller, but there are plenty of resources left to implement a subset of the operating system, middleware and application components directly on the FPGA. Instead of creating this partition early in the design phase, we advocate a flexible hardware/software boundary enabling "late binding" of components to the soft core or the hardware fabric. This increases the complexity of the design space mandating sophisticated tool support. The paper describes a toolchain that helps manage this complexity. The two main tools are a domain-specific modeling environment and a symbolic design-space exploration tool.

Categories and Subject Descriptors

C.2.4 [Computer-Communications Networks]: Distributed Systems; D.2.2 [Software Engineering]: Design Tools and Techniques

General Terms

Design, Algorithms, Experimentation

Keywords

Wireless Sensors, FPGA, Design Space Exploration

1. INTRODUCTION

The prevailing approach to short-range, low-power wireless communication hardware typically used in wireless sensor networks (WSN) is to combine integrated radio transceivers with general purpose microcontrollers. However, microcontrollers with their resource constraints significantly constrain

performance. Experimenting with novel protocols or innovative applications that require higher performance than microcontrollers offer, can be cumbersome and tedious because it often requires developing custom hardware.

Software Defined Radios (SDR), with their high computing power and easy configurability, overcome these limitations but they bring other problems. They typically require wall power and a dedicated PC, and they are expensive and physically bulky. As a result, SDRs are not suitable for low-power, large scale, or mobile experimentation. The flexibility and performance of SDRs come from their FPGA-based architecture. While the active power draw of FPGAs is reasonable, especially considering their performance, it is hardware duty cycling that ultimately determines the lifetime of a WSN node. However, traditional FPGAs do not allow efficient duty cycling due to their slow and power hungry startup. In contrast, the recent appearance of Flash-based FPGAs eliminates this limitation. These devices can wake up in a microsecond with no significant in-rush current spike, and they preserve their configuration and register values even in deep sleep. Therefore, they present a great opportunity to build a reconfigurable, low-power wireless node that enables a new class of protocols and applications.

Such a platform could exhibit two orders of magnitude higher performance than microcontroller-based nodes, depending on the application, while drawing about the same power. This performance improvement comes at a cost. The new platform presents new challenges because joint design of hardware configuration and soft-core hosted embedded software is a significant departure from traditional wireless protocol and application development. There is a clear need for innovative design methodologies, since boundaries between application, services, operating system, and hardware are blurred, therefore traditional system design techniques are challenged in this domain. We argue that it is the application that dictates whether a certain functionality is implemented in software or hardware – whether multiplexing computation over time or parallelizing it over space is more appropriate. We present a toolchain that supports uniform handling of reusable hardware and software components, and promotes late binding of components to a particular realization. Since the flexible hardware-software boundary generates an expanded application design space with increased complexity, the key component is an integrated design space exploration tool that aids in finding the right combination of hardware or software implementation of components that meets the application requirements.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SESENA'10, May 3, 2010, Cape Town, South Africa.

Copyright © 2010 ACM 978-1-60558-969-5/10/05...\$10.00.

2. HARDWARE PLATFORM

The novel hardware architecture is built around a low-power Flash-based FPGA device (Actel IGLOO family), as shown in Figure 1. The soft microprocessor core (ARM Cortex-M1) – synthesized on the FPGA fabric – provides an easy to use computational platform capable of running standard applications written in C. High level sequential logic and other system integration tasks can be implemented just like on any existing processor-based platforms. By using a soft core – instead of a discrete microcontroller – we can tailor the processor (speed, capabilities) to the application and provide tighter integration with other IP cores.

The real strength and novelty of the platform, however, lies in the remaining part of the FPGA fabric. This is where key elements of the physical and media access layers of the radio communication can be implemented as IP cores. Also, it can be populated by the usual peripheral modules (UARTs, timers, I²C, A/D drivers, GPIO) based on the demands of the application. Finally, high performance signal processing – or other computationally intensive tasks – can be accelerated by providing them as IP cores. At the machine level, the soft processor communicates with other peripherals and custom cores via a well-defined register interface. Higher level software interfaces and thin proxies can be used to hide the differences between a software-based implementation and its hardware-based alternative.

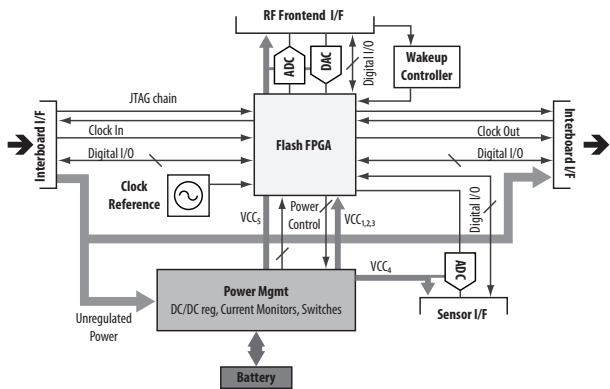


Figure 1: Novel hardware architecture for wireless sensor nodes.

Efficient power management is the cornerstone of this low-power platform. The fine-grained control over the different power supply networks provides a straightforward way to save power while the FPGA is turned on and the clock networks are active. However, truly power aware WSN applications need to duty cycle the central processing unit also. Unlike traditional FPGAs, Flash technology enables very low standby currents and rapid wake-up.

3. SOFTWARE DEVELOPMENT

While we have seen tremendous progress in architectures and programming models for low-power wireless networks in the past decade, application development targeting soft-core enabled FPGAs still poses significant challenges. Expertise in both hardware description languages and embedded software tools is required. Furthermore, awareness of low-level

details must be coupled with a high-level view of the system. Of particular importance is the partitioning of functionality between hardware and software, and the interfaces between the two domains. Migrating functionality from software to hardware, and vice versa, is cumbersome today and often lacks efficient tool support and limited reuse of software or hardware descriptions result in frequent reimplementations of the same functionality.

The requirements for joint development methodologies for reconfigurable, low-power hardware and soft-core hosted embedded software are significantly different from those for low-power wireless applications. Boundaries between application, services, operating system, and hardware are blurred, therefore traditional system design techniques do not work in this domain. There is a clear need for a joint hardware-software design methodology that allows for application-specific optimizations, supports cross-layer protocols, enables propagating low-level concerns to high-level design, and promotes reuse. Lessons learned in hardware-software codesign [4] unanimously support this claim.

The development methodology presented here promotes late binding of components, aiming to avoid ‘dead ends’ in the development process resulting from committing to a poor hardware-software decomposition too early in the design cycle when the true bottlenecks and tradeoffs are still unknown. A central idea of our approach is the flexible hardware/software boundary. The key enabling technology here is design space exploration: given a set of alternative implementations (components, services, modules, IP cores) that meet the same functional requirements, the toolchain finds the combination of components that fulfill the application requirements. Design space exploration relies on component property annotations such as gate count, memory usage, power draw, sampling rate, latency, jitter, etc.

The unified hardware-software architecture relies on concepts such as rich interfaces, compositionality, and uniform treatment of hardware and software components. The architecture utilizes the following concepts:

Unified hardware-software components. Application, middleware and operating system functionality are implemented in components. The architecture treats software and hardware components uniformly, therefore software and hardware implementations are interchangeable as long as they use and/or provide the same interfaces. Components are composable, meaning that a set of components can be grouped into a composite component. A composition may include wirings between the child components’ interfaces or from an interface used/provided by the composition to an interface used/provided by a child. This heterogeneous hierarchy is the key enabler of the flexible hardware-software boundary.

Rich interfaces. Components interact via messaging through interfaces. In software, these messages are function calls, while in hardware, they manifest themselves as signals; both of which are treated uniformly in interface definitions. An interface is a collection of named and strongly typed inputs and outputs that are related to the same service. Interfaces are bi-directional: a component that provides an interface must be prepared to handle all of inputs defined in the interface, and may produce outputs that, in turn, must be handled by any components using the interface. Interface definitions are annotated by contract specifications: invariants, pre- and postconditions on the parameters, as well as

the dynamic contracts that provide constraints on the ordering of messages passing across the interface (e.g. *init* must precede *send*, *send* not allowed after *stop*, etc.)

Transcendent types. Type definitions must be portable across the hardware-software boundary. Therefore, it is imperative that the architecture provide a unified type system. The type system includes boolean, integer, floating point and composite types (structures, unions, arrays) which have well defined bindings in both the software and in the hardware description languages.

Adapters. Hardware and software components are linked together using adapters. Adapters are reusable, bidirectional hardware-software bridges that translate between function calls and hardware signals. Adapters are the key elements of hardware-software interaction, and provide a clean way to hide the complexities of the hardware-software boundary.

Creating a toolchain that supports a unified hardware-software development process is challenging. Synthesis tools are supplied by the FPGA vendor and provide support for the hardware description language to FPGA path, and there exist C compilers for most soft MCU cores. Also, there is a broad spectrum of operating systems and programming frameworks (TinyOS [6], Contiki [5], MANTIS OS [2], etc.) targeting low-end microcontrollers that can be easily adapted to run on soft processor cores. In our current work, we decided to utilize TinyOS for multiple reasons. It is probably the most widely used operating system in the WSN domain. It has an elegant components-based architecture. And finally, we have several years of experience working with it.

Traditionally, hardware-software partitioning decisions, and the corresponding interfaces, are determined early in the design process, so these tools are used independently and side-by-side. Our architecture, however, dictates a significant departure from this design process, that allows for deferring architectural decisions to late in the design. We believe that a novel, high-level tool seamlessly integrating existing tools that allows for uniform handling of software and hardware components is required to meet the requirements of the proposed architecture. In our architecture, this role is played by the Generic Modeling Environment (GME) [7], a configurable modeling, model analysis and system generation tool suite. GME and the set of related tools implement Model Integrated Computing (MIC). MIC is predicated on the notion that many times different engineering or scientific domains require their own domain-specific system representation. Hence, MIC supports the rapid generation of domain-specific modeling languages and their supporting tool infrastructure utilizing formal metamodels.

We created the metamodels defining the modeling language and configured GME to support modeling WSN application on the new hardware platform. Component specifications (interfaces, composition, and interceptors) are captured in the high-level primarily visual modeling language. Their actual implementations of the (non-composite) components are programmed by the user. To allow for flexibility at the hardware-software boundary, certain components will have both hardware (HDL) and software (C) implementations. Resource requirements and performance attributes of the different component implementations (e.g. gate count, memory usage, power draw, sampling rate, latency, jitter) are also captured in the models.

The models are utilized by symbolic design space exploration tool based on DESERT [9], that aids the programmer in choosing from alternative hardware and software implementations of particular components. Clearly, the introduction of alternatives into the design space causes a combinatorial explosion in the number of possible designs, making system optimization a challenging task. DESERT tackles this problem by transforming the resource requirements and performance attributes of components as well as application requirements and hardware resource limits into constraints. DESERT then carries out symbolic constraint satisfaction to prune the design space based on user specified constraints on the overall system properties. As a result, the user is offered a set of concrete designs, with components bound to particular realizations, that meet the given constraints.

3.1 GRATIS++: An MIC Approach

A TinyOS application is a hierarchical component assembly where component configurations, i.e. wiring specifications, interface declarations and module implementations are specified in numerous text files. Graphical representation of the same information increases the readability and understandability of the application architecture and helps in avoiding configuration errors, such as the omission of the wiring specification of one or more interfaces of a component.

Our previous work, the Graphical Development Environment for TinyOS (GRATIS) [11] is a typical application of Model Integrated Computing (MIC) in general, and the Generic Modeling Environment (GME), in particular [7]. GME is a metaprogrammable toolkit for creating domain-specific modeling environments. GME metamodels specify the modeling language of the application domain. They are used to automatically configure GME for the domain, that is, to create a modeling environment that has native support of the target modeling language.

GME models take the form of graphical, multi-aspect, attributed entity-relationship diagrams. Their syntax is defined by the metamodels specified in a UML class diagram-based notation. The static semantics of a model are specified by OCL constraints that are also part of the metamodels. They are enforced by a built-in constraint manager during model building time. The dynamic semantics are applied by the model translators, i.e. by the process of translating the models to source code, configuration files, database schema or any other artifact the given application domain calls for.

This approach fits component-based software development very nicely. The interface of the individual components can be modeled along with a link to their implementation. The model editor can enforce the composition rules to make sure that only valid component assemblies are allowed. More sophisticated analysis can be performed by interfacing to outside tools. Finally, model translators can generate the glue code that ties the final system together. For this work, we are enhancing Gratis by 1) adding hardware components that are to be mapped directly to FPGA, 2) adding the ability to specify alternative implementations for the same functional component and 3) interfacing it with DESERT, our design space exploration tool. We call the enhanced tool GRATIS++.

Metamodels. The metamodel of GRATIS++, shown in Figure 2, defines the mapping of TinyOS concepts to GME concepts. The three basic building blocks of Gratis mod-

els are *interfaces*, *modules* and *configurations*. Modules and configurations can have variants captured as *moduleAlternative* and *configurationAlternative*, respectively. An interface consists of a set of *events* and *commands*. Both events and commands are functions. The return type is captured by a textual attribute, while the *arguments* are modeled with contained objects each having its own type declaration. A module contains a set of interface references (*interface_ref*). A reference is a graphical object that points to another object contained elsewhere in the model hierarchy. This is captured in the metamodel by a directed connection pointing from the interface to the interface_ref metamodel in Figure 2. Interfaces are declared at the global level and modules do not contain them directly; they just refer to their declaration through the use of references. This allows multiple modules using and/or implementing the same interface declarations.

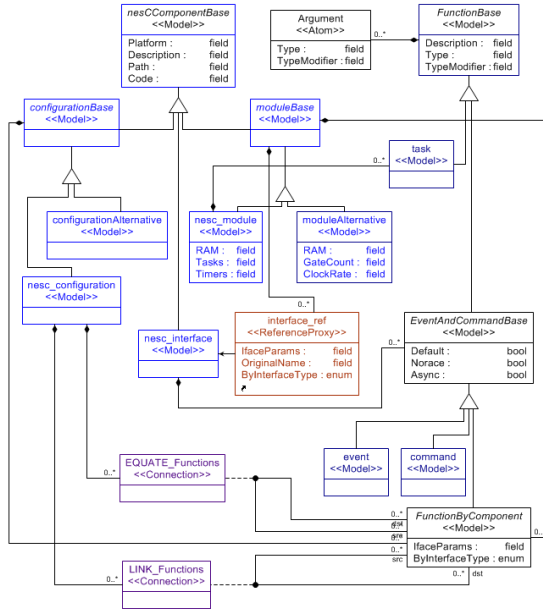


Figure 2: Partial and simplified metamodel for capturing the design space of mixed hardware and software components.

Similarly, configurations contain references to interfaces, modules and other configurations (not shown in Figure 2 for clarity). Interface references contained in modules and configurations appear as ports in higher level configurations. Component wiring specifications are expressed in Gratis++ as connections between interfaces and/or interface ports in configurations. In fact, two different kinds of connections are used in configurations. A *LINK* specifies that a component uses an interface that another provides. An *EQUATE* connection specifies that the interface the given configuration uses/provides is delegated down to a contained component that either implements it or delegates it further down the component hierarchy.

Data Type Modeling. Data types in Gratis++ are modeled similarly to the MILAN system [8]. It allows the specification of both simple and composite types. Simple types, such as floats and integers, specify their representation size, i.e. the number of bits used. Composite types can contain simple types and other composite types. At-

tributes of the fields specify extra information such as array size or signed/unsigned type. All data types supported by nesC programming language can be modeled in Gratis++. Preexisting data types, specified in a library for example, can also be modeled. To describe the entire type system of a given application, all the necessary data types and their relations need to be modeled. If a given simple type can be converted to another without loss of precision (or with a loss of precision that is acceptable for the given application), they need to be connected with a directed connection. If a given simple or composite type can be converted to another with a conversion function, then they need to be connected together through a converter model that specifies the conversion function in the target programming languages. This way, the data type models form a directed, possibly disconnected, graph. A directed path from a node to another one means that there is a valid conversion from the source data type to the destination one. The model interpreters insert the necessary conversion functions automatically. Furthermore, correct typing is enforced during model building time. This is accomplished by a set of constraints that only allow connecting components whose types are compatible.

Interface Modeling. Traditional programming languages and interface description methods—such as CORBA IDL—capture only the type aspects of software components. The access points of a given component are enumerated along with their accepted and returned parameter types in terms of values and domains. nesC is no exception to this: component interfaces are defined by a set of function declarations. Compatibility checking provided by compilers guarantees that the user of a function provides the required parameters and handles the returned value in a type-safe manner. Even in trivial applications, the access points of a software component are not isolated, dependencies and complex relationships might impose additional constraints on the use of their services. Typical patterns—such as *initialization before use*—can be found in almost every component. A component providing communication services may have more restrictions that are inherent in the communication protocol. Even if the legal order and dependencies of the function calls are described in the documentation of the component as informal rules, automatic tools and formal methods cannot be developed to verify these constraints.

We have developed a modeling language based on Interface Automata [1] that captures the dynamic aspects of component interfaces and enables us to describe more complex behavior [11]. Gratis++ also includes a translator that targets a model checking tool called SPIN [3] that verifies interface compatibility across the entire design.

System Generation. The only information captured textually in Gratis++ is the code of module implementations, either in C or VHDL. Because we capture the design space of the application through the use of alternatives, a selection needs to be made which alternative implementation is to be utilized in a given design instance. This can be done manually in the modeling environment or with the help of DESERT. Once a particular design point is picked, the translator generates all the nesC files containing interface, module and configuration specifications automatically. It also inserts the necessary adapters at the hardware/software boundary and assembles the final VHDL files.

Keeping the graphical models and the corresponding nesC files in synch is a challenge, especially because a large code

base of TinyOS components exists in text form only. Therefore, the Gratis model translator is bi-directional; not only does it generate the nesC files from graphical models, but it is also capable of parsing existing source files and building the corresponding models automatically. The main use of this parsing feature is to automatically generate the graphical equivalent of the TinyOS system components and to provide them as a library to the user in the Gratis++ environment. This library can then be refreshed when new TinyOS versions become available without any modifications to existing graphical application models.




Figure 3: Example application model

Example. An example model specifying one of the alternatives for the detection component of an acoustic shooter localization application [10] is shown in Figure 3. The main window shows the structure of the model in the Hierarchy aspect. The Constraint and Interface Automata aspects are not shown, but they can be selected by clicking on a tab in the bottom center window, the Partbrowser showing the available parts for this kind of model. The bottom right window shows the hierarchical structure of the entire application and TinyOS, over a thousand components. The bottom left window shows the attributes of the selected model. It is here where the user can pick this option from the available alternatives.

3.2 Design Space Exploration

Given the flexibility in defining design alternatives and configuration parameters, the design spaces for the systems represented can be extremely large. However, it is expected that only a subset of these designs will satisfy all the con-

straints and, hence, meet the design goals. Thus, a design space exploration method is desired to be able to rapidly navigate, and prune this large design space to select feasible design alternatives, and configuration parameters, that satisfy the user-defined constraints. Given the size of the design space, and the complexity of the analysis, DESERT, a powerful, scalable Ordered Binary Decision Diagram-based (OBDD) design space exploration tool was developed.

The design space exploration method relies on a symbolic Boolean representation of the space. A binary encoding is defined over the member elements of this space. The entire space can be symbolically represented as a conjunction over the Boolean representations of individual elements. OBDDs represent Boolean functions as directed acyclic graphs in a memory efficient format. The operations over these functions are implemented as graph algorithms, thus rendering "manipulation" of the space fast and efficient. Logical (compositional) constraints can be solved with ease with this symbolic Boolean representation. The logical relation expressed in the constraint over the elements of the design space is simply transformed to a logical relation between the Boolean representations of these elements. The resultant expression represents symbolically the "constrained" design space. Performance constraints can also be solved, however the mapping is non-trivial [9].

The power of this approach is the fact that it obviates the need for exhaustive combinatorial enumeration of all design choices. The entire design space can be symbolically evaluated without enumerating individual design points, thus rendering the approach highly scalable for exploring large design spaces. However, in large design spaces with many constraints simultaneously applied, an exponential explosion of the OBDD can occur. To address this problem, the constraint processing is done hierarchically with constraints scoped to a particular level; i.e. constraints are applied to sub-spaces first, pruning them to the extent possible and then progressing upwards in the hierarchy. This technique is very effective when there are a large number of constraints with a limited scope.

The design space exploration step, progresses by applying the constraints one at a time. Each constraint application results in a pruning of the space. Moreover, the pruned design space contains only the designs that are "correct" with respect to the applied constraints. When the initial design space is reduced to a manageable number of designs, the designer can progress to the next step of design simulation. Notice that some conflicting constraints may result in the elimination of the design space altogether, i.e. no design satisfies all the constraints simultaneously. In this case, some of the constraints must be relaxed.

DESERT has a bi-directional interface to Gratis++. A model translator generates a representation of the design-space and all constraints in an XML file. This file is parsed in DESERT and a tree-based view of the space is presented to the user in a GUI. The user can either apply all constraints at once or explore the space by selectively applying some of the constraints. DESERT always displays the size of the resulting pruned space. Once the user is satisfied with the results, i.e. the size of the space is manageable and all the important constraints have been applied, the resulting space is shipped back to Gratis++. If the size of the space is larger than one, the user needs to select manually from the still available choices.

3.3 Design Flow

The tool architecture is shown in Figure 4. When the user starts a new design he has a library of TinyOS and possibly HDL component models already available in Gratis++. Any new application-specific components need to be modeled and implemented in nesC and/or VHDL. These components also need to be characterized for resource usage and performance. Then the user can model the entire application in Gratis++. If there are alternative implementations available for certain components, the design-space needs to be pruned either manually or through the use of DESERT. The user then chooses a particular design from the pruned set, that is, a concrete binding of components to hardware or software, which is then mapped to a concrete application image.

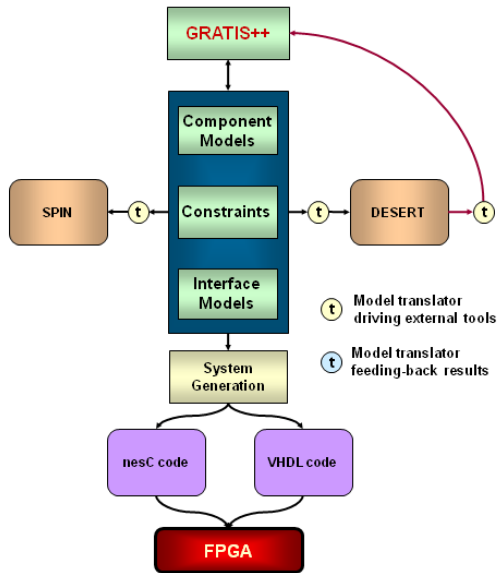


Figure 4: Example application model

Mapping of the component specification of an application involves the following steps: 1. Component hierarchy is flattened, resulting in a component graph that contains non-composite components and the connections of their interfaces. 2. Static interface compatibilities are checked (function/signal names, types, etc.). Optionally, a model translator generates a representation of the system utilizing the Interface Automata models and feeds it to SPIN where dynamic compatibility checks are performed. 3. Interface connections are mapped to (a) hardware configurations connecting hardware components, (b) software configurations connecting software components, or (c) adapters connecting a software component to a hardware component. 4. nesC code is assembled from software components, VHDL code is assembled from hardware components. 5. The HDL code is handed over to the FPGA synthesis toolchain, the nesC code to the corresponding compiler. 7. The hardware is programmed with the resulting images.

4. CONCLUSIONS

The paper presented a work-in-progress system to develop a toolset to support application development for a novel

FPGA-based sensor node platform. The foundation of the approach, Model Integrated Computing and its supporting tools, GME and DESERT in particular, are mature and have been proved in many domains in industry and academic research alike. However, Gratis++ is only partially functional at this stage. Significant effort remains to fully implement its VHDL capability. There are also hurdles as far as user adoption is concerned. Our experience with the original GRATIS environment was that only hardcore engineers program in TinyOS and they want to stick to their text-based programming tools. The hope is that the increased complexity brought about by the flexible hardware/software boundary and the capability to deal with design spaces not just point solutions will provide enough incentive for users to switch to our model integrated approach.

5. ACKNOWLEDGMENTS

This research was partially supported by ARO MURI grant W911NF-06-1-0076 and the TÁMOP-4.2.2/08/1/2008-0008 program of the Hungarian National Development Agency.

6. REFERENCES

- [1] L. Alfaro and T. A. Henzinger. Interface automata. *Ninth Annual Symposium on Foundations of Software Engineering (FSE)*, 2001.
- [2] S. Bhatti, J. Carlson, H. Dai, J. Deng, J. Rose, A. Sheth, B. Shucker, C. Gruenwald, A. Torgerson, and R. Han. Mantis os: an embedded multithreaded operating system for wireless micro sensor platforms. *Mob. Netw. Appl.*, 10(4):563–579, 2005.
- [3] H.-L. Chang, J.-B. Tian, T.-T. Lai, H.-H. Chu, and P. Huang. Spinning beacons for precise indoor localization. In *Proc. of ACM Sensys*, 2008.
- [4] K. Compton and S. Hauck. Reconfigurable computing: a survey of systems and software. *ACM Comput. Surv.*, 34(2):171–210, 2002.
- [5] A. Dunkels, B. Grönvall, and T. Voigt. Contiki - a lightweight and flexible operating system for tiny networked sensors. *EmNetSI*, nov 2004.
- [6] J. Hill, R. Szewczyk, A. Woo, S. Hollar, D. Culler, and K. Pister. System architecture directions for networked sensors. In *Proc. of ASPLOS-IX*, Nov. 2000.
- [7] A. Ledeczi, A. Bakay, M. Maroti, P. Volgyesi, G. Nordstrom, J. Sprinkle, and G. Karsai. Composing domain-specific design environments. *IEEE Computer*, pages 44–51, 2001.
- [8] A. Ledeczi, J. Davis, S. Neema, and A. Aggraval. Modeling methodology for integrated simulation of embedded systems. *ACM Transactions on Modeling and Computer Simulation*, 13, 2003.
- [9] S. Neema. System level synthesis of adaptive computing systems. *Ph.D. Thesis*, 2001.
- [10] P. Volgyesi, G. Balogh, A. Nadas, C. Nash, and A. Ledeczi. Shooter localization and weapon classification with soldier-wearable networked sensors. *5th International Conference on Mobile Systems, Applications, and Services (MobiSys)*, 2007.
- [11] P. Volgyesi, M. Maroti, S. Dora, E. Osses, and A. Ledeczi. Software composition and verification for sensor networks. *Journal of Science of Computer Programming (Elsevier)*, 56(1–2), 2005.