PART B

COLLABORATIVE PROJECT

Table of Contents

B1 Concept and objectives, novelty and foundational character, specific contribution to S&T, S/T methodology and work plan

B1.1 Concepts and Objectives

Overall Goal. This project proposes an energy-aware system development approach covering hardware, software and the run-time environment. The central goal is to make *energy usage transparent* through the system layers, thus enabling optimizations both during code development and at run-time.

The project work packages will develop novel program analysis and energy modelling techniques. Tools incorporating these techniques will enable energy optimizations both during code development and at run-time, helping to promote energy efficiency to a first-class software design objective. The project will also develop a concept of optimality and a set of benchmarks allowing measurement of energy efficiency with respect to the minimal energy achievable by optimal utilization of existing hardware.

Motivation and Novelty. Lack of energy transparency in today's system development tools means that much of the potential energy saving available from power-efficient hardware is wasted. The project departs from the approach of today's systems and development tools because energy transparency is at odds with a basic principle in modern software engineering – the desire to abstract away machine-level details in high-level code in the interests of portability, understandability and software reuse. By contrast, energy transparency requires making visible the effects of energy-saving features of modern processors. The project targets outcome (c) of Objective ICT-2011.9.8, namely to address software models and programming methodologies supporting the strive for the energetic limit (e.g. energy cost awareness or exploiting the trade-off between energy and performance/precision).

Key Techniques to Achieve Energy Transparency The realization of the project vision of energy transparency relies on innovation in two main areas.

- *Program analysis* concerns the derivation of energy information from software, which may refer to whole programs, individual procedures, statement blocks or lines of code, programs in high-level languages as well as machine code. Static analysis concerns deriving such information without executing the program, while dynamic analysis obtains information on energy usage from running programs. In general, dynamic analysis infers precise information for a finite number of execution traces, while static analysis infers approximated information for all possible execution traces. Thus, it is a major challenge to develop novel analysis techniques that bring together the best of static and dynamic approaches for each system layer and also to infer information that relates energy consumption with precision and performance.
- Advances in combined *hardware-software energy modelling* are needed to support effective program analysis. Traditionally, energy models are restricted to the level of machine instructions, and in particular to assign energy consumption values or ranges to individual machine instructions. We will develop methods of constructing energy models of high-level programming abstractions like blocks of code, functions, APIs and interfaces, building on combinations of semantic interpretation and information extracted at run-time. We will also study how to make higher-level models more precise incorporating information about the execution context, like execution history, and program and processor state.

Promoting Energy Efficiency to a First-Class Design Goal

- *Optimization: Towards Energy Optimality* A recent estimate by Intel states that energy savings by a factor of 3 to 5 could be achieved using software optimizations alone [22]. One of the keys to achieving this saving lies with equipping all stages of the software lifecycle with information on the energy-saving capabilities of the underlying hardware and the energy costs of performing various operations with that hardware. The project will establish the maximum energy saving for a given hardware/software combination and apply optimizations that push the energy efficiency of this system towards this limit. This will include annotations and development methods utilized by software engineers, transformations applied during compilation, and task management strategies used at run-time. This will enable software engineers to tackle energy-efficiency directly, rather than as a side-effect of other efforts. Figure 1 summarises the vision of optimality towards the limit, both for a given hardware platform and with respect to the theoretical limit.
- *Verifying Energy Budgets* Verification compares actual system properties with required properties in order to prove them or to detect inconsistencies. Being able to certify that a given energy budget is met, while maintaining a given level of quality of service (regarding precision or performance) is an important challenge. First, energy related information depends on different system and data parameters, so that a rich energy consumption semantics, that represents energy as a function of such parameters, is needed. Second, complex function comparison operations are needed in order to compare analysis information against specifications. The result of such comparison should ideally give preconditions under which a given specification is met or not. Another challenge is to find the causes of an energy consumption assertion violation, and help the system developer to redesign it.

Project Objectives

1. Establish energy modelling (WP2) and analysis techniques (WP3) that support energy transparency through the system layers.
2. Develop prototype tools enabling engineers to understand and quantify the impact of design decisions on energy (WP1).
3. Enable software engineers to express and verify properties of the system throughout its lifecycle relevant to energy (including precision and performance) (WP3).
4. Develop optimization techniques both at design and run-time enabled by energy transparency (WP4).
5. Develop energy benchmarking techniques (WP5) and a concept of energy optimality for a given application and hardware platform against which project case studies can be evaluated (WP6).

B1.2 Novelty and foundational character

Why achieving energy transparency is difficult Energy transparency is an ambitious target, since the natural trend in software engineering is to move software design further and further from the machine level towards more abstract concepts, with high-level languages, abstract data types and classes and reusable generic libraries, executed by layers of interpretation or compilation. This means that the level at which energy is actually consumed in a system – transistor activity and state preservation in silicon – is far removed from the level at which designers and system users usually think

and work. Static analysis of energy consumption of a high-level program thus involves either combining several stages of analysis, or analysing the program with respect to more complex semantics than usual. Similarly, dynamic energy analysis faces the inverse problem of interpreting measured energy at the compiled machine code level in terms of the high-level program source.

A further challenge is that energy usage is typically dependent on program input and other external factors including hardware parameters like processor clock speed; the results of static analysis of energy consumption are thus required to be parameterised by these external variables whose values might not be known at analysis time.

Energy efficiency as a first-class software design goal Achieving energy transparency throughout a system will open up the possibility of promoting energy efficiency to a *first-class software design principle and goal*. This promotion implies that desired energy usage can become part of the specification of an application, and can be considered right from the start of the design process and throughout all stages of software development. In particular, the possibility of formally verifying that a program meets an energy budget becomes feasible. At present, by contrast, energy consumption is measured in the final implementation, and if a program does not meet some energy target, the only choice is to reimplement the program and try again, potentially at great cost.

A whole-systems approach to energy management Furthermore, energy transparency will enable a range of run-time optimizations that are inconceivable at present, deciding dynamically how available energy in a system can be utilized optimally. Energy transparency enables to combine predictive models of energy usage of individual programs in the context of overall system usage.

As a rough analogy, it will be possible to apply ideas similar to those in "smart grids" to the energy management of a computer system. The run-time system of a software built using techniques from energy transparency will have access to models of energy-usage and can effectively exploit a range of dynamic controls and scheduling possibilities to reduce overall energy usage. While there are current research efforts directed at drastically reducing the energy consumption when a system is idling, this is only one approach to energy saving. The "whole-systems" enabled by energy transparency not only allows to increase idling periods and improve idle tasks, but also allows to design systems where energy management is improved by judiciously amortizing and smoothing the use of resources.

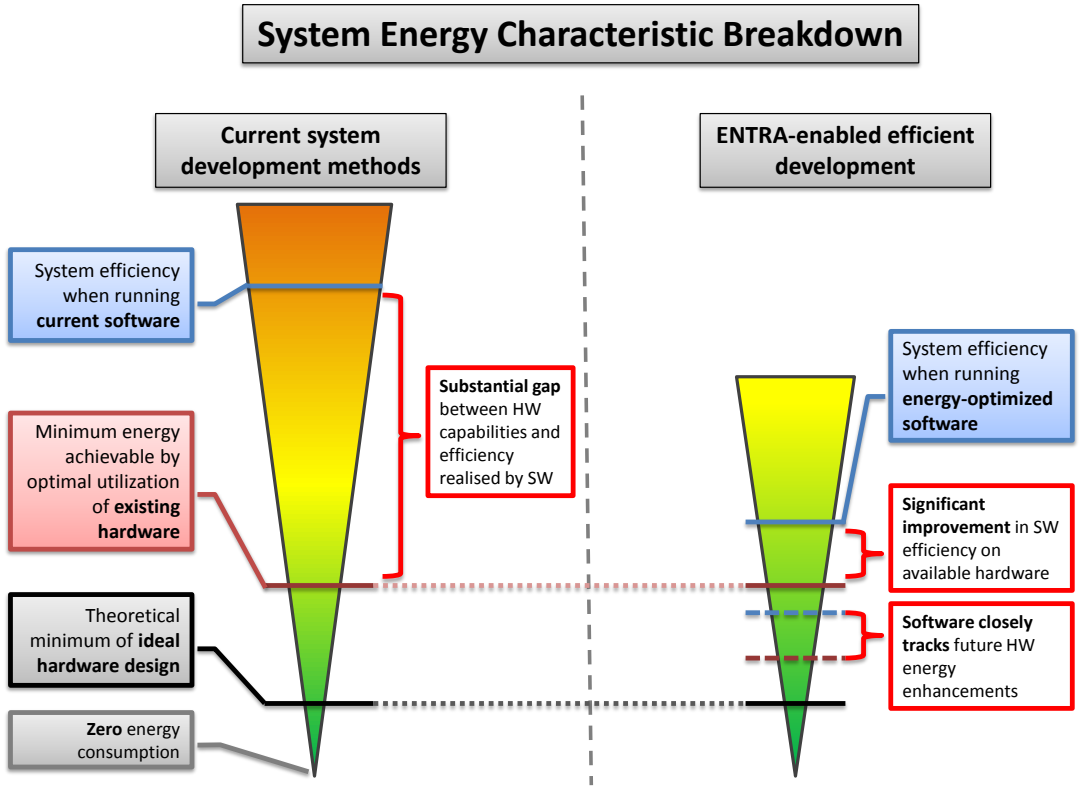

Figure 1: ENTRA Vision of Energy Optimization towards the Limit

B1.3 Specific contribution to progress in science and technology

In this section we discuss the main scientific challenges and advances needed to make energy-efficiency a first-class design goal.

B1.3.1 The goal of designing energy-efficient systems

Why is it a problem? Today's software engineering approaches produce systems that are far from energy-optimal, even though energy consumption is a major concern, especially in embedded systems. Traditionally, researchers and engineers work within one or perhaps two layers of the system stack with very limited overlap, e.g. software engineers, computer architects or hardware designers. Energy-efficient computing, however, is a challenge that requires considering the entire system stack vertically (as opposed to horizontally), from application software and algorithms, via programming languages, compilers, instruction sets and micro architectures, down to the design and manufacture of the hardware.

In an attempt to address this, various software-controllable energy-saving features have been developed, such as DVFS, where power modes can be selected via embedded software drivers or application software. However, many if not most applications, especially if designed in high-level languages, are simply unable to take advantage of these features and so the application will not see any low-power benefits. The designer of the software has no way to understand or exploit the available controls.

While hardware can be designed to save a modest amount of energy, the potential for savings are far greater at the higher levels of abstraction in the system stack [22]. The greatest savings are expected from energy-efficient software. Addressing the challenge of energy-efficient computing requires energy transparency over the entire system stack.

B1.3.2 Is energy transparency possible?

The essential requirement is that energy usage should be visible at the level at which software is designed or used. One could imagine an over-simplified view of energy transparency as an interactive development environment that colours source code according to how much energy it consumes – red for the energy-expensive code and green for the "cool" parts. Despite the over-simplification we can use it to identify the essential components of a potential solution that enables the energy consumption at hardware level to be immediately visible at source code level. These components are:

- analysis of programs with respect to energy usage, performance and precision;
- energy modelling at low and high level;
- a common assertion language as an integrating framework;

Program analysis. Static analysis of source code derives information about run-time behaviour without actually running the code. Among other run-time properties, analysis of resource usage of computations has been the subject of previous research. Typically, execution time and memory usage are analysed, but more recently analysis of other resources including energy, bits sent or received, number of accesses to a database, or monetary units spent have also been the subject of static analysis. Furthermore, information from static analysis plays an increasingly central role in optimisation, debugging and verification tools.

Abstract interpretation as a generic analysis framework. The set of possible computations of a program is usually extremely large, and therefore practical automatic analysis tools have to make approximations to program behaviour. A systematic approach to static analysis has been developed over the past 30 years, based on the principles of abstract interpretation [16]. Abstract interpretation provides an approach for implementing a sound approximate execution of a program from a specification of the concrete semantics of the language of the program together with a so-called abstraction function. The abstraction function typically abstracts away things that are not relevant to the analysis in question. For example, in an analysis focussing on energy usage, an abstraction function would abstract away aspects of a state not relevant to energy consumption.

Abstract interpretation has been applied to programs ranging from assembler code and embedded systems code through to high-level declarative languages and specification languages. Tools based on abstract interpretation are being applied in industrial applications (see e.g. [20]).

Time and complexity analysis. The problem of automatically estimating the amount of time a program will take to run, in relation to the size of the input, is sometimes called complexity analysis or cost analysis and will form a basis for energy analysis. Techniques exist for inferring the complexity in terms of (an upper bound on) the number of execution steps, reductions, resolution steps or function calls (depending on the programming paradigm) of a given program. Often, this involves analysing a step-counting program transformed from the original program. The established approach is to set up recurrence relations and obtain solutions of them as closed form complexity cost functions. Relevant previous work includes [18, 17, 19, 54, 9, 28, 17, 18, 19, 49, 32, 2, 1]. It is notable that the results of such analyses are *parameterized* expressions, that is, they express the number of steps as a function of input data size.

In some applications step-counting is not sufficiently precise, especially where there are real-time constraints, and accurate timing information is needed. A number of static analyses are aimed at worst case execution time (WCET), usually for imperative languages in different application domains. These have produced precise timing models (see, e.g., [91, 87, 8, 23, 39, 8] and their references). However, in contrast to cost analysis, these and related methods do not infer cost functions on input data sizes but rather absolute maximum execution times, and they generally require the manual annotation of loops to express an upper-bound on the number of iterations. A first attempt to combine dependency of WCET and data size for PIC programs was reported in [34].

Energy analysis. Formal models of energy usage are not so well-developed compared with timing models, especially in high-level languages. Some work has been performed for the energy consumption of Java Bytecode programs, based on a simple low-level model of the energy consumed by the execution of each Java Bytecode. The analysis techniques are similar to those for complexity analysis. Relevant work includes [65, 66, 64, 48].

Other approaches have been developed for estimating the energy consumption of programs (e.g. [83, 48]) based on measuring actual consumption at run-time for large sets of random inputs. This approach has the limitation that it applies only to the class of programs that have been measured.

Energy analysis of parallel programs. The energy analysis of parallel programs is of particular interest, since the widespread use of multi-core processors is partly driven by the need to control the energy-performance trade-off. Korthikanti et al. present in [46] an analysis of the energy consumption of a number of particular parallel algorithms executed on shared memory multi-core processors, which is used to establish the optimal number of cores to minimize the energy consumed by a given algorithm for a specific problem size while satisfying a given performance requirement. The analysis is not automatic, and it is a challenge to develop automatic analyses. In [45, 44], the authors argue that examining the relation between the performance of parallel algorithms and their energy requirements on multi-core processors may be facilitated by analyzing some scalability metrics. For a given problem instance (i.e., the execution of a given algorithm for a given input size), and a fixed performance requirement, *energy scalability under iso-performance* provides the optimal number of cores required to minimize the energy consumption.

Energy analysis of precision. The precision of computations is a highly relevant factor when considering energy optimizations that trade energy for precision. Floating point computations break most algebraic properties of numbers (such as associativity or commutativity), and introduce rounding errors at each step of the computation. This introduces problems for program analysis. Precision analysis must take into account the floating point computation norm [40]. Some analysis by abstract interpretation can take such rounding errors into account [15], but the rounding error is always assumed to be a worst case, and although the analysis is quite efficient, it only gives bounds on the output, not the imprecision of the computation. Other analyzers based on abstract interpretation [30, 60] give sound bounds on the difference between an idealized computation using real numbers and the actual floating point computation. It is even possible to compute program transformations to improve the precision with respect to the real number semantics [59].

Other analyses. Other resources relevant to energy consumption will be considered. A general framework for analysis of user-defined resources is defined in [66]. It infers platform-dependent resources by including a one-time profiling phase of a given platform in order to determine the values of certain parameters for that platform. The framework has also been particularized to the estimation of execution times in [62, 61]. The approach in [61] is applicable to programs running on a bytecode-based abstract machine.

Static analysis for energy use can be directed also at subsidiary properties that are critical for understanding resource use in a program. Examples of such properties are patterns of use of shared

resources such as memory and communication channels, and the use of energy-expensive operations such as communication. This is a relatively unexplored area for static analysis; a recent relevant work is [58].

Energy modelling at low and high level Energy models will form the basis of energy analysis techniques and will allow energy use of a program to be estimated without necessarily running it on existing hardware.

Low-level modelling of energy. Estimated energy use of a design can be obtained based on simulating its netlist using the program as simulation stimulus (test vector). The low-level energy model underlying such simulations are toggle counts that can then be characterized with respect to energy consumption based on the fabrication process data provided by the chip manufacturer (this typically is closely protected Intellectual Property). While the latter approach does not require the hardware to be manufactured, allowing energy consumption values to be obtained early in the system development process, the associated simulations are computationally extremely expensive. For example, simulating 10M cycles on a 10M gate design can easily take a week (if not two) even if running on a huge server farm. Significant reductions, from weeks to hours, can be obtained using expensive dedicated emulation hardware.

Various approaches have been developed to establish power dissipation models for hardware at abstraction levels above netlist. Typical examples include models at the micro architectural cycle-accurate level [11] and the instruction level [88]. The latter links the software to the hardware as the Instruction Set Architecture (ISA) is the target of the compilation process, which translates a source code program into an executable sequence of instructions, also referred to as assembler code. For software this is typically the "lowest" level of abstraction, while for hardware design this is a "high" level of abstraction. In fact, the ISA serves as the functional specification in the hardware design process.

Higher-level energy modelling. More recent approaches [71] promote using several levels of abstraction, effectively providing multi-granularity models which offer designers a choice in trading off accuracy versus simulation speed. These models most typically are integrated into the design-flow in order to speed up simulation time with reasonable success. However, *formal analysis* of such higher-level models, used instead of simulation, can completely eliminate or more significantly reduce effort currently invested into simulation. It should be noticed that in reality, the base cost of an instruction can vary depending on the value of operands and, more generally, depending on the state of the processor at the time the instruction is executed (the execution history). Further work on incorporating such dependencies into the model are required.

Assertion language integrating analysis and modelling Energy models must be given a suitable representation within analysis algorithms. Models can take many forms, from simple lists associating instructions with a fixed energy cost, to functions applied to a whole section of code and taking various parameters expressing input data and execution environment.

At the same time, results derived by analysis have to be represented in a form both suitable for manipulation in the analysis algorithm and for interpretation by the user. Information on energy usage derived by analysis can also be compared with user-supplied assertions, and thus the results should be in a form readable by designers.

This takes the form of an *assertion language* that binds together models, analysis results, and user assertions.

Various available assertion formats. For hardware design, where power has long replaced performance as a prime design goal, the Common Power Format [84] provides hardware designers with a file format

to specify power-saving features of the design. These specifications are then used by synthesis and verification tools. This permits hardware designers to express the design intent with respect to power saving at design time. The use of assertions to express functional hardware properties to be verified either formally or during simulation is now standard practice in industry [25].

At the software level, assertion languages to express program properties to be checked statically or at runtime have been proposed for several declarative and object-oriented programming languages (e.g., Eiffel [63], Racket[24] and Ciao [35]). Assertions allow the users to express contracts between program components, typically talking about functional properties, and are the basis for the so called "design by contract" method [63]. The use of a multi-purpose common assertion language for integrated program verification, analysis and optimization has been proposed in the context of a multi-paradigm (declarative) language [35], which also allows to express some non-functional properties.

Currently, such languages do not permit non-functional properties, such as timing and energy budgets, to be integrated into the code from the outset. That is why, for software, timing and energy are a consequence of the design rather than a design goal. XC [89], the language designed to best utilize the features of the XMOS XCore processor, provides a mechanism via pragmas for programmers to annotate code with timing requirements, which can then be checked for single threaded programs only using the XMOS Timing Analyzer [93].

To lift timing and energy requirements to first class design goals for software development, a suitable assertion language must provide programmers with expressive constructs to specify these constraints at the same time as the code is being developed: this language must be integrated with the language in which energy models and analysis results are represented.

B1.3.3 What optimizations are enabled by energy transparency?

Program transformations are studied in many contexts. In this project the aim is to optimize with respect to energy consumption. In general the task is twofold: firstly to recognize the opportunity for a transformation, namely the semantic conditions that allow one piece of code to be replaced by another, and secondly the selection or generation of the new code that replaces the old, reducing energy consumption.

The analyses discussed above will be used to identify opportunities for optimizing transformations. Typically the conditions are obtained from a global analysis, establishing that some condition holds in all (a possibly unbounded number) of possible executions of code.

Program Specialization and Partial Evaluation. In many applications, especially when code is developed in high-level languages or from library modules, code is executed only in a restricted context, with constraints arising implicitly from the initial state and the call structure of the program. Detecting such constraints is in general a non-trivial task requiring a global analysis to propagate the effects of input restrictions through the program. *Specialization* of code with respect to constraints can enable redundancies to be removed, with considerable savings of time and energy.

Partial evaluation and related specialization techniques using statically known input have been researched since the 1970s and numerous tools have been developed for a variety of programming languages. A bibliography of applications can be found at `http://readscheme.org/partial-eval/`. In particular, tools have been developed for imperative and object-oriented languages [27, 14, 82], for logic languages [52, 70] and for functional languages [33, 42, 13]. Frameworks integrating specialization with program analysis have been defined [74, 50], and program analyses aimed at improving specialization have also been developed, e.g. [26].

Specialization of interpreters. Recent work in both analysis and specialization makes use of *metalanguage reflection*, allowing an analyser/specializer for one language M (the metalanguage) to be applied

to programs in another language L (the target language). This approach has been used for example for partial evaluation of Java bytecode [29] and PIC assembly code [34] both of which used a partial evaluator for constraint logic programs. It has also been applied for analysis of high-level specifications in Petri nets [53] and analysis of linear hybrid automata [6]. It is now also well understood how to obtain optimal specialization of interpreters [51, 7]. It is expected that this interpretation technique will be adapted in the project for handling higher-level languages and execution strategies.

Local optimizations. A useful class of optimizations, that is common in low-level but not high-level code, is the replacement of some code by equivalent code that is more power-efficient. At the machine level, lookup tables typically record short sequences of assembly language instructions that can be replaced by equivalent more efficient sequences. These ideas could be extended to high-level code, where individual or sequences of library function calls could be replaced. Such tables can also be accumulated by a programming environment, to save inferring the same optimizations many times.

Parallelism - transformations and analyses. The exploitation of parallelism to reduce energy consumption is receiving increasing attention and will be researched in the project. The general idea is to split sequential tasks into independent subtasks that can be executed by different cores running at a lower frequency, yielding lower net energy use for the same task. For example, the energy consumption analysis proposed by Korthikanti et al. [46, 45, 44] can be used to determine the optimal number of cores to minimize the energy consumed.

It is important to ensure (functional) correctness of parallel execution (in the sense that the parallel program must behave the same as the corresponding sequential program) [12, 36]. This is often a matter of proving the independence of the subtasks using static analysis. It is also necessary to ensure (practical) efficiency [57, 56]. For the latter, we must take into account the overhead associated to parallel execution (e.g., due to communication, or parallel task creation) and the one associated to changing clock frequency in order to ensure that such costs do not exceed the benefits of parallel execution.

Trading off Quality-of-Service against energy Optimizations that save energy, while possibly losing precision, performance or some other quality-of-service measure will be studied. Some techniques from the literature exist and will form the basis of research in the project, augmented by energy analyses. PowerDial [37] is a state-of-the-art *autotuner* that converts static configuration parameters that already exist in a program into dynamic "knobs" that can be tuned at run-time. This technique is automatic but requires the user to provide training data, an output abstraction and identify a set of parameters (static configuration parameters) in the program.

Other approaches use program transformations that indirectly and dynamically reduce energy consumption. Hoffmann et al. proposed SpeedPress, a framework designed for exploiting performance-accuracy trade-offs using a technique called loop perforation [79, 80, 81]. Loop perforation transforms the program loops to execute only a subset of the iterations in the original computation. This technique is automatic but requires the user to provide training data, an output abstraction and bounds on output distortion.

Similarly, a number of other mechanisms have been developed for dynamically varying application behaviour to maximize performance subject to an accuracy constraint or vice versa. Specific mechanisms include multiple selectable implementations of a given component or components [4, 5, 95]; sampling inputs to reductions [95], skipping tasks in parallel programs [75, 76]; approximate function memoization (returning a previously computed value when the arguments of the new function call are close to the arguments of the previous function call) [77]; and approximate data types (data types that return approximate results for operations) [78].

Power/Energy management at OS level(DVFS) with scheduling policies
At the level of the Operating System (OS), energy efficiency techniques based on *dynamic voltage and frequency scaling* (DVFS) have been explored in the research community [3, 85, 90].

In [90], models were developed based on parameters such as memory requests per cycle and instructions per cycle, which can be counted using the performance-monitoring unit (PMU) available in most processors. By predicting a workload's response to a change in frequency, a more energy-efficient frequency can be chosen at each scheduler invocation. Snowdon [85] enhanced this approach by developing a technique to choose automatically the best model parameters from the hundreds of possible events that modern PMUs can measure. His framework uses predictive models for execution time and power consumption. These provide an estimate of the energy consumed by the CPU when executing a task. Choosing parameters and their coefficients requires a one-time model-characterization on the system that is being modelled.

In [3], the optimal frequency selection is considered on multicore processors in the presense of resource contention, since cores of a chip share resources such as caches and memory interfaces. Combining tasks that run best at a certain frequency does not pay off if it leads to resource contention. The work analyzed scheduling for avoiding resource contention and for optimal frequency selection. It argues that the two are opposing goals, and that scheduling to avoid resource contention is crucial both in terms of performance and energy efficiency. This work proposes scheduling policies, using activity vectors for representing resource utilization, that reduce contention significantly by co-scheduling tasks with complementary demands.

B1.3.4 What are the implications for tools and methods?

Verification Verification is the process used to demonstrate the correctness of a design with respect to the requirements and specification. Traditionally, verification has concentrated on establishing functional correctness. In the hardware community, the scope of verification has recently been extended to include non-functional properties, i.e. verifying that a design satisfies specification constraints beyond functional correctness, such as low-power features [84].

Verification of energy-related properties. Energy transparency opens up the possibility of tools that allow the software engineer to verify assertions about energy usage. Verification based on static analysis, particularly techniques based on abstract interpretation, consists essentially of comparing derived properties against specified or desired properties. Even though static analysis in general derives safe (over-)approximations of program behaviour, many typical properties, especially safety properties, can be proved. Most work on verification based on static analysis focusses on the proving of conformance to functional properties and typically results in the outcome true/false/unknown. However, in the case of properties such as resource usage, it is necessary to extend analysis to produce results which include properties such as upper- and lower-bounds on execution time and usage of energy, memory, or other types of resources. For example, it may be possible to say that a certain result holds if the input data size is in a given range.

The consortium has experience in the development of state-of-the-art frameworks for (static) verification of general resource usage program properties, based on static analysis [55, 35, 10]. Although initial results in this topic are quite promising, there are however open challenges that remain to be solved, as for example, dealing with richer resource-usage requirements (and hence analysis), and techniques for comparing analysis information against specifications in order to prove them.

A related problem is debugging of energy-related "bugs". Some previous work on debugging of resource usage properties exists [72, 35] and will be extended in the project.

Energy as a first-class design goal. There is already evidence that energy usage is becoming a major software design goal in some areas (see for example articles in the ISLPED conference series `http://www.islped.org/`). The areas in which some success has been achieved are either where design is close to hardware, such as embedded systems, or where there are very specific use-cases and computation patterns, as in data-centres.

Energy transparency will enable us to extend the range of applications in which energy is a first-class design goal. Progress will depend on the precision and scalability of energy analyses, and on the quality of the energy models available.

B1.3.5 Further research beyond the project

The project is expected to establish the feasibility of energy transparency, implement prototype tools and investigate the optimizations enabled on a range of case studies. Many areas of future work will still exist. Among these are:

- Hardware design for energy transparency. How to design hardware to facilitate analysis of energy usage and hence energy transparency.
- Deeper transformations and optimizations. As multicore processors proliferate the problem of automatic parallelization will become more urgent in order to achieve potential energy savings. Although this project will investigate some parallelizing transformations, many more challenges will remain.

B1.3.6 Project Strategy and Research Areas

The preceding discussion identified key areas of research relevant to achieving and exploiting energy transparency. The following summary outlines the approach to a number of research problems to be taken in the project, and the associated work packages.

XCore architecture. Many techniques researched in the project will be processor-independent, but in some cases we need to restrict ourselves to a single architecture, for example when modelling an ISA. We have chosen to use an XMOS XCore for those experiments because it is a scalable multicore, can run different cores at different speeds, and is event-driven at architecture-level, enabling power-efficient implementations.

- **Energy Modelling (Low level): Work Package 2.** ISA-level models will be established for the XCore, initially associating instructions with average values. A significant advancement of the state of the art with respect to modelling will take advantage of the fact that formal analysis can also be performed on the basis of cost *functions*, rather than just cost values. Such cost functions can capture dependencies that influence the energy consumption of an instruction (such as operand values or execution history). The consortium has expertise in formal modelling and verification of designs, including designs of industrial complexity, at different levels of abstraction [94, 92, 43, 21, 41, 38].
- **Energy Modelling (High level): Work Package 2.** The modelling of energy usage at higher levels goes beyond the state of the art, and will be achieved using semantic mapping techniques on the one hand and dynamic profiling on the other. Explicit interpretive models of higher-level languages and execution models (including parallelism) will be used to lift low-level energy models to more abstract structures.

- **Common Assertion Language: Work Package 2.** The project will advance the state of the art by defining a common assertion language for expressing energy and timing properties in languages used by the embedded software design community. In particular, this language will allow the expression complex specifications including energy consumption functions which depend on data properties (such as data size), other environmental properties (such as clock frequency and voltage) and inter-module contracts. The assertion language will be multi-purpose, used for analysis, optimization, and verification and debugging of embedded systems, allowing (energy) information flow between different components of the integrated tool set.
- **Energy analysis: Work Package 3.** We will focus on the continued adaptation of advanced cost and WCET analysis techniques to energy use, and on development of techniques for inferring upper- and lower-bounds on energy use. Furthermore, we will integrate timing analysis for parallel programs with energy use analyses to derive energy use for multi-core applications. Existing and new techniques for energy analysis will be applied to the main target languages used in the proof-of-concept for the project, namely, the XC language and XCore assembler. Our research will make advances in integrating analysis of energy use in parallel programs with analysis of access of shared resources. All three aspects will also advance the state of the art in identification and classification of applications with respect to scalability of analysis.
- **Verification and debugging of energy properties: Work Package 3.** Our work on verification focuses on developing novel automatic verification and debugging techniques and tools based on static analysis that support rich specifications about energy usage (as well as timing and precision) properties of programs, and apply them to industrial languages. We will focus on static verification, and will produce novel techniques for the comparison of analysis information against specifications involving energy, timing and precision properties. To complement verification, we will produce novel techniques for debugging the energy usage of programs.
- **Program transformations: Work Package 4.** In the project we will further develop resource-aware specialization in general and energy-aware partial evaluation in particular. In particular for energy-aware specialization, new empirical insights will be gained and new techniques developed. These techniques will further tighten the link between abstract interpretation and program specialization. We will also explore parallelization techniques to automatically derive energy efficient parallel versions, while giving support to the integration of both, manual and automatic techniques. The analysis results of parallel programs will be applied to exploit XC and XCore language features to parallelize code.
- **Dynamic program optimizations: Work Package 4.** Energy transparency will allow a radical re-interpretation and extension of techniques available in the literature, such as the techniques discussed earlier for trading off energy usage against precision or speed of computation, and run-time scheduling which takes account of global energy consumption of a set of energy-consuming tasks.We will explore integration of dynamic transformations such as autotuning and code perforation with advanced energy models developed in Work Package 2.
- **Benchmarking and Optimality: Work Package 5.** A new notion of optimality will be defined and explored, based on the idea of the minimal energy achievable by optimal utilization of existing hardware. This will provide a pragmatic and realistic approach to benchmarking and is adaptable to all platforms.

B1.3.7 Related Projects

ASAP: Advanced Specialization and Analysis for Pervasive Computing (2002-2005) was an EU FP5 FET project (38059) to automate as much as possible the development of sophisticated

and reliable software for pervasive computing platforms using high-level languages and analysis, verification, and specialization techniques. The results included a novel toolkit based on the ideas above available as open-source code. Partners included the precursor to IMDEA, University of Bristol, HHU Düsseldorf and the University of Roskilde.

SAFT: Static Analysis with Finite Tree Automata. 2007-2010. Funded by Dansk Forskningsråd for Natur og Univers (FNU = Danish Natural Science Research Council) under grant FNU-272-06-0574. Partner: RUC.

MERIT/COMVERS: Resource-Aware and Verifiable Mobile Computing. 2005-2009. Funded by the Spanish Ministry of Science and Education (MEC) under grant TIN2005-09207-C03-01. Partner: IMDEA.

ES_PASS: Embedded Software Product-based ASSurance. 2007-2009. EU ITEA2 cluster of EUREKA Program, Project number 06042, Spanish PROFIT grant FIT-340005-2007-14. Partner: IMDEA.

DOVES: Development Of Verifiable and Efficient Software. 2009-2013. Funded by the Spanish Ministry of Science and Innovation (MICINN) under grant 2008-05624/TIN. Partner: IMDEA.

Numeric and Symbolic Abstractions for Software Model Checking. 2011-2013. Funded by Dansk Forskningsråd for Natur og Univers (FNU = Danish Natural Science Research Council) under grant FNU-10-084290. Partner: RUC.

Formal Modelling and Analysis of the XCore Instruction Set Architecture Design. 2010-2011. Funded by the Knowledge Transfer Secondment Scheme of the UK Engineering and Physical Sciences Research Council under grant number EP/H500316/1. Partner: UNIVBRIS and XMOS.

Understanding the State of the Art in Power-Aware System Design. 2010-2011. Funded by the Royal Academy of Engineering under agreement number 10465/258. Partner: UNIVBRIS and XMOS.

B1.4 S/T Methodology and Associated Work Plan

B1.4.1 Overall strategy and general description

ENTRA aims to make energy usage transparent throughout a system, promoting energy efficiency to a first-class goal in design and system management.

The project strategy can be visualized as in Figure 2. The central project goal of *energy transparency* relies on scientific advances in two main areas: modelling of energy usage from silicon through to high-level language constructs (WP2: Energy Modelling Through the System Layers) and program analysis, both static and dynamic (WP3: Analysis and Verification).

Achieving the goal of energy transparency enables the project to study novel program optimizations that reduce energy usage at both design-time and run-time (WP4: Optimization) and to investigate how energy transparency in tools and methods enables energy-aware software engineering (WP1: Energy-Aware Software Engineering).

Benchmarking suites (WP5: Establishing the Benchmarks) and design case studies (WP6: Evaluation) will be developed in the project to support the optimization and evaluate the software engineering work packages respectively. Finally, WP7 (Dissemination, Collaboration and Exploitation) covers the consideration of how to present the project results to a wider audience and initiate follow-up activities, and WP8 (Management) covers project management.

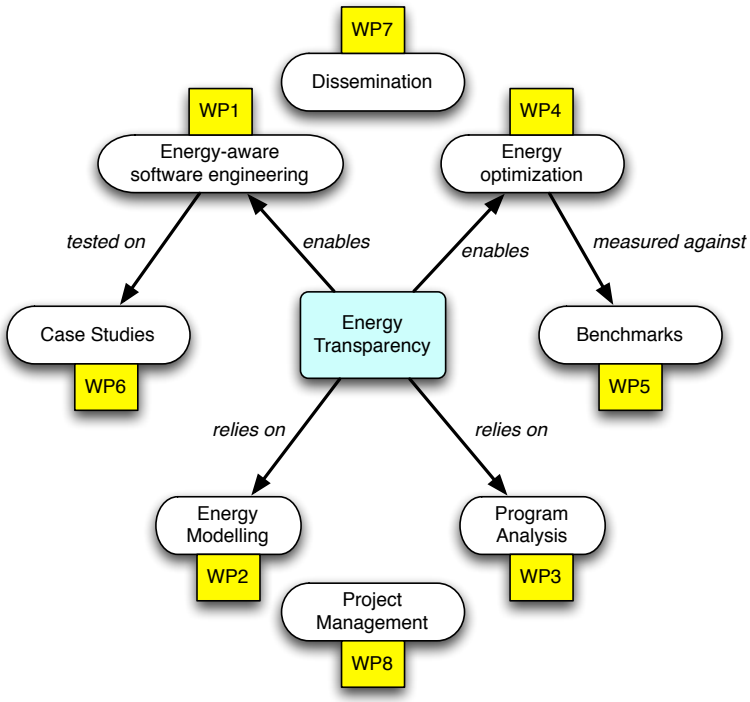

Figure 2: ENTRA Work Packages

Project Stages The project strategy is to follow four stages that successively integrate the fundamental components of energy transparency. A complete list of milestones, including those concerned with Management and Dissemination, can be found in Part A. Some additional description is added below, explaining how it is established that each (technical/scientific) milestone has been reached. The deliverables are described in the list of deliverables of Part A.

1. Define a common *assertion language* that integrates the fundamental components of energy transparency, namely program analysis, energy models and the software engineer's view of the energy usage of a program. This stage concludes with **Milestone 1** at Month 12. This milestone will be established by scenarios showing the use of the assertion language described in Deliverables D2.1 and D3.1, in which energy models and analysis results are expressed in the common assertion language.

2. Energy modelling and program analysis tools will be integrated, leading to a first energy analysis prototype. This stage concludes with **Milestone 2** at Month 18. This milestone will be reached when the lower layer energy models of WP2 are connected to the higher layer analysis of WP3, and the first energy consumption estimations for XC programs are inferred (Deliverables D2.2 and D3.2).

3. Optimization prototype tools will be applied based on the energy analysis, and results compared with idealized benchmark optimizations. This stage concludes with **Milestone 3** at Month 24. This milestone will be established at a project working meeting where a prototype optimization tool will be demonstrated on at least one of the project case studies.

4. The initial experiments and prototype tools will be extended, results evaluated, and advanced topics such as energy contract verification and dynamic scheduling will be investigated. This stage concludes at **project completion** at Month 36.

The work packages have been identified to address the specific objectives to be achieved by the ENTRA project. Individual work package effort allocation is based on a cooperative estimate of the needed person months, jointly made by the consortium members during project preparation and based on their practical experience from collaborations in previous joint projects. Every work package has a work package leader who will be responsible for the work package progress, deliverables and milestones. The work package leader reports to the project coordinator. Work package leaders have been nominated according to skills and experience of partners. Each task is labelled with the participating project partners. The heaviest work packages in terms of effort are WP3: Analysis and Verification and WP4: Optimization, since these require significant advances over the state of the art.

Dependencies Figure 3 shows the key task dependencies. These tasks are described in the work package description table of Part A, and their timing is given in Section B1.4.2. In some cases the dependency links are labelled by a deliverable, in particular where the dependent task receives input from an ongoing preceding task via some intermediate report or prototype. Apart from the defined dependencies, regular communication will ensure flow of ideas and feedback from one WP to another. The consortium is small and cohesive, and the project team leaders will foster a culture of communication and collaboration, leading to greater exchange of ideas and progress reports, especially at quarterly project meetings. In particular, tasks in WP2, WP3 and WP4 will continuously obtain feedback from each other.

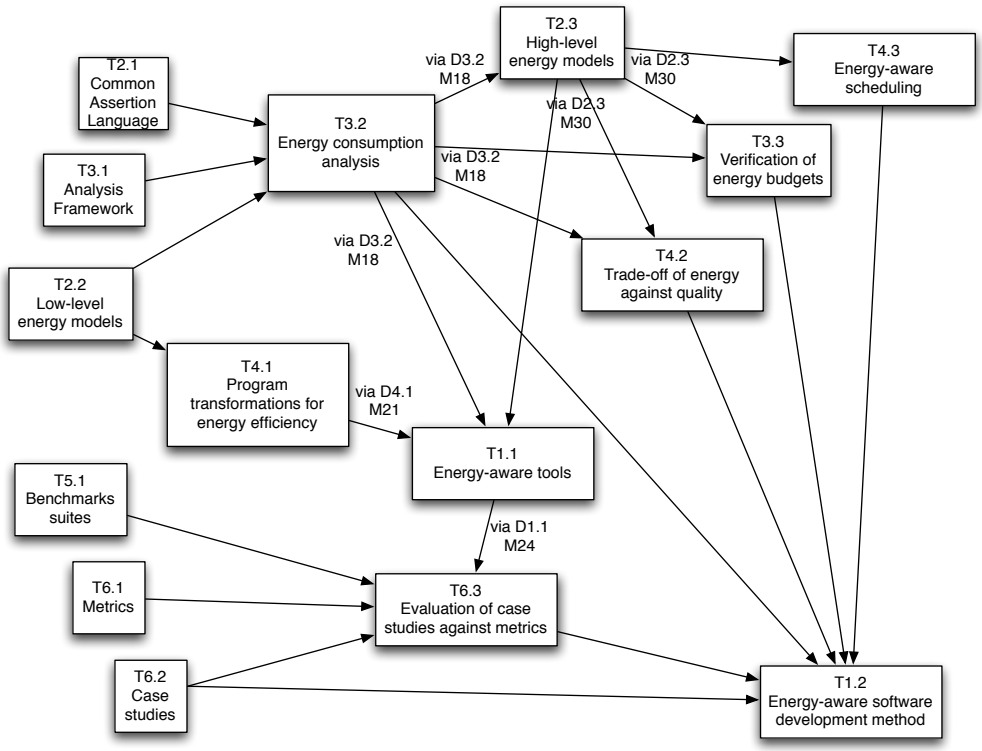

Figure 3: Dependencies

Risks and contingencies Risks can be classified into technical risks and organisational/commercial risks. Risk management of the project is based on assessing the risks and planning preventative actions (actions taken a priori to reduce the chance of a risk happening) and ameliorating actions (actions to be taken if a risk does occur to reduce its effect).

Technical risk is part of the nature of adventurous and innovative research. Organizational risk is inherently present in projects that involve international collaboration. The strategy in ENTRA is to reduce the likelihood of a risk occurring and to put measures in place to significantly limit the impact of the risk in the unlikely event that it does occur.

The following sections list the technical and organizational risks identified by the ENTRA consortium together with the preventative/ameliorating actions planned.

Technical Risks

- **Risk of not meeting the milestones:** Risk is associated with not meeting the outputs specified for the 3 project milestones.

  Milestones have been placed strategically so that task outputs are available to reach the milestone. Partners are aware of the milestone objectives and of their importance to reach the next phase of the project. The work is structured in such a way that it permits alternative paths to reach technical objectives, this significantly reduces the risk of not meeting objectives.
- **Risk of techniques and methods not being mature enough to be adaptable:** ENTRA is primarily focused on providing proof of concept solutions and on assessing

viability. XMOS will assess the techniques and methods during the project, especially during evaluation and adapt what is deemed to be mature.

- **Risk of techniques and methods not being mature enough to be integrated into a fully working prototype:** There is a high technical risk that not all of the tools developed will be mature enough to fully integrate into the toolset at the end of the project. In this case, the prototype will be more limited in that the output of a phase may need manual processing before it is passed to the next phase. Thus the techniques and methods will still be demonstrated, but not as automatically (and hence user friendly) as ideally desirable.
- **Risk of complexity and related scalability of models:** There is a risk in the scalability of models - single core models are low risk but multi-core models are a higher risk; having a progression from one to the other reduces the overall risks.

Organisational etc. Risks

- **Staff recruitment and associated delays in project start:** There are a number of potential problems in recruiting staff especially in academic organisations. Academic organisations are frequently constrained by not being able to recruit until contracts are in place, which may lead to staff not being hired in time for the start of the project.

  In ENTRA, a lower effort profile has intentionally been planned at the start for these partners for the first few months to allow for this, and thus avoid tasks, especially those with tight dependencies, starting late and causing project delays from early in the project.
- **Change of interests:** There are always risks in partners changing their focus and the research no longer being aligned fully with their commercial interests.

  In ENTRA we have reduced this as far as possible by having an industrial partner for whom the results of this project are a key to their core business offering.
- **Failure to come to an Consortium Agreement:** The partners may not reach consensus over the Consortium Agreement.

  This is a very small risk because several of the partners in ENTRA have previously worked together in collaborative projects. During proposal writing partners agreed on the DESCA FP7 consortium agreement as a basis. The actual agreement will be agreed upon and signed at the start of the project.
- **Financial viability risk:** A small risk is associated with partners not being able to finance their involvement.

  The academic partners are mature institutions and financially stable. XMOS has been in business for more than 7 years and is financially sound and in a position to cover its commitments in ENTRA.

B1.4.2 Timing of work packages and their components

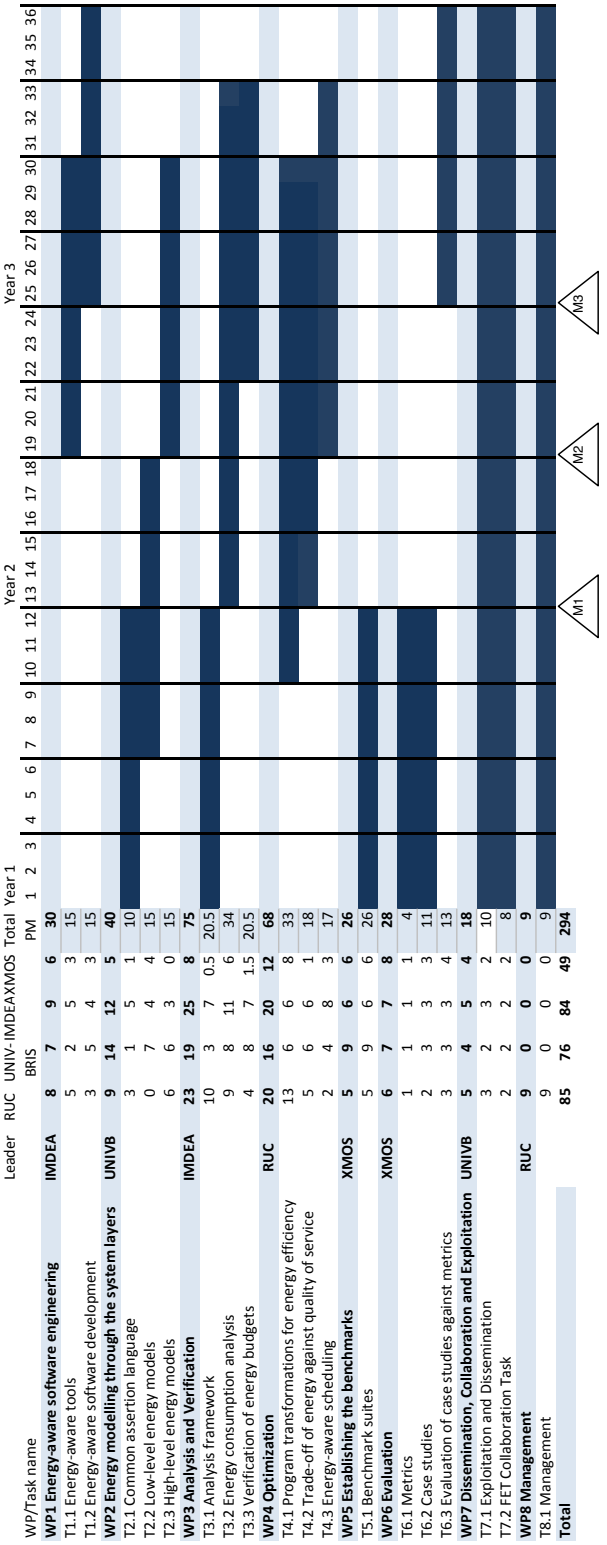

Figure 4: Gantt Chart including Work Packages, Tasks and Milestones

B2 Implementation

B2.1 Management Structure and Procedures

The general management of the project is based on the following principles:

1. Management should be consistent with the size of the project.
2. Decision making should be quick and simple, but respect all partners' interests.
3. Decision making should be by consensus as far as possible.
4. Where consensus cannot be reached, a simple process should be used.

A **Consortium Agreement** based on the DESCA FP7 model (`http://www.desca-fp7.eu/`) defines the details of the project management structure and procedures and the handling of financial matters. The main points from this agreement are summarised in this section but the Consortium Agreement is the definitive reference.

Management consists of the following formal "actors":

- The Project Coordinator (RUC).
- The Project Coordination Committee.

Project Coordinator: The Project Coordinator (RUC) (whose team leader is referred to as the Project Manager below) is responsible for the overall management and coordination of the entire research project. A special emphasis within its responsibilities is to ensure in cooperation with WP Leaders the overall integration of the individual WPs and also to chair the project board, the Project Coordination Committee. Furthermore, the Project Coordinator is the only official channel through which interaction with the European Commission takes place, especially with regard to the submission of deliverables, aspects related to third parties and the consortium.

Project Coordination Committee: The Project Coordination Committee (PCC) consists of one representative of each beneficiary. The PCC will meet at least once a year and at any time upon written request by one or more beneficiaries. The PCC is the highest decision board and its main task is project governance. It will have the overall responsibility for all technical, financial, legal, administrative, ethical, and dissemination issues of the project. For this reason, the PCC will monitor and assess the actual progress of the project and make amendments if necessary.

In particular the PCC will be responsible for the following tasks:

- Content, finances and intellectual property rights
  - Proposals for changes to Annex I of the EC-GA to be agreed by the European Commission;
  - Changes to the Consortium Plan (including the Consortium Budget);
  - Changes to the definition of and access rights to Background.
- Evolution of the Consortium
  - Entry or withdrawal of parties to or from the Consortium and approval of the conditions for such changes;

- Declaration of a Party to be a Defaulting Party and consequent measures to be taken;
- Proposal to the European Commission for a change of the Coordinator;
- Proposal to the European Commission for suspension of all or part of the Project;
- Proposal to the European Commission for termination of the Project and the Consortium Agreement.

Voting Rules and Quorum: Decisions of the PCC require the presence or representation of two-thirds (2/3) of its Members. Each Member shall have one vote. Defaulting Parties may not vote. Decisions shall be taken by a majority of two-thirds (2/3) of the votes.

Responsibilities of Beneficiaries: Each beneficiary will nominate a local team leader who will be responsible for the planning, execution and controlling of that beneficiary's activities in the project, in compliance with the decisions of the PCC. The beneficiary team leader's responsibilities are:

- Provision of interim reports on project resource usage every six months to the Coordinator;
- Administration and scientific coordination activities;
- Implementation of all beneficiary action plans;
- Developing and maintaining a culture of communication and reporting among the team members;
- Creation of efficient team structures to minimize the number of meetings while being flexible.

Project Meetings

The project has planned for at least one physical meeting per year. The timing of these will vary slightly according to the work plan, so that they can be aligned with the best points for the collaboration between tasks.

In addition, teleconference calls or online meetings every four months will ensure that progress is tracked and any issues are picked up quickly and dealt with at the earliest possible opportunity.

Where possible, meetings will be arranged alongside existing meetings and conferences to minimise costs and travel time.

Conflict Resolution will follow the following procedure:

Stage 1 - this is an informal stage where the Project Coordinator will discuss individually with the parties and together with the parties to determine the precise nature of the conflict and explore possible solutions to the conflict.

If the conflict cannot be resolved, then the process is to go to the next stage.

Stage 2 - this is a formal stage. The parties will bring the conflict to the Project Coordination Committee, which is the formal decision making body in the project.

If the PCC cannot resolve the conflict to the satisfaction of the parties, then the process is to go to the next (and final) stage.

Stage 3 - this is the external stage. The conflict will now be resolved according to the Consortium Agreement, which will include an option of arbitration.

Project Coordinator The Project Coordinator will be John Gallagher, who has over 20 years' experience of European collaborative research projects, and has worked in both academic and industrial institutions. He has been the team leader in 6 European projects. He will be assisted by administrative project manager Dixi Louise Strand.

Project Work Package Leaders

The following are the assigned work package leaders:

Work Package	Leader
WP1 Energy-Aware Software Engineering	P. López-Garcia
WP2 Energy Modelling Through the System Layers	K. Eder
WP3 Analysis and Verification	P. López-Garcia
WP4 Optimization	J. Gallagher
WP5 Establishing the Benchmarks	H. Muller
WP6 Evaluation	H. Muller
WP7 Dissemination, Collaboration and Exploitation	K. Eder
WP8 Management	J. Gallagher

Quality Assurance

Quality assurance is simple and will consist of peer review of technical deliverables for accuracy and conformity with the project workplan. The Project Manager will be responsible for ensuring the timely delivery of deliverables to the EC and the external reviewers.

Publication Vetting

The Project Manager will vet all material intended for publication to check that none of it is prejudicial to the industrial beneficiary's needs to protect intellectual property as a precautionary measure.

In cases where a potential risk is found, the Project Manager will chair discussions with the beneficiaries involved to resolve it by changes to the material or an agreed reasonable delay to allow IP protection e.g. patents to be filed without prejudice. If resolution is not possible informally, then the normal dispute resolution procedure will be followed.

Risk Management

The project manager will keep a risk register updated every quarter. In addition, the project manager will keep a watching brief on industry developments that may affect the project exploitation. Both of these will be reported on in the management reports.

B2.2 Beneficiaries

B2.2.1 Roskilde University

Roskilde University (RUC)) is a Danish state university founded in 1972 with the objective of providing research and education at the highest level in the fields of natural science, social science and the humanities. Activities at Roskilde University are organised in six departments (institutes) and degrees are offered in a wide range of subjects at Bachelor, Master and Ph.D. levels. The University is the daily place of work for around 8000 students, 700 researchers and lecturers as well as 250 technical/administrative staff. Classical thoroughness combined with problem-oriented interdisciplinary approaches characterizes research at Roskilde University. The Department of Communication, Business and Information Technologies (CBIT) is an innovative and interdisciplinary university environment with over 100 academic staff. CBIT combines information technology, the social sciences and humanities in the study of the interplay between communication, business and information technology in theory and in practice. CBIT contains seven research groups including the "Programming, Logic and Intelligent Systems" group (PLIS). The PLIS group (http://plis.ruc.dk) contains 7 full-time researchers plus a number of PhD students and postdocs. The group investigates foundations, tools and

languages for the development of adaptable, reliable, human-oriented computer systems. The group's research covers theoretical foundations, languages, tools and semantic models together with application areas. The main research topics are: Programming languages and tools; Knowledge-based systems and intelligent interaction with systems; Logic and knowledge representation. Expertise and Related Activities The group is active in topics around static analysis of programs and verification. Recent and current relevant projects carried out within the PLIS group include "ASAP: Advanced Specialization and Analysis for Pervasive Systems," funded by the EU Framework 5 programme, 2002-2006; "SAFT: Static Analysis with Finite Tree Automata;" funded by the Danish Natural Science Research Council, 2007-2010 and "Numeric and Symbolic Abstractions in Software Model Checking," funded by the Danish Natural Science Research Council, 2011-2013. The PLIS group has built up tools and expertise in program analysis and specialisation tools, including the CHA and Tattoo online tools (http://saft.ruc.dk/CHA, http://saft.ruc.dk/Tattoo) for numeric and symbolic analysis of constraint logic programs. These have been used previously for analysis of microprocessor code and real-time controllers. The group has expertise in integrating state-of-the-art solver libraries (SAT, SMT, BDD, polyhedra, octagons, etc.) into program analysis tools.

Key personnel:

John Gallagher (leader of Roskilde project team). Professor in Computer Science. From 1990 to 2002, he was lecturer and then senior lecturer at the Roskilde University, UK. He obtained his PhD in 1983 in Trinity College Dublin and has had research positions at the Weizmann Institute, Israel, Katholieke Universiteit Leuven, Belgium and in industry in Germany. He has been active in the area of logic in programming, knowledge representations, program analysis and software development tools for 20 years, and has over 50 peer-reviewed papers in this area, particularly on analysis and specialization of constraint logic programs. He is in the steering committee of the ACM PEPM conference series, the executive committee of the Association of Logic Programming and is an Area Editorial for the Journal of the Theory and Practice of Logic Programming.

John will coordinate the whole project, as well as coordinating WP4. He will work mainly in WP1, WP3 and WP4 and co-supervise PhD students and research assistants hired in the project.

Mads Rosendahl. Associate Professor. Research topics include abstract interpretation, complexity analysis, set constraints and program specialisation.

Mads will contribute mainly to WP3, and will co-supervise PhD students and assistants working in WP1 and WP4.

Morten Rhiger. Associate Professor. Research topics include programming language implementations, including interpreters, compilers, and run-time code generation; type systems and type safety; multi-stage programming; denotational, operational, and monadic semantics; functional and other higher-order programming languages.

Morten will contribute mainly to WP1 and will co-supervise PhD students and assistants working in other WPs.

Dixi Louise Strand. Special advisor. She obtained her PhD from Computer Science, Roskilde University in 2003. Dixi has experience as the administrative research manager of a large collaborative project on innovation in virtual worlds funded by the Danish Strategic Research Council (2009-2011). Dixi has furthermore participated in two other strategic research projects, funded by the Danish Research Council, on IT design of web services (2000-2003) and IT-based collaboration in health care (2004-2007).

Dixi will assist the Project Coordinator and the Project Coordination Committee as project administration manager and will coordinate and contribute mainly to WP8.

B2.2.2 University of Bristol

The University of Bristol (UNIVBRIS) is renowned internationally for its excellence in research across a wide range of disciplines, regularly being ranked amongst the top Universities in the world. It is a member of the Worldwide Universities Network and a member of the prestigious Russell Group of Universities in the UK. The Merchant Venturers School of Engineering (MVSE) at the University of Bristol has nine highly specialized Research Groups and several world-leading Research Centres. The research on Design Verification within the recently established Microelectronics Research Group concentrates on specification, verification and analysis techniques which allow designers to define a design and to verify/explore its behaviour in terms of functional correctness, performance, power consumption and energy efficiency. The University of Bristol was the first university in the UK to train Computer Science undergraduates in the skills expected from professional verification engineers. Research in the Microelectronics group is conducted in close collaboration with the local and international semiconductor design and EDA industry where group members have developed strong links to companies such as large multi-nationals including Infineon, ST Microelectronics, Broadcom, as well as Cadence and Mentor Graphics, and many smaller companies such as NVIDIA, Picochip, Xintronix and XMOS. Research and industrial collaboration in the MVSE are closely linked to teaching activities. The Department of Computer Science runs MEng/BSc and MSc programmes; most relevant to ENTRA are the MSc in Advanced Microelectronic Systems Engineering and also the undergraduate MEng in Computer Science and Electronics. The University of Bristol also hosts a world-leading Systems Centre where Sustainability is a major research topic.

Key personnel:

Dr Kerstin Eder is Reader and Research Fellow in Computer Science. She holds a degree from the Technical University Dresden (Dipl.-Informatikerin) and an MSc in Artificial Intelligence from the University of Bristol. Her PhD focused on Computational Logic. Dr Eder's research activities are focused on Design Verification including both formal methods and simulation-based approaches. She has authored over 30 technical publications in fully refereed international conferences, workshops and journals and has been investigator on 7 externally funded research projects raising in excess of £1.8M in grant income. Her most recent projects are on formalizing processor instruction sets [92, 94] and on functional verification of adaptive robotics [31]. Since 2003 she has been advising the UK National Microelectronics Institute (NMI) where she is the lead academic in the NMI Design Verification Roadmapping initiative. Dr Eder initiated the introduction of the MSc in Advanced Microelectronics Systems Engineering at Bristol and has been Programme Director until 2010. In March 2007 she was awarded a prize for "Excellence in Engineering" by the Royal Academy of Engineering. Based on Dr Eder's expertise in Design Verification the University of Bristol was selected as a Lead University to head the Cadence Academic Network in the area of Advanced Verification Methodology. In October 2010 Dr Eder won a one-year Royal Academy of Engineering Industrial Secondment to XMOS with the aim to understand industrial power-efficient design methods. Dr Eder is an active member of the Systems Centre at the University of Bristol.

In January 2011 Dr Eder set up the **Energy-Aware COmputing (EACO) Initiative**[1] at Bristol. It brings together researchers and engineers with interests in energy-aware computing. EACO aims to go significantly beyond the state of the art by engaging both communities in collaborative research projects at the leading edge of science and technology. The EACO workshop series has attracted participation from leading industrialists and academics - locally, nationally and internationally. Dr Eder is now leading the EACO initiative toward the first set of collaborative research projects to be launched on 18 April 2012.

[1] Further and up-to-date information on the EACO workshop series, the Industrial Partners and the latest news can be found online at http://www.cs.bris.ac.uk/Research/Micro/eaco.jsp.

Based on her profile and expertise both in energy modelling and verification, Dr Eder will coordinate the research in WP2 and the work in WP7. She will also contribute towards WP3, especially on verification, and participate in WP1 and WP4 tasks. Dr Eder will be the principal supervisor of the Research Assistants to be employed in Bristol during the project.

Dr Nunez-Yanez is a Senior Lecturer in Electronic Circuit Design. He obtained a BSc in Industrial Engineering and an MSc in Microelectronics Engineering from the University of La Coruna (Spain) and University Politecnica the Catalunya (UPC, Spain) respectively. He holds a PhD from Loughborough University. His expertise is in the area of reconfigurable computing using FPGAs, computer architecture and signal processing. During 2005 he was a Marie Curie Fellow at ST Microelectronics, Italy, investigating optimial couplings between a RISC processor and a dynamically reconfigurable fabric [68]. Dr Nunez-Yanez pioneered the area of lossless data compression using variable order Markov models in hardware to support optimal statistical compression of multimedia data without any loss of quality and real time processing [67]. In addition, he is active in the areas of high-performance reconfigurable computing and on-chip communication architectures exploring fault-tolerant computing paradigms. During 2011 he was an Industrial Fellow at ARM Ltd supported by the Royal Society, investigating techniques for high-level modelling and optimization of energy consumption in complex system-on-chips. He also continues work in efficient video/signal processing algorithms targeted to hardware architectures [69] and Bio-inspired computing [96].

Based on his profile and expertise, especially in high-level energy modelling, Dr Nunez-Yanez will contribute mainly towards the research in WP2 and WP5. Dr Nunez-Yanez will co-supervise the Research Assistants to be employed in Bristol during the project.

Dr Simon Hollis is a Lecturer in Computer Science. He holds an MA and a PhD in Computer Science from the University of Cambridge. His PhD focused on energy-saving interconnects, and his research interests range through computer architecture, novel clocking schemes, on-chip networks and parallel programming paradigms. In all his work, energy has been a primary focus, and has led to the creation of the energy-efficient 'RasP' and 'Skip-link' networks. He is currently building a communication-oriented 480 processor system, built on the XMOS platform. Dr Hollis has published 16 papers in internationally refereed conferences and journals, and teaches two undergraduate and post-graduate courses on embedded system and ISA design. He has hosted the UK Asynchronous Design Forum, dedicated to novel approaches to logic design.

Based on his profile and expertise, especially in computer architecture and parallel computing, Dr Hollis will contribute mainly towards the research in WP2, WP4 and WP6. He will also participate in WP3 tasks. Dr Hollis will co-supervise the Research Assistants to be employed in Bristol during the project.

Steve Kerrison holds an MEng Hons in Computer Systems Engineering. He received the Toshiba sponsored award for "Best Computer Systems Engineering Student" from the Department of Computer Science. His areas of interest are embedded systems development and computer architectures. He is now a PhD student with research interests in methods for energy-aware software implementations. During 2011 he has been working jointly with Dr Eder at XMOS to examine architectural and micro architectural features of processors that permit software control of power consumption for strictly-timed, real-time applications.

Steve will be contributing towards energy modelling in WP2, analysis in WP3 and benchmarking in WP5.

B2.2.3 IMDEA Software Institute

IMDEA Software Institute is a non-profit Foundation. This legal status provides appropriate guarantees for the management of public funds while allowing flexibility and simplicity in administrative procedures as well as in hiring, for both researchers and technical staff. IMDEA Software Institute is part of IMDEA, the Madrid Institute of Advanced Studies, a network of international research centers in the Madrid region for research of excellence in areas of high economic impact. The main goal of IMDEA Software Institute is to perform research of excellence for the cost-effective development of high quality software products with sophisticated functionality. In order to achieve this goal, the Institute is gathering a critical mass of world-wide, top class researchers and providing them with an ideal research environment. Search and selection procedures follow international standards and the recommendations of the European Charter for Researchers, and are coordinated by the Institute in consultation with its Scientific Advisory Board. The main selection criterion is the demonstrated ability and commitment of the candidate to perform research in the area of the Institute. Although the research carried out at the Institute is organized in a number of lines or themes, the interactions and synergies among the different groups are very strong. This is due to the fact that this research shares not only a high-level objective - developing science and technology for the cost effective development of high quality software products with sophisticated functionality- but also a common focus: the search for rigorous approaches to software development that can be implemented or integrated in practical tools. Also, these interactions and synergies are natural, and almost necessary, given the current size of the Institute. For these reasons, to judge the training capacity of the group that proposes this project, it is only fair to consider the Institute as a whole. IMDEA has expertise in static analysis (including resource usage analysis), and verification. Succesful experiences in transferring static analysis and verification techniques to the critical embedded systems industry have been carried out within the project ES_PASS "Embedded Software Product-based ASSurance." IMDEA has also expertise in programming languages design and implementation, declarative languages, compilers, and tool implementation and integration (c.f., the Ciao/CiaoPP system http://clip.dia.fi.upm.es/Software/Ciao).

Key personnel:

Pedro López-García received a MS degree and a Ph.D. in Computer Science from the Technical University of Madrid (UPM), Spain in 1994 and 2000, respectively. In May 28, 2008 he got a Scientific Researcher position at the Spanish Council for Scientific Research (CSIC) and joined the IMDEA Software Institute. Immediately prior to this position, he held associate and assistant professor positions at UPM and was deputy director of the Artificial Intelligence unit at the Computer Science Department. He has published about 40 refereed scientific papers (50% of them at conferences and journals of high or very high impact.) He has also coordinated the international project ES_PASS and participated as a researcher in many other national and international projects. His main areas of interest include automatic analysis and verification of non-functional global program properties such as resource usage (user defined, execution time, memory, etc.), non-failure and determinism; performance debugging; (automatic) granularity analysis/control for parallel and distributed computing; profiling; unit-testing; type systems; constraint and logic programming. He is one of the main authors of the Ciao/CiaoPP system, an advanced program development environment that performs automatic analysis, verification, debugging and optimization.

Given his profile and broad experience, Pedro will mainly coordinate the work in WP1 and WP3, participate in WP2 tasks, and co-supervise PhD students and assistant researchers hired in the project.

Manuel Hermenegildo is the Director of the Institute, and also full professor of Computer Science at the Universidad Politénica de Madrid (UPM). Previously to joining IMDEA Software he held the P. of Asturias Endowed Chair in Information Science and Technology at the U. of New Mexico, USA. He has also been coordinator and/or principal investigator of many national and international projects.

He has received the Julio Rey Pastor Spanish National Prize in Mathematics and Information

Science and Technology and the Aritmel National prize in Computer Science. Manuel Hermenegildo has been involved in the supervision of 13 Ph.D. students (of which 11 have already completed their Ph.D.s) at UPM, The University of New Mexico, and The University of Texas at Austin. The group expertise is enriched by the experience provided by the rest of the Institute's scientific staff and by the growing list of international distinguished visitors that come to the Institute for medium to long-term research stays (Peter Stuckey, Alan Mycroft, and Javier Esparza, among others).

The project in general, and in particular WP1, WP3 and WP4 tasks will benefit form his ample experience and vision in program analysis (including abstract interpretation), verification, optimization (including automatic parallelisation) and advanced software engineering tool development. He will also co-supervise PhD students and assistant researchers hired in the project.

César Sánchez received his M.S. and Ph.D. degrees from Stanford University in 2001 and 2007. Previously to join the IMDEA Software Institute in 2008, he was a post-doc at University of California at Santa Cruz. In 2009, Dr. Sánchez was appointed as Scientific Researcher by the Spanish Council for Scientific Research (CSIC), to support his research at the IMDEA Software Institute. César has published many scientific papers in conferences of high impact in the areas of formal methods and distributed systems. César Sánchez is currently supervising 2 Ph.D. students. His research activities focus on formal methods for reactive systems with emphasis on the development and verification of concurrent, embedded and distributed systems. His foundational research includes the temporal verification of concurrent datatypes, runtime verification, and enhancements of linear temporal logics. He is interested in the study of hard problems in the area of concurrency like non-functional properties and properties beyond safety and liveness, asynchronous programs, distributed systems and datatypes In parallel, he is collaborating with industrial partners from the aerospace and embedded sectors to aid in the adoption of formal techniques for software development and validation. Current projects include the interactive formal generation of parallel software for satellite image processing, and the synthesis of advanced online debuggers for testing embedded software.

In this project, César will contribute mainly to WP4 tasks, as well as WP6 tasks, given his ample expertise in embedded systems. He will also co-supervise PhD students and assistant researchers hired in the project.

Laurent Mauborgne received his Ph.D. in Computer Science from École Polytechnique, France, in 1999, and an Habilitation à diriger les recherches from University Paris-Dauphine (France) in 2007. He has been assistant professor at École normale superieure, Paris, since 2000, and associate director of computer science studies there since 2006. He was also part-time professor at École polytechnique. He was invited to spend a year at IMDEA Software in August 2009. He published 16 refereed papers in international conferences and 3 papers in reviews. He gave courses in research summer schools and participated in the European projects DAEDALUS and ES_PASS. He was program committee member of the Static Analysis Symposium for 4 years. He is one of the authors of the Astrée analyzer, a tool that proved the absence of run-time errors in critical avionic codes. The research of Laurent Mauborgne is focused on static analysis of programs and abstract interpretation. The goal is to develop theoretical as well as practical tools to analyze the behaviours of programs. This includes proving safety or temporal properties, optimizing compilation and computing resource usage. Among the recent subjects, he studied the cooperative combination of analyzes in different frameworks.

In this project, Laurent will contribute mainly to WP1 and WP3 tasks, as well as WP4 and WP6 tasks, given his ample expertise in analysis and tool development, one of them currently being comercialised and used in the embedded software development sector. He will also co-supervise PhD students and assistant researchers hired in the project.

In general, each PhD student or assistant researcher to be hired in the project will be co-supervised

by two of the above key researchers.

B2.2.4 XMOS Limited

XMOS Limited is a fabless semiconductor company that was founded in 2005 in Bristol, United Kingdom. XMOS has attracted investments from top tier venture capitalist firms Amadeus Capital, DFJ-Esprit and Foundation Capital. XMOS develops technology for an event-driven multi-threaded processor engine. It has with tightly integrated pin I/O and timing control, allowing many hardware protocols to be implemented in software. The XMOS architecture enables systems to be constructed from multiple XCore processors connected by communication links. Every XMOS device includes one or more XCores and a high-speed low- latency switch. The switch is used to route messages between the XCores on the chip, and to route messages between chips via the links. The architecture is scalable and any number of XCores can be connected together. As XCores are added, computational performance increases, memory increases, communication throughput increases and event-handling throughput increases. XCores can be combined on chips, substrates, boards or in distributed systems. The architecture is event-driven and is designed to minimize energy. XCores and threads do not consume processing resources when waiting for events. The XCore instruction encoding gives rise to very compact programs, optimizing use of memory and minimizing power to access instructions. Direct transfer of data between thread registers and channels or ports avoids transferring data via memory. XMOS has developed two generations of silicon (the XS1-G and XS1-L), and a complete tool chain that targets the silicon. The tool chain is based on XC, a deterministic concurrent programming language that complements sequential capabilities of C and C++. XC provides explicit control of concurrency, communication, timing and I/O.

Key personnel:

Henk Muller is the Principal Technologist at XMOS since January 2008. In that role he has been involved in the design and implementation of audio and other real time protocols, and the design and implementation of the XS1-L architecture. One of his tasks is to keep a close eye on the interaction between applications, compilers, and silicon. Prior to working in Industry, Henk worked in academia for 20 years in the fields of Computer Architecture, Compilers, and Ubiquitous Computing. Henk was the Bristol Principal Investigator on the ASAP project on program specialisation, and on the 10M EPSRC Equator project on merging the digital and the physical. He holds a doctorate from the University of Amsterdam, has contributed to more than 80 papers and 5 patents.

Henk will contribute mainly to WP4 and WP6 and will supervise XMOS staff in WP1 and WP5.

David May is the Chief Technology Officer at XMOS. He is one of the founders of XMOS and also holds a chair at the University of Bristol. David is known for his work on the Transputer in the 1980s, when he was chief architect at Inmos. In this role, he created the Transputer architecture and designed the associated programming language Occam. Occam was one of the first industrial strength concurrent and real time language. In 2005 he co-founded XMOS and designed the XCore architecture and defined the XC programming language. The XCore went for sale in 2007 with an XC compiler. David holds two fellowships: in 1991 David was elected a Fellow of The Royal Society, and in 2010, he was elected a Fellow of the Royal Academy of Engineering.

David will contribute to WP6.

B2.3 Consortium as a Whole

The ENTRA consortium brings together a unique set of competencies. These bridge several traditionally separated disciplines in the wider subject areas of computer science, semiconductor design

and electronic design automation. It is the combination of the techniques and methods from these disciplines that underpin the fundamental vision of ENTRA.

The consortium brings together:

- Roskilde University. The PLIS group investigates foundations, tools and languages for the development of adaptable, reliable, human-oriented computer systems. The group's research covers theoretical foundations, languages, tools and semantic models together with application areas. Static analysis and program transformation are particular areas of expertise.
- The University of Bristol, the lead university in the Cadence Academic Network in Europe on Advanced Design Verification Methodology. It has unique expertise in formally modelling and analyzing designs of industrial complexity at various levels of abstraction from software down to silicon. It is well connected to the large cluster of semiconductor design companies in the local area. Research at Bristol informs teaching and hence offers a fast route to bringing leading edge research into the educational curriculum. UNIVBRIS also hosts the Energy-Aware COmputing (EACO) initiative - a platform to engage with industry to promote the ENTRA project and its research results, and a route to exploitation.
- The IMDEA Software Institute has ample expertise in areas related to the project such as static analysis (including resource usage analysis), verification, parallelization and programming languages design and implementation. IMDEA has also experience in transferring static analysis and verification tools and techniques to industry.
- XMOS, an SME with a revolutionary event driven multi-core processor family for embedded systems, along with tools for ensuring that hard timing constraints can be met.

Most of the project partners have worked together in the ASAP project, and the University of Bristol is the lead institution for the Cadence academic network and has close ties; this includes the Cadence tools chain being available for academic use. There are also close ties between XMOS and the University of Bristol. John Gallagher from Roskilde University has ongoing collaboration with IMDEA Software Institute as a part-time research professor.

The following table covers the beneficiaries' principle skills and background, and where they are most focussed in the project.

Expertise	**WPs**	**RUC**	**UNIVBRIS**	**IMDEA**	**XMOS**
Formal Specification	2, 3	X	X	X	
Energy Models	2, 4		X		X
Program Analysis	3				
Resource Usage Analysis				X	
Control/Data Flow Analysis		X			
Optimisation	4				
Low-level Optimization			X		
High-level Optimization		X		X	
Verification	2, 3	X	X	X	
Tool Integration	1			X	X
Evaluation	1, 5, 6				X
Emb. System Design	5, 6		X		X
Sustainability of ICT	6, 7		X		

i) Sub-contracting:

No sub-contracting is planned, other than the audit certificates under management.

ii) Other countries:

No beneficiaries from other countries are involved.

B3 Impact

B3.1 Strategic Impact

B3.1.1 Transformational Impact on Science, Technology and/or Society

In this section we first identify the transformational effect that *Energy Transparency through the Layers in the ICT System Stack* will have on Science, Technology and on Society. We then present, for each impact aspect, the measures we will take and the connections we can mobilize to ensure that the best achievable impact is actually realized.

Whole Systems Energy Transparency is a radical departure from existing system design practice where the lack of energy transparency puts a stop to the savings that could technically be made. Energy Transparency will thus have a significant impact on Science, Technology and on Society. It bridges the gap that currently exists between modern hardware, that uses lower and lower switching power per bit processed, and software, by enabling software to take full advantage of the energy savings realized in hardware. Energy transparency is the key principle that allows reducing software energy consumption to the limits achievable by modern hardware. This is where the biggest savings can be made in system energy consumption.

It is important to note that the modelling, analysis, optimization and verification methods to achieve energy transparency as proposed in ENTRA are in principle generic. They are applicable throughout the various layers within a system. While the focus during ENTRA will be on embedded systems, which is motivated through the industrial partner, XMOS, in the consortium, the techniques to achieve energy transparency can equally be applied to "end-to-end" applications, digital media etc.

B3.1.1.1 Transformational Impact on Science

Whole Systems Energy Transparency will significantly transform science by creating the new integrated discipline **Energy-Efficient Computing** and by stimulating new research and innovation in existing research disciplines.

- Energy transparency will bridge the gap between two research communities, hardware design and software engineering. Advances in one area will have an "instantaneous" effect in the other, and thus on the whole system energy consumption. Energy transparency will allow a bidirectional feedback between both communities, so that, instead of having two research areas advancing at different speeds, the new integrated research discipline of **Energy-Efficient Computing** will emerge. Fundamental research is needed to gain full understanding of the impact of design decisions in one area on the energy consumption of the other, and to drive the development of new holistic energy-efficient system design methods.

Existing research disciplines that will be transformed by exploiting Energy Transparency throughout system layers include, but are not limited to:

- **Algorithms:** Energy transparency will stimulate new research to re-evaluate algorithms for energy efficiency and to develop new energy-efficient algorithms. Finding the best "hardware fit" for an algorithm is currently a challenging task that requires intricate familiarity with the power-saving features of the target hardware on which an algorithm will be executed. This is further complicated by the performance focused optimizations of compilers and at micro-architectural level. Some of these are counter-productive when it comes to optimizing for energy consumption. Energy transparency is the key for making a fundamental breakthrough to advance this field.
- **Computational Complexity:** The metrics traditionally used to determine the computational complexity of an algorithm in terms of the input data size are not precise enough to capture energy consumption. For instance, the number of data items does not capture the amount of switching in the hardware caused by the data stream. Energy transparency from silicon to software will stimulate new research to re-evaluate traditional metrics for computational complexity and to develop new metrics that capture energy consumption of computations.
- **From Real-time Programming to "Real-Energy" Programming:** Energy transparency opens the opportunity for a new dimension in the design space of systems, especially for embedded applications. Fundamental research is needed to fully understand the tradeoffs to be made and to develop novel holistic system design methods that optimize for energy efficiency. Research into new extensions to programming languages and to compiler tool chains will be required to make energy consumption a first class software design goal.
- **Electronic Design Automation for Hardware Design:** Energy transparency will enable the provision of early "use cases" that indicate energy consumption budgets so hardware can be developed to meet these. Fundamental research is needed to establish how best to utilize this information and how to integrate it into existing hardware development EDA tool flows.
- **High-Performance Computing (HPC):** Energy consumption is a major concern in High-Performance Computing. Energy transparency enables significant advances in HPC where new research will lead e.g. to better understanding of the energy consumption requirements of computations and therefore to more accurate cost models for HPC services.
- **Sustainability of ICT:** Energy transparency will help researchers understand what needs to be done to achieve a sustainable level of computing related energy consumption. For instance, it has been estimated [86] that cloud-computing related energy consumption has been increasing by 14% per year, while Internet traffic has increased by 50% per year. To provide the 2030 Global Middle Classes with western standards of online media could require 18% (as opposed to 2%) of world energy at 2010 efficiency levels.

  This is unsustainable; a solution needs to be found to prevent this gap from growing. The fundamental question now is whether a techno-fix is sufficient (i.e. making computing more energy efficient) or whether end-user behaviour also needs to change to achieve a sustainable level of computing-related energy consumption. To answer this question, researchers need to understand the limits achievable by making computing as energy efficient as possible. This, in turn, will help them understand the impact that can be expected from a techno-fix alone. It is very likely that end user behaviour will also have to change to keep computing-related energy consumption within sustainable levels. Energy transparency provides key information to help answer these important research questions.

The ENTRA consortium is well placed to take a leading role in realizing the transformational effects of the project results on the above mentioned areas in Science because it brings together

internationally leading research universities and centers with excellent track records in dissemination of research results. All academic partners are well engaged with the scientific research community, specific dissemination plans are detailed in Section B3.2.

Energy efficiency in a broad sense is a highly interdisciplinary field; below we list activities and connections beyond those directly relevant to the research programme in ENTRA for each academic partner.

- The research within the Microelectronics Research Group at UNIVBRIS has a strong focus on energy-efficient computing from silicon to software. At UNIVBRIS this area is a high-priority research area and the university has committed to further strengthen this activity by supporting the EACO initiative which has put UNIVBRIS into a nationally and internationally leading position in this field. The EACO initiative at Bristol also includes leading researchers in Algorithms and Complexity as well as members from the Sustainable Computing research group and the Systems Centre, all share a common interest in energy-efficient computing and in advancing the state of the art in this area. UNIVBRIS also operates an HPC facility and has HPC experts that are closely related to the ENTRA consortium through membership in the Microelectronics Research Group.
- RUC's organisation and study programmes foster interdisciplinarity in research and teaching and thus RUC will develop future research in energy-aware design of computer systems in a multi-disciplinary context. Research areas will concern not only energy-efficient computing systems themselves, but also the design of the total environment of which computer systems are a part, including its energy aspects. Two institutes at RUC, CBIT (Department of Communication, Business and Information Technologies) and ENSPAC (Department of Environmental, Social and Spatial Change) already pursue a variety of research themes around design and modelling in whole-systems contexts.
- IMDEA Software Institute is part of IMDEA, the Madrid Institute of Advanced Studies, a network of international research centers in the Madrid region for research of excellence in areas of high economic impact. One of such centers is the IMDEA Energy Institute, whose goal is to promote R&D activities in energy related themes, with main emphasis in those topics related to renewable energy and clean energy technologies. IMDEA has also very close connections with a Spanish HPC facility and research center, located in the same "International Campus of Excellence in Research and Technology Transfer" awarded science and technology park as IMDEA. The Madrid Supercomputing and Visualization Center (CESVIMA), hosts the "Magerit" Supercomputer, the more powerful and only Spanish supercomputer, which has reached the second best position in history on the TOP500 list of world's most powerful supercomputers. IMDEA is also a member of the INES cooperative net, and the Madrid+d Foundation, that integrate major Spanish research and technology agents, for which IT efficiency is a main topic in their strategic agendas.

B3.1.1.2 Transformational Impact on Technology

Whole Systems Energy Transparency will fundamentally transform our methods of system design and engineering in the future. The availability of energy consumption information will lay the foundations for a **new industry, Energy-Efficient Computing,** with Europe at the forefront of this new technology wave. In addition, exploiting Energy Transparency will also give **significant competitive advantage to existing ICT sectors** from hardware design to application development, from embedded systems to High-Performance Computing.

- Energy Transparency through the layers of the system stack will enable system designers and engineers to use energy consumption information during system design and develop-

ment. This creates the opportunity for **Energy-Efficient Computing**, a *new industry* focused on the design and development of software, hardware or entire computing systems that optimize energy consumption of computation by achieving an optimal fit of algorithms to hardware through early exploitation of energy consumption information at all layers in the system stack.

This is a fundamental shift from existing practice and constitutes progress that is significantly beyond the state of the art. Energy Transparency enables system designers and engineers to drive energy consumption of computing to the limits achievable by the latest hardware designs. More importantly, it allows them to keep up with and to take advantage of the latest savings realized though fundamental advances within the entire system stack, e.g. at the physical level or via novel paradigms of computation.

The UK government has recently estimated the global market for this emerging new industry to be $50 billion in 2020[2]. Europe has a strong hardware and software industry, especially in the embedded systems sector. Energy-Efficient Computing enabled by Whole System Energy Transparency will fundamentally transform the way this industry operates and thereby give this industry a competitive advantage that fosters Europe's leadership in this area globally.

Existing technology sectors that will be transformed by exploiting Energy Transparency throughout system layers include, but are not limited to:

- **Software Engineering:** Energy transparency will enable software engineers to achieve a better fit of programs to hardware. Currently, energy-aware software design can only be achieved if software engineers are intimately familiar with the power-saving features of the hardware they target. Energy transparency will enable software engineers to understand the impact of their design decisions on energy consumption without the need to be hardware literate and without the "trial and error" approach that is currently used to compensate for this.

  Energy transparency thus adds a new dimension to the software development process that exposes energy consumption information to software engineers while maintaining existing software engineering strategies such as use of high-level programming languages.

- **Hardware Design:** Energy transparency will enable hardware designers to gain understanding of how the software uses the hardware. Different applications have different needs, and the lack of transparency creates uncertainty for hardware designers about how the software impacts the system. This uncertainty has reportedly led to hardware design projects being canceled because of the lack of certainty to meet the energy budget of the final product.

  Energy transparency will transform the hardware design sector by enabling system-driven energy consumption analysis that allows hardware designers to fine-tune hardware to best meet the needs of the target application software.

- **High-Performance Computing (HPC):** HPC solves advanced computations and is a major consumer of computation-related energy today. Energy transparency will enable HPC providers and users to gain a better understanding of the energy consumption requirements of HPC applications. A fundamental shift from predominantly *performance-oriented* considerations to include also *energy consumption* will be possible through energy transparency of both the HPC hardware and application software.

[2] page 29 in http://www.bis.gov.uk/assets/biscore/innovation/docs/i/11-1387-innovation-and-research-strategy-for-growth.pdf

- **ICT Service Provision:** Through energy transparency, reporting energy consumption on a whole systems level will become feasible. Once all applications (end-user, network, server) know about the energy consumption of individual API calls, a very precise accounting becomes possible. This level of transparency in turn enables system-wide market mechanism based competition of services for energy (resources). Currently, energy consumption of individual services is not transparent, i.e. "unknown", and hence, users cannot choose between services based on energy consumption.

The ENTRA consortium is well connected to the stakeholders in the technology sector and to policy makers. In particular, the EACO initiative led by UNIVBRIS provides a platform for ENTRA to engage with industry to promote the ground-breaking transformations that can be achieved through Energy Transparency. XMOS itself is a hi-tec SME and as part of the ENTRA consortium will take advantage of the research results within ENTRA.

National governmental agencies in the UK, Spain and in Denmark have developed strategies towards Energy-Efficient Computing and "Green IT":

- The UK Department for Business, Innovation and Skills (BIS, http://www.bis.gov.uk/) together with the UK Technology Strategy Board (TSB, http://www.innovateuk.org/) have recently prioritized energy-efficient computing within their Innovation Strategy[3].
- INES (http://www.ines.org.es/), a scientific and technological cooperative net that integrates the main Spanish agents (companies, universities, technological centers, etc.) in Software and Services, and the Madrid+d Foundation (http://www.madrimasd.org/), that also integrates the main agents in the Madrid Region and Spain, have included energy efficiency of computing (and IT in general) as one of the key challenges in their strategic agendas. The Madrid+d Foundation has produced a technology watch report on Green IT[4].
- Denmark has a long tradition of being at the forefront in the development of energy-efficient technologies. The Danish national IT and Telecom Agency launched a funding programme for "Green IT" in 2008, aiming to "... contribute to positioning Denmark on the forefront of Green IT, partly by bringing about a greener lifecycle for IT solutions - from development, production and use to disposal - and partly by strengthening research and the development of solutions applying IT to reduce our environmental footprint.

At the European level there are similar initiatives and uniting efforts in this direction. Clearly, there is an opportunity and an urgent need for fundamental research in this area to enable a transformational effect and the associated significant growth in this emerging technology area. The ENTRA consortium will use every opportunity to promote ENTRA results at national and at European level, e.g. by participating at "Show and Tell" events organized by these agencies.

In addition, UNIVBRIS is already working with the Carbon Trust in the UK to help them in specifying energy accounting standards for ICT services, which currently can only be done through very rough estimation. Similar activities are being undertaken at the European level by the European Telecommunications Standards Institute (ETSI) or globally by the International Telecommunication Union (ITU), the United Nations specialized agency for information and communication technologies. The ENTRA consortium will use its existing contacts and actively seek new connections to ensure these organizations are kept abreast with the ground-breaking research results from the ENTRA project.

[3] http://www.bis.gov.uk/assets/biscore/innovation/docs/i/11-1387-innovation-and-research-strategy-for-growth.pdf

[4] http://www.madrimasd.org/informacionidi/biblioteca/publicacion/doc/VT/VT19_green_IT_tecnologias_eficiencia_energetica_sistemas_TI.pdf

Finally, **education** is a fundamental cornerstone that safeguards the transformational impact on Technology we expect whole-systems energy transparency will have. The academic partners in the ENTRA consortium have all committed to include Energy-Efficient Computing based on energy transparency into their teaching and research programmes to educate the next generation of software/hardware and systems engineers who will develop the energy-efficient ICT infrastructure that will be required to sustain the growing need for computing power in the future. The teaching materials for a new course on Energy-Efficient Computing will be made publically available enabling other universities to introduce the course. The ENTRA consortium firmly believes that Energy-Efficient Computing is of crucial importance over the coming decades, and that engineers should obtain fundamental education in this area.

B3.1.1.3 Transformational Impact on Society

The societal impact of lowering energy consumption of computing or "Green IT" is generally well understood. In this section we focus on the transformational effects on Society specifically from introducing **Energy Transparency** in this context.

Whole Systems Energy Transparency will transform society because it exposes energy consumption information of computation to governments, policy makers and to the consumers.

- **Raised end user awareness:** Imagine a world where ICT products, from applications running on your smartphone, your web browser and accounting software on your desk top, to the SatNav in your car or the data mining package you use for market analysis at work etc, all these come with energy ratings similar to the ratings used for kitchen white goods or to label light bulbs. This will be possible when energy consumption of computing is fully transparent.

  Energy transparency permits end users to differentiate between functionally equivalent ICT products based on their energy consumption. This opens up the opportunity to **rate ICT products according to their energy consumption.** In turn, such ratings empower end users to demand more energy-efficient options, thus promoting energy consumption targets to first class ICT product requirements. This creates new demand that further drives competition, research and innovation in **Energy-Efficient Computing** in the future.
- **Policy makers** will be able to use energy consumption information to set crisp targets that can be monitored and regulated. Several organisations are currently implementing energy accounting standards for ICT services (e.g. the UK Carbon Trust, ETSI and ITU). These are only usable for crude product labeling because the models and mechanism for data collection are too coarse. Cross-layer energy modelling and transparency as proposed by ENTRA will enable the development of much more rigorous standards supporting policy and decision makers.

The ENTRA consortium will engage policy makers at national and at European level to guarantee the transformational impact of Energy Transparency on Society can be realized in practice. At UNIVBRIS the Sustainable Computing Group and the Systems Centre are already engaged with UK and European policy makers. This provides the consortium with an important link to ensure societal impact. Furthermore, the consortium will participate at public engagement events and at outreach events for researchers at national and at European level to promote the results of the project.

B3.1.2 Expected Impacts Listed in the Work Programme

The expected impacts listed in the work programme are as follows:

i) Understanding of theoretical limits of energy efficiency in computation (e.g. energy dissipation, thermodynamic and quantum physics limits).
ii) Foundations for computing technologies with negligible energy consumption.
iii) Reduction of the environmental impact caused by the energy consumption of ICT.

ENTRA contributes towards all three impact targets. The specific contributions are outlined in the following sections.

B3.1.2.1 Understanding of theoretical limits of energy efficiency in computation

ENTRA takes a pragmatic perspective on the "theoretical limit of energy efficiency in computation", which we define to mean the *minimal energy achievable by optimal utilization of existing hardware*. The benchmarks to be defined in WP5: Establishing the Benchmarks perform a task on modern hardware using all energy-saving features to full effect; this may need to be achieved through manual encoding.

These benchmarks will establish the "bottom line", i.e. the limit which we will use in ENTRA for comparison. In fact, any code generated using a toolchain will be less (or at best equally) energy efficient than custom-made code. This is the cost of modern software engineering methods that use high-level programming languages and paradigms coupled with the automation provided by toolchains to produce executable code - a long accepted tradeoff.

It is the lack of energy transparency in this process that is causing the substantial gap between hardware capabilities and energy efficiency realized by software. ENTRA is focused on closing this gap by exposing energy consumption transparently through the system layers and using this information to analyse, verify and optimize software for energy efficiency. Energy transparency is the key to realizing the huge energy savings that can be achieved through energy-efficient software.

While research into establishing the theoretical limits helps to set targets for energy-efficient computing, it does not resolve the practical challenges arising out of the lack of energy transparency between system layers that currently prevent us from reaching these. This is where the biggest gains can be made in terms of practical impact and where the ENTRA project will make a significant contribution.

B3.1.2.2 Foundations for computing technologies with negligible energy consumption

Whole-Systems Energy Transparency lays a foundation for energy-efficient system design without which any novel devices or computing paradigms will fall short of realizing their true potential once integrated into the system stack. The techniques and methods developed in ENTRA are generic and effective irrespective of the underlying devices or computing paradigms. Energy Transparency adds considerable value to the system design process because it permits the efficiency gains enabled by any new ground-breaking innovations in computing technologies to propagate through the system layers.

This particular impact will be provided through our work in WP2, WP3 and WP4 and their subsequent dissemination through WP7.

B3.1.2.3 Reduction of the environmental impact caused by the energy consumption of ICT

ENTRA's vision is to enable energy-efficient system design, especially energy-efficient software engineering, through analysis, verification and optimization based on whole-system energy transparency. Experts from Intel [22] estimate that energy-efficient software can realize savings that are three to five times of what is currently possible with conventional software. These savings can be made purely by software taking better control of the energy-saving features of hardware. The ENTRA project is aiming to realize exactly these savings - closing the gap that currently exists between software and hardware.

The energy savings realized by ENTRA undoubtedly contribute towards significantly reducing the environmental impact caused by the energy consumption of ICT. The demand for ICT use globally, however, grows much faster, and with it the ICT-related energy consumption, than the savings that are being made by increasingly energy-efficient ICT provision. Whether energy-efficient computing will be enough to achieve a sustainable level of computing-related energy consumption in the future is currently a hot topic of investigation by experts in Sustainability of ICT [86, 47]. The latest research by sustainability experts at UNIVBRIS suggests [73] that it is very likely that end-user behaviour will also have to change to keep computing-related energy consumption within sustainable levels. As described in Section **3.1.1.1**, under **Sustainability of ICT**, energy transparency provides key information to help researchers understand what needs to be done to achieve a sustainable level of computing related energy consumption.

We must therefore ensure that our dissemination efforts in WP7 reach industry and researchers with active interests in whole-systems design and user behaviour. The sustainability research within the Systems Centre at UNIVBRIS links the ENTRA consortium to this community.

The expertise necessary to successfully conduct the research in ENTRA can only be amassed by forming a consortium from partners that are internationally leading in the respective research areas that ENTRA draws from. Each partner contributes a unique set of skills, knowledge, techniques, infrastructure and connections; individual partner profiles are provided in Section B2.2 and a summary of the expertise per partner is given in Section B2.3. The ENTRA consortium evidently is equipped with a level and breadth of expertise that is not possible to find on a national nor on a local level. ENTRA necessitates cross-country collaboration to bring together the best of European expertise to create competitive advantage and to establish leadership for European research and industry, ultimately to the benefit of European society.

The ENTRA consortium has a significant number of strong links with major energy-related collaborative activities, including research and technological groups, organizations, cooperation networks, projects, etc., both at national and international level. Some of these have been already mentioned (e.g., in Sections B3.1.1 and B2.2). Within ENTRA, we will foster and strengthen

such links, which will increase the impact and dissemination activities and allow the project consortium members to keep abreast of developments in related areas.

The impact of the project would be affected in the unlikely event of a revolution in the energy industry, providing limitless, cheap, renewable energy. While this in itself would be a positive development for the environment, it would certainly reduce the urgency of energy-efficiency in ICT.

B3.2 Plan for the use and Dissemination of Foreground

A number of different activities are proposed to communicate project results and to ensure their exploitation. Special care has been taken to cover a wide range of media and audiences. This maximizes the effectiveness of dissemination and exploitation. The following table summarises the proposed activities over a timeline:

	Year 1		**Year 2**		**Year 3**		**Year 4**	
Activity	1H	2H	1H	2H	1H	2H	1H	2H
Web site		■	■	■	■	■		
Technical press		■		■		■		
Journals				■	■	■		
Conferences				■	■	■		
Workshops		■		■		■		
Education						■		
Summer School						■		
Technology Transfer						■		
Exploitation								

B3.2.1 Dissemination

The ENTRA consortium considers dissemination as a very important route to exploitation and thereby to achieving the fundamental transformational impact that Energy Transparency will have on science, technology and on society. The dissemination plan is structured into two phases: the first phase terminates after 18 months, and the second phase covers the second half of the project. The dissemination strategy of the ENTRA consortium aims to achieve the following objectives:

- raising awareness by promoting the project, communicating the project vision and the importance that energy transparency will play in minimizing energy consumption of computing to the limit,
- contributing to knowledge and understanding by publishing the project results to selected target audiences within the international academic and engineering communities, and
- engaging with stakeholders in science, technology and society as a route to secure uptake and exploitation.

Project web site:

The project web site will feature a public and a private area. The publicly accessible pages shall inform the European citizens and generally any third parties interested in Energy-Efficient Computing on the vision, approach, current status and fundamental results of the ENTRA project. The private area

shall be restricted to the consortium in the first instance. This will be used by the project partners to collect, review and disseminate research related knowledge and data.

Impact: The public project website will lead to higher awareness of the vision of the ENTRA project, its consortium partners, their associated expertise, research collaborations and commercial activities. This will give third parties, such as interested businesses and academic colleagues, an interface through which to interact with the ENTRA consortium, thus widening dissemination, stimulating new R&D collaboration and increasing the potential for knowledge/technology transfer as well as exploitation.

Technical press and Electronic Publishing media:

Press releases are an important instrument to inform practitioners on a national and international level about new advances on energy transparency and on the research results achieved within ENTRA. The consortium will utilize the Press Offices of the academic partners to issue ENTRA-related press articles. UNIVBRIS, e.g. has previously published news items with ElectronicsWeekly and ElectronicsTalk, both reach a large community of electronics design professionals EU-wide.

Electronic publishing media such as email newsletters and blogs are increasingly popular instruments to reach the engineering and academic community. The ENTRA consortium will disseminate the project vision and research results to relevant newsletters such as the "Low Power Engineering" newsletter of the *Chip Design Magazine* and actively contribute towards professional blogs on publisher sites to further increase exposure of the ENTRA project.

Academic publications in Conferences, Workshops and Journals:

In the first phase, research is mainly explorative and will not yet have produced many firm or fundamental results. Thus, the dissemination activities are focused on making potential users of the results aware at an early stage, so that the dissemination in the second phase can be targeted. This also includes alerting the academic community and the industrial sector of the ENTRA project. The EACO workshop initiative at UNIVBRIS will be used to to disseminate "work in progress" at an early stage. The consortium will also seek opportunities to submit posters to workshops (other than EACO) to promote the project internationally.

Target conferences to publish project results mainly during the second phase include (but are not limited to):

- **DAC:** Design Automation Conference, held annually in the US for industry and academia.
- **DATE:** Design, Automation and Test in Europe, held annually in Europe.
- **CAV:** Computer Aided Verification, held annually and alternating between the US and Europe.
- **SAS:** International Static Analysis Symposium, held annually, mostly in Europe.
- **PEPM:** Partial Evaluation and Program Manipulation Symposium/Workshop, held annually.
- **TACAS:** Tools And Algorithms For The Construction And Analysis Of Systems, is an annual forum for researchers, developers and users interested in rigorously based tools and algorithms for the construction and analysis of systems.
- **IEEE SCAM:** Source Code Analysis and Manipulation, held annually and alternating between the US and Europe.
- **VMCAI:** Verification, Model Checking, and Abstract Interpretation, held annually and alternating between the US and Europe.
- **ICECCS:** IEEE International Conference on Engineering of Complex Computer Systems.
- **FM:** International Symposium on Formal Methods.
- **IGCC:** International Green Computing Conference.

There are also a number of conferences on parallelism and multi-core, such as Euro-Par, PaCT and SPAA (ACM Symposium on Parallelism in Algorithms and Architectures) where the results of ENTRA

can be promoted to stimulate innovation in these areas based on energy transparency.

Project partners have had papers in the past at the majority of the above conferences.

Journals targeted include ACM TODAES, the Science of Computer Programming, ACM TOPLAS, and New Generation Computing, amongst others.

Education material and Summer School:

Educating the next generation of system engineers is a fundamental cornerstone that safeguards the transformational impact we expect whole-systems energy transparency will have. The ENTRA consortium will develop material for a new MSc unit on Energy-Efficient Computing based on energy transparency. This unit can be integrated into the curriculum at the academic partners, e.g. the MSc in Advanced Microelectronic Systems Engineering at UNIVBRIS or the FIRST Research School (www.first.dk), which organizes a broad series of PhD-training activities within theoretical computer science and fundamental software technologies in the Copenhagen area where RUC is a partner. Teaching materials will be made available on the ENTRA web site for public access.

To keep the wider educational community up to date with developments in energy-efficient computing and energy transparency as an enabling technology, the ENTRA consortium will organize a summer school on "Whole-Systems Energy Transparency for Optimizing Energy Efficiency of Computing". This summer school will be run in the last 6 months of the project and, if successful, repeated after the project. UNIVBRIS already has an inter-disciplinary "Systems Centre" dedicated to advanced research into Systems Engineering with strong industrial collaborations. This provides one platform on which to run and organize the summer school(s).

Contribution towards Standards:

While ENTRA aims to provide a proof of concept, the basic principles of energy modelling, in particular the features of the Common Assertion Language, may make valuable contributions towards future standards for specifying energy consumption intent, requirements or actual values. A similar standardization effort has led to the "Standard for Design and Verification of Low Power Integrated Circuits", commonly known as "Unified Power Format (UPF)" and used for specifying power intent in power optimization of electronic design automation; it was published in March 2009 as IEEE Std 1801-2009. UNIVBRIS is well connected in the EDA community to make contributions towards standardization should this become relevant.

Open access:

In addition to Article II.30.4, beneficiaries shall deposit an electronic copy of the published version or the final manuscript accepted for publication of a scientific publication relating to foreground published before or after the final report in an institutional or subject-based repository at the moment of publication. Beneficiaries are required to make their best efforts to ensure that this electronic copy becomes freely and electronically available to anyone through this repository: a) immediately if the scientific publication is published "open access", i.e. if an electronic version is also available free of charge via the publisher, or b) within 6 months of publication.

B3.2.2 Exploitation

The development of a full Exploitation Plan (draft and final version) is part of the ENTRA work programme. This section lists the routes to exploitation of the ENTRA project results per partner.

Roskilde University

Roskilde University will use the results of the ENTRA project in teaching, research and the development of further collaborations with academia and industry. RUC's computer science PhD programme

(DMIT) is affiliated to the FIRST Research School (www.first.dk), which organises research training activities among universities in the greater Copenhagen area. Roskilde University also promotes research-driven teaching in undergraduate programmes with emphasis on project work, and it is intended to involve students in projects related to ENTRA. Finally the results of the project will be used to build new collaboration with industries in the Danish "Green IT" initiative.

University of Bristol

The University of Bristol leads a series of international workshops on *Energy-Aware COmputing (EACO)*[5]. The EACO Initiative brings together researchers and engineers with interests in energy-aware computing to exchange knowledge and to engage in collaborative research to advance the state of the art. EACO is an excellent route to exploitation for the ENTRA consortium. Internationally leading companies are engaged in EACO, including[6] ARM, IBM, ST Ericsson, ST Microelectronics, NVIDIA, Picochip, Cadence and XMOS.

UNIVBRIS also heads the Cadence Academic Network in Europe in the area of Advanced Verification Methodology and will use the links within this network to disseminate ENTRA research results. The undergraduate MEng in Computer Science and Electronics and the MSc in Advanced Microelectronic Systems Engineering are programs at Bristol that will directly benefit from the expertise of the staff involved in ENTRA. Students will be given the opportunity to undertake projects within the scope of ENTRA. In addition, UNIVBRIS regularly gives presentations at local industry events such as DVClub in the UK; the results of ENTRA will be presented at these events as a route to exploitation. UNIVBRIS has strong links to the local industry and regular hosts seminars by industrial visitors. These will be used to promote ENTRA research results and to actively explore new opportunities for future collaborations to stimulate innovation and to transfer knowledge and/or technology from ENTRA into industry.

IMDEA Software Institute

The IMDEA Software Institute regularly disseminates research results via presentations at international fora, and publications in journals and conference proceedings of high impact. The Institute has also close links to industry, including companies in the embedded software sector (e.g., DEIMOS and GTD), and has carried out successful technology transfer to industry. Forms of collaboration with industry include research projects, the participation of company staff in Institute activities, joint scholarships for doctorate or masters work, and transfer of research personnel trained by the institute to companies. Moreover, IMDEA is located in the Montegancedo Science and Technology Park, which has been the only campus in Spain to receive a "Campus of Excellence in Research and Technology Transfer" award in the Information and Communications Technologies area from the Ministry of Science and Innovation. This brings together an excellent framework for the transfer of the project's results to industry.

XMOS

XMOS will evaluate the technology and prototypes produced in the project for inclusion into its toolchain. The XMOS toolchain is available for free to XMOS customers; the toolchain is an essential requirement for customers, and chip sales are directly related to the quality of the toolchain provided. The original XMOS toolchain comprised an IDE with compiler, linker, simulator, and debugger, and was released in 2007 together with first silicon. In 2009 the toolchain was extended with a timing analyser, the XTA. The XTA enabled customers to, at compile-time, verify that threads would adhere to timing constraints.

The research in ENTRA enables three additions to the toolchain:

[5] http://www.cs.bris.ac.uk/Research/Micro/eaco.jsp
[6] http://www.cs.bris.ac.uk/Research/Micro/eaco-4.jsp

- A multi-threaded timing analyser.
- An energy analyser.
- A compiler that performs basic specialisation.

The precise tools that will be added will depend on what parts of the research are successful and the maturity of the prototypes developed. Although the tools are free, adding these features to the XMOS toolchain will significantly enhance the software design experience. This will lead to better design with XMOS chips, and enhance the competitive position of XMOS.

Creating Opportunities for Uptake and Technology Transfer

The ENTRA consortium will participate at several workshops aimed mainly at industrial participants to seek opportunities for uptake of the ENTRA vision and for knowledge and technology transfer to industry. At this stage we envision the following opportunities:

- **Silicon Southwest (SSW):** This is a network of over 100 companies connected to the microelectronics sector within 100 km of Bristol, UK, and is the second largest cluster of silicon design companies worldwide (after Silicon Valley in California). SSW includes a significant number of application developers.
- **NMI networking events:** The UK National Microelectronics Network organizes regular networking events for the microelectronics industry in the UK. The events on energy efficient computing or on verification provide excellent platforms to promote ENTRA.
- One of the EU Joint Undertaking Initiatives e.g. ARTEMIS or ENIAC. Their brokerage events provide ideal opportunities to promote ENTRA.
- **European CDN Live! event:** This is organized by the EDA company Cadence brings together EDA users, developers, and industry experts to network, share best practices on critical design and verification issues, and discover new techniques for realizing advanced silicon, SoCs, and systems.
- **University Booth at DAC or DATE:** These provide opportunities for academic institutions to demonstrate the latest research results alongside a major trade show that accompanies these prestigious international conferences.

B3.2.3 Management of Knowledge and Intellectual Property

Intellectual property is owned by the originator.

The availability of the XMOS design tools and simulator for free do not confer any rights to the Intellectual Property contained within them.

At the start of the project, each partner will provide the project manager with a list of intellectual property that is the background to their work in the project.

The project manager will periodically review the work in the project, especially work in joint tasks, for possible issues over ownership, and discuss these with the relevant partners to ensure there is clarity over who owns the IP developed.

The Consortium Agreement will detail how to Intellectual Property disputes will be resolved.

B4 Ethical Issues

Not applicable.

References

[1] E. Albert, P. Arenas, S. Genaim, and G. Puebla. Closed-Form Upper Bounds in Static Cost Analysis. *Journal of Automated Reasoning*, 46(2):161–203, February 2011.

[2] E. Albert, P. Arenas, S. Genaim, G. Puebla, and D. Zanardini. Cost Analysis of Object-Oriented Bytecode Programs. *Theoretical Computer Science*, 413(1):142–159, 2012.

[3] Jan Stoess Andreas Merkel and Frank Bellosa. Resource-conscious scheduling for energy efficiency on multicore processors. In *Proc. of EuroSys'10*, 2010.

[4] J. Ansel, C. Chan, Y. L. Wong, M. Olszewski, Q. Zhao, A. Edelman, and S. Amarasinghe. Petabricks: a language and compiler for algorithmic choice. In *Proc. of PLDI'09*, 2009.

[5] Woongki Baek and Trishul M. Chilimbi. Green: a framework for supporting energy-conscious programming using controlled approximation. In *PLDI*, pages 198–209, 2010.

[6] Gourinath Banda and John P. Gallagher. Constraint-based abstract semantics for temporal logic: A direct approach to design and implementation. In Edmund M. Clarke and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning - 16th International Conference, LPAR-16*, volume 6355 of *Lecture Notes in Computer Science*, pages 27–45. Springer, 2010.

[7] Steve Barker, Michael Leuschel, and Mauricio Varea. Efficient and flexible access control via jones-optimal logic program specialisation. *Higher-Order and Symbolic Computation*, 21(1-2):5–35, 2008.

[8] I. Bate, G. Bernat, and P. Puschner. Java Virtual-Machine Support for Portable Worst-Case Execution-Time Analysis. In *5th IEEE International Symposium on Object-oriented Real-time Distributed Computing, Washington, DC, USA*, Apr. 2002.

[9] R. Benzinger. Automated Higher-Order Complexity Analysis. *Theoretical Computer Science*, 318(1-2), 2004.

[10] B. Blanchet, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. A Static Analyzer for Large Safety-Critical Software. In *Proc. of PLDI'03*. ACM Press, 2003.

[11] David Brooks, Vivek Tiwari, and Margaret Martonosi. Wattch: a framework for architectural-level power analysis and optimizations. In *Proceedings of the 27th annual international symposium on Computer architecture*, ISCA '00, pages 83–94, New York, NY, USA, 2000. ACM.

[12] J. Chassin and P. Codognet. Parallel Logic Programming Systems. *Computing Surveys*, 26(3):295–336, September 1994.

[13] Charles Consel. A tour of schism: A partial evaluation system for higher-order applicative languages. In *PEPM*, pages 145–154, 1993.

[14] Charles Consel, Julia L. Lawall, and Anne-Françoise Le Meur. A tour of tempo: a program specializer for the C language. *Sci. Comput. Program.*, 52:341–370, 2004.

[15] P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. Combination of abstractions in the Astrée static analyzer. In M. Okada and I. Satoh, editors, *Proc. of the 11th Annual Asian Computing Science Conference (ASIAN'06)*, volume 4435 of *LNCS*, pages 272–300, Tokyo, Japan, 6–8 Dec. 2006. Springer (Berlin).

[16] Patrick Cousot and Radhia Cousot. Abstract interpretation frameworks. *J. Log. Comput.*, 2(4):511–547, 1992.

[17] S. K. Debray and N. W. Lin. Cost Analysis of Logic Programs. *ACM Transactions on Programming Languages and Systems*, 15(5):826–875, November 1993.

[18] S. K. Debray, N.-W. Lin, and M. Hermenegildo. Task Granularity Analysis in Logic Programs. In *Proc. of the 1990 ACM Conf. on Programming Language Design and Implementation*, pages 174–188. ACM Press, June 1990.

[19] S. K. Debray, P. López-García, M. Hermenegildo, and N.-W. Lin. Lower Bound Cost Estimation for Logic Programs. In *1997 International Logic Programming Symposium*, pages 291–305. MIT Press, Cambridge, MA, October 1997.

[20] David Delmas and Jean Souyris. Astrée: From research to industry. In Hanne Riis Nielson and Gilberto Filé, editors, *Static Analysis, 14th International Symposium, SAS 2007, Kongens Lyngby, Denmark, August 22-24, 2007, Proceedings*, volume 4634 of *Lecture Notes in Computer Science*, pages 437–451. Springer, 2007.

[21] Kerstin Eder and Geoff Barrett. Achieving maximum performance: A method for the verification of interlocked pipeline control logic. In *Proceedings of the 39th Design Automation Conference*, pages 135–140. ACM, June 2002.

[22] C. Edwards. Lack of Software Support marks the Low Power Scorecard at DAC. ElectronicsWeekly (http://www.electronicsweekly.com), No. 2472, 15-21 June 2011.

[23] Jochen Eisinger, Ilia Polian, Bernd Becker, Alexander Metzner, Stephan Thesing, and Reinhard Wilhelm. Automatic identification of timing anomalies for cycle-accurate worst-case execution time analysis. In *Proceedings of IEEE Workshop on Design & Diagnostics of Electronic Circuits & Systems (DDECS)*, pages 15–20. IEEE Computer Society, April 2006.

[24] Robert Bruce Findler, Shu-Yu Guo, and Anne Rogers. Lazy contract checking for immutable data structures. In Olaf Chitil, Zoltán Horváth, and Viktória Zsók, editors, *Implementation and Application of Functional Languages*, pages 111–128. Springer-Verlag, Berlin, Heidelberg, 2008.

[25] Harry D. Foster, Adam C. Krolnik, and David J. Lacey. *Assertion-based Design*. Kluwer Academic Publishers, 2nd edition edition, 2004.

[26] John P. Gallagher and Julio C. Peralta. Regular tree languages as an abstract domain in program specialisation. *Higher-Order and Symbolic Computation*, 14(2-3):143–172, 2001.

[27] Arne J. Glenstrup, Henning Makholm, and Jens P. Secher. C-MIX: Specialization of C programs. In John Hatcliff, Torben Æ. Mogensen, and Peter Thiemann, editors, *Partial Evaluation - Practice and Theory*, volume 1706 of *Lecture Notes in Computer Science*, pages 108–154. Springer, 1999.

[28] G. Gómez and Y. A. Liu. Automatic Time-Bound Analysis for a Higher-Order Language. In *Proceedings of the ACM SIGPLAN 2002 Workshop on Partial Evaluation and Semantics-Based Program Manipulation*, pages 75–88. ACM Press, 2002.

[29] Miguel Gómez-Zamalloa, Elvira Albert, and Germán Puebla. Decompilation of Java bytecode to Prolog by partial evaluation. *Information & Software Technology*, 51(10):1409–1427, 2009.

[30] É. Goubault. Static analyses of floating-point operations. In *Proc. of the 8th Int. Static Analysis Symposium (SAS'01)*, volume 2126 of *LNCS*, pages 234–259. Springer (Berlin), 2001.

[31] Elena Corina Grigore, Kerstin Eder, Alexander Lenz, Sergey Skachek, Anthony G. Pipe, and Chris Melhuish. Towards safe human-robot interaction. In Roderich Groß, Lyuba Alboul, Chris Melhuish, Mark Witkowski, Tony J. Prescott, and Jacques Penders, editors, *TAROS*, volume 6856 of *Lecture Notes in Computer Science*, pages 323–335. Springer, 2011.

[32] Sumit Gulwani, Krishna K. Mehra, and Trishul M. Chilimbi. Speed: precise and efficient static estimation of program computational complexity. In Zhong Shao and Benjamin C. Pierce, editors, *Proceedings of the 36th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2009*, pages 127–139. ACM, 2009.

[33] Simon Helsen and Peter Thiemann. Polymorphic specialization for ML. *ACM Trans. Program. Lang. Syst.*, 26(4):652–701, 2004.

[34] Kim S. Henriksen and John P. Gallagher. Abstract interpretation of PIC programs through logic programming. In *Sixth IEEE International Workshop on Source Code Analysis and Manipulation (SCAM 2006)*, pages 184–196. IEEE Computer Society, 2006.

[35] M. Hermenegildo, G. Puebla, F. Bueno, and P. López-García. Integrated Program Debugging, Verification, and Optimization Using Abstract Interpretation (and The Ciao System Preprocessor). *Science of Computer Programming*, 58(1–2), 2005.

[36] M. Hermenegildo and F. Rossi. Strict and Non-Strict Independent And-Parallelism in Logic Programs: Correctness, Efficiency, and Compile-Time Conditions. *Journal of Logic Programming*, 22(1):1–45, 1995.

[37] H. Hoffmann, S. Sidiroglou, M. Carbin, S. Misailovic, A. Agarwal, and M. Rinard. Dynamic knobs for responsive power-aware computing. In *Proc. of ASPLOS'11*. ACM Press, 2011.

[38] S.J. Hollis and C.R. Jackson. Skip the analysis: Self-optimising networks-on-chip (invited paper). In *Proc. International Symposium on Electronic System Design (ISED)*, 2010.

[39] E. Yu-Shing Hu, A. J. Wellings, and G. Bernat. Deriving Java Virtual Machine Timing Models for Portable Worst-Case Execution Time Analysis. In *On The Move to Meaningful Internet Systems 2003: OTM 2003 Workshops*, volume 2889 of *LNCS*, pages 411–424. Springer, October 2003.

[40] IEEE Computer Society. IEEE standard for binary floating-point arithmetic. Technical report, ANSI/IEEE Std. 745-1985, 1985.

[41] C. R. Jackson and S.J. Hollis. Skip-links: A dynamically reconfiguring topology for energy-efficient nocs. In *NoCs International Symposium on System-on-Chip*, 2010.

[42] Jesper Jørgensen. SIMILIX: A self-applicable partial evaluator for Scheme. In John Hatcliff, Torben Æ. Mogensen, and Peter Thiemann, editors, *Partial Evaluation - Practice and Theory*, volume 1706 of *Lecture Notes in Computer Science*, pages 83–107. Springer, 1999.

[43] Suleiman Abu Kharmeh, Kerstin Eder, and David May. Formal analysis of a programmable performance-critical processor communication interface. In *Proceedings of the 10th International Workshop on Automated Verification of Critical Systems (AVoCS 2010)*. tbc, October 2010.

[44] Vijay Anand Korthikanti and Gul Agha. Energy-bounded scalability analysis of parallel algorithms. In *Technical Report, Department of Computer Science, University of Illinois at Urbana Champaign*, 2009.

[45] Vijay Anand Korthikanti and Gul Agha. Energy-performance trade-off analysis of parallel algorithms. In *USENIX Workshop on Hot Topics in Parallelism (HotPar)*, 2010.

[46] Vijay Anand Korthikanti and Gul Agha. Towards optimizing energy costs of algorithms for shared memory architectures. In *Proceedings of the 22nd ACM symposium on Parallelism in algorithms and architectures*, SPAA '10, pages 157–165, New York, NY, USA, 2010. ACM.

[47] P.J. Krause, K. Craig-Wood, and N. Craig-Wood. Green ICT: Oxymoron, or call to innovation? In *Proc. Green IT*, Singapore, 2010.

[48] Sébastien Lafond and Johan Lilius. Energy consumption analysis for two embedded Java virtual machines. *J. Syst. Archit.*, 53(5-6):328–337, 2007.

[49] D. Le Metayer. ACE: An Automatic Complexity Evaluator. *ACM Transactions on Programming Languages and Systems*, 10(2), 1988.

[50] Michael Leuschel. A framework for the integration of partial evaluation and abstract interpretation of logic programs. *ACM Trans. Program. Lang. Syst.*, 26(3):413–463, 2004.

[51] Michael Leuschel, Stephen Craig, Maurice Bruynooghe, and Wim Vanhoof. Specializing interpreters using offline partial deduction. In Maurice Bruynooghe and Kung-Kiu Lau, editors, *Program Development in Computational Logic*, LNCS 3049, pages 341–376. Springer-Verlag, 2004.

[52] Michael Leuschel, Daniel Elphick, Mauricio Varea, Stephen-John Craig, and Marc Fontaine. The Ecce and Logen partial evaluators and their web interfaces. In John Hatcliff and Frank Tip, editors, *Proceedings of the 2006 ACM SIGPLAN Workshop on Partial Evaluation and Semantics-based Program Manipulation*, pages 88–94. ACM, 2006.

[53] Michael Leuschel and Helko Lehmann. Solving coverability problems of petri nets by partial deduction. In *PPDP*, pages 268–279, 2000.

[54] P. López-García. *Non-failure Analysis and Granularity Control in Parallel Execution of Logic Programs.* PhD thesis, Universidad Politécnica de Madrid (UPM), Facultad Informatica UPM, 28660-Boadilla del Monte, Madrid-Spain, June 2000.

[55] P. López-García, L. Darmawan, and F. Bueno. A Framework for Verification and Debugging of Resource Usage Properties. In M. Hermenegildo and T. Schaub, editors, *Technical Communications of the 26th Int'l. Conference on Logic Programming (ICLP'10)*, volume 7 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 104–113, Dagstuhl, Germany, July 2010. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik.

[56] P. López-García and M. Hermenegildo. Efficient Term Size Computation for Granularity Control. In *ICLP'95*, 1995.

[57] P. López-García, M. Hermenegildo, and S. K. Debray. A Methodology for Granularity Based Control of Parallelism in Logic Programs. *Journal of Symbolic Computation, Special Issue on Parallel Symbolic Computation*, 21(4–6):715–734, 1996.

[58] Mingsong Lv, Nan Guan, Wang Yi, and Ge Yu. Combining abstract interpretation with model checking for timing analysis of multicore software. In *31st IEEE Real-Time Systems Symposium, San Diego*, 2010.

[59] M. Martel. Enhancing the implementation of mathematical formulas for fixed-point and floating-point arithmetics. *Formal Methods in System Design*, 35(3):265–278, Dec. 2009.

[60] Matthieu Martel. Semantics of roundoff error propagation in finite precision calculations. *Higher-Order and Symbolic Computation*, 19(1):7–30, 2006.

[61] E. Mera, P. López-García, M. Carro, and M. Hermenegildo. Towards Execution Time Estimation in Abstract Machine-Based Languages. In *10th Int'l. ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming (PPDP'08)*, pages 174–184. ACM Press, July 2008.

[62] E. Mera, P. López-García, G. Puebla, M. Carro, and M. Hermenegildo. Combining Static Analysis and Profiling for Estimating Execution Times. In *Ninth International Symposium on Practical Aspects of Declarative Languages (PADL'07)*, number 4354 in LNCS, pages 140–154. Springer-Verlag, January 2007.

[63] Bertrand Meyer. *Eiffel: the language*. Prentice-Hall, Inc., Upper Saddle River, NJ, USA, 1992.

[64] J. Navas, M. Méndez-Lojo, and M. Hermenegildo. Safe Upper-bounds Inference of Energy Consumption for Java Bytecode Applications. In *The Sixth NASA Langley Formal Methods Workshop (LFM 08)*, April 2008. Extended Abstract.

[65] J. Navas, M. Méndez-Lojo, and M. Hermenegildo. User-Definable Resource Usage Bounds Analysis for Java Bytecode. In *Proceedings of the Workshop on Bytecode Semantics, Verification, Analysis and Transformation (BYTECODE'09)*, volume 253 of *Electronic Notes in Theoretical Computer Science*, pages 6–86. Elsevier - North Holland, March 2009.

[66] J. Navas, E. Mera, P. López-García, and M. Hermenegildo. User-Definable Resource Bounds Analysis for Logic Programs. In *23rd International Conference on Logic Programming (ICLP'07)*, volume 4670 of *Lecture Notes in Computer Science*. Springer, 2007.

[67] J.L Nunez-Yanez and V.A. Chouliaras. A configurable statistical lossless compression core based on variable order modelling and arithmetic coding. *IEEE Transactions on Computers*, 54(11):1345–1359, Nov 2005.

[68] J.L Nunez-Yanez, V.A. Chouliaras, D. Alfonso, and F. Rovati. Hardware assisted rate distortion optimization with embedded cabac accelerator for the h.264 advanced video codec. *IEEE Transactions on Consumer Electronics*, 52(2):590–598, May 2006.

[69] Jose L. Nunez-Yanez, T. Spiteri, and G. Vafiadis. Multi-standard reconfigurable motion estimation processor for hybrid video codecs. *IET Computers & Digital Techniques*, 5(2):73–85, March 2011.

[70] Claudio Ochoa and Germán Puebla. Poly-controlled partial evaluation in practice. In G. Ramalingam and Eelco Visser, editors, *Proceedings of the 2007 ACM SIGPLAN Workshop on Partial Evaluation and Semantics-based Program Manipulation*, pages 164–173. ACM, 2007.

[71] Young-Hwan Park, Sudeep Pasricha, Fadi J. Kurdahi, and Nikil Dutt. Methodology for multi-granularity embedded processor power model generation for an esl design flow. In *Proceedings of the 6th IEEE/ACM/IFIP international conference on Hardware/Software codesign and system synthesis*, CODES+ISSS '08, pages 255–260, New York, NY, USA, 2008. ACM.

[72] P. Pietrzak and M. Hermenegildo. Automatic Binding-related Error Diagnosis in Logic Programs. In *International Conference on Logic Programming (ICLP'07)*, number 4670 in LNCS, pages 333–347. Springer-Verlag, September 2007.

[73] C. Preist and P. Shabajee. Energy use in the media cloud: Behaviour change, or technofix? In *Cloud Computing Technology and Science (CloudCom), 2010 IEEE Second International Conference on*, pages 581 –586, 30 2010-dec. 3 2010.

[74] Germán Puebla, Manuel V. Hermenegildo, and John P. Gallagher. An integration of partial evaluation in a generic abstract interpretation framework. In *Proceedings of the 1999 ACM SIGPLAN Workshop on Partial Evaluation and Semantics-based Program Manipulation*, 1999.

[75] M. Rinard. Probabilistic accuracy bounds for fault-tolerant computations that discard tasks. In *Proc. of ICS'06*. ACM Press, 2006.

[76] M. Rinard. Using early phase termination to eliminate load imbalances at barrier synchronization points. In *Proc. of OOPSLA'07*. ACM Press, 2007.

[77] R. Lublinerman S. Chaudhuri, S. Gulwani and S. Navidpour. Proving programs robust. In *Proc. of FSE'11*, 2012.

[78] Adrian Sampson, Werner Dietl, Emily Fortuna, Danushen Gnanapragasam, Luis Ceze, and Dan Grossman. Enerj: approximate data types for safe and general low-power computation. In *Proc. of PLDI'11*. ACM Press, 2011.

[79] Daniel M. Roy Sasa Misailovic and Martin C. Rinard. Probabilistically accurate program transformations. In *Proc. of SAS'11*, volume 6887 of *Lecture Notes in Computer Science*. Springer, 2011.

[80] Henry Hoffmann Sasa Misailovic, Stelios Sidiroglou and Martin Rinard. Quality of service profiling. In *Proc. of ICSE'10*. ACM Press, 2010.

[81] Henry Hoffmann Sasa Misailovic, Stelios Sidiroglou and Martin Rinard. Managing performance vs. accuracy trade-offs with loop perforation. In *Proc. of FSE'11*. ACM Press, 2011.

[82] Ulrik Pagh Schultz, Julia L. Lawall, and Charles Consel. Automatic program specialization for Java. *ACM Trans. Program. Lang. Syst.*, 25(4):452–499, 2003.

[83] Chiyoung Seo, Sam Malek, and Nenad Medvidovic. An energy consumption framework for distributed Java-based systems. In *ASE '07: Proceedings of the twenty-second IEEE/ACM international conference on Automated software engineering*, pages 421–424, New York, NY, USA, 2007. ACM.

[84] Silicon Integration Initiative (Si2). *A Practical Guide to Low-Power Design*. `http://www.si2.org/?page=1061`.

[85] David C. Snowdon, Etienne Le Sueur, Stefan M. Petters, and Gernot Heiser. Koala: a platform for os-level power management. In *Proc. of EuroSys'09*. ACM Press, 2009.

[86] C. Taylor and J. Koomey. Estimating energy use and greenhouse gas emissions of Internet advertising. Working paper for IMC2. Retrieved from http://evanmills.lbl.gov/commentary/docs/carbonemissions.pdf, February 2008.

[87] Lothar Thiele and Reinhard Wilhelm. Design for time-predictability. In *Perspectives Workshop: Design of Systems with Predictable Behaviour, 16.-19. November 2003*, volume 03471 of *Dagstuhl Seminar Proceedings*. IBFI, Schloss Dagstuhl, Germany, 2004.

[88] Vivek Tiwari, Sharad Malik, Andrew Wolfe, and Mike Tien-Chien Lee. Instruction level power analysis and optimization of software. *The Journal of VLSI Signal Processing*, 13:223–238, 1996. 10.1007/BF01130407.

[89] Douglas Watt. *Programming XC on XMOS Devices*. XMOS Ltd., 2009.

[90] Andreas Weissel and Frank Bellosa. Process cruise control-event-driven clock scaling for dynamic power management. In *Proceedings of the International Conference on Compilers, Architecture and Synthesis for Embedded Systems (CASES 2002)*, Grenoble, France, October 8–11 2002.

[91] Reinhard Wilhelm. Timing Analysis and Timing Predictability. In Frank S. de Boer, Marcello M. Bonsangue, Susanne Graf, and Willem P. de Roever, editors, *Formal Methods for Components and Objects, Third International Symposium (FMCO)*, volume 3657 of *LNCS, Revised Lectures*, pages 317–323. Springer, 2004.

[92] Steve Wright and Kerstin Eder. Using event-b to construct instruction set architectures. *Formal Aspects of Computing*, 23(1):73–79, January 2010.

[93] XMOS, editor. *XMOS Timing Analyzer Whitepaper*. XMOS Ltd., 2010.

[94] Fangfang Yuan, Stephen Wright, Kerstin Eder, and David May. Managing complexity through abstraction: A refinement-based approach to formalize instruction set architectures. In Shengchao Qin and Zongyan Qiu, editors, *ICFEM*, volume 6991 of *Lecture Notes in Computer Science*, pages 585–600. Springer, 2011.

[95] J. Kelner Z. Zhu, S. Misailovic and M. Rinard. Optimal accuracy aware transformations for approximate computations. In *Proc. of POPL'12*, 2012.

[96] Yiwei Zhang, Jose Nunez-Yanez, Joe P. McGeehan, Stephen Kelly, and Ed Regan. Biophysically accurate foating point neuroprocessors for reconfigurable logic. *IEEE Transactions on Computers*, Dec 2011. http://doi.ieeecomputersociety.org/ 10.1109/TC.2011.257.