Analysis of Multi-Objective Evolutionary Algorithms to Optimize Dynamic Data Types in Embedded Systems

J. Ignacio Hidalgo, José L. Risco-Martín, David Atienza, Juan Lanchares
Department of Computer Architecture and Automation
Complutense University of Madrid
Prof. José García Santesmases, s/n, 28040 Madrid, Spain
{hidalgo, jlrisco, datienza, julandan}@dacya.ucm.es

ABSTRACT

New multimedia embedded applications are increasingly dynamic, and rely on Dynamically-allocated Data Types (DDTs) to store their data. The optimization of DDTs for each target embedded system is a time-consuming process due to the large design space of possible DDTs implementations. Thus, suitable exploration methods for embedded design metrics (memory accesses, memory usage and power consumption) need to be developed. In this work we present a detailed analysis of the characteristics of different types of Multi-Objective Evolutionary Algorithms (MOEAs) to tackle the optimization of DDTs in multimedia applications and compare them with other state-of-the-art heuristics. Our results with state-of-the-art MOEAs in two object-oriented multimedia embedded applications show that more sophisticated MOEAs (SPEA2 and NSGA-II) offer better solutions than simple schemes (VEGA). Moreover, the suitable sophisticated scheme varies according to the available exploration time, namely, NSGA-II outperforms SPEA2 in the first set of solutions (300-500 generations), while SPEA2 offers better solutions afterwards.

Categories and Subject Descriptors

C.3 [**Computer Systems Organization**]: Real-time and embedded systems; I.2.8 [**Artificial Intelligence**]: Problem Solving, Control Methods, Search—*Heuristic methods*

General Terms

Design, Performance

Keywords

Multi-Objective Optimization, Pareto Optimal Front, Evolutionary Computation, Embedded Systems Design

1. INTRODUCTION

Latest multimedia embedded devices are enhancing its capabilities and are currently able to run applications re-served to powerful desktop computers few years ago (e.g., 3D games, video players). As a result, one of the most important problems designers face nowadays is the integration of a great amount of applications coming from the general-purpose domain in a compact and highly-constrained device. One major task of this porting process is the optimization of the dynamic memory subsystem. Thus, the designer must choose among a number of possible dynamically-allocated data structures or *Dynamic Data Types (DDTs)* implementations [1] (dynamic arrays, linked lists, etc.) the best one in each case, according to the specific restrictions of the target device and typical embedded design metrics, such as memory accesses, memory usage and energy consumption [2].

This task is typically performed using a pseudo-exhaustive evaluation of the design space of DDTs, including multiple executions of the application, to attain a Pareto's front [5], which tries to cover all the optimal implementation points for the aforementioned required design metrics. The construction of this Pareto's front is a very time-consuming process, sometimes even unaffordable [5], [12].

In this paper we analyze the potential benefits of the different state-of-the-art *Multi-Objective Evolutionary Algorithms (MOEAs)* [6] to explore the design space of DDT implementation. Our experiments in two real-life dynamic embedded applications using three representative MOEAs, namely *Vector Evaluated Genetic Algorithm (VEGA)* [14], *Strength Pareto Evolutionary Algorithm 2 (SPEA2)* [19] and *Non-dominated Sorting Genetic Algorithm II (NSGA-II)* [7], show that they can achieve significant speed-ups (up to 12000× faster in the case of VDrift and 3800× faster in the case of a 3D Physics engine) with respect to other traditional heuristics to achieve optimal DDT implementations for multiple metrics. Furthermore, our results indicate that the selection of the most suitable type of MOEA depends on the design constraints and devoted exploration time.

The rest of the paper is organized as follows. Section 2 reviews the work related to memory optimizations for embedded systems. In Section 3, we present our multi-objective optimization process. A description of the three representative MOEAs algorithms used in this work is given in Section 4. Section 5 details our experimental setup. In Section 6, we discuss the obtained results in two real-life embedded applications. Finally, Section 7 summarizes the main conclusions of this paper.

2. RELATED WORK

Up today extensive work has been performed in the field of embedded memory subsystem optimization. [12] and [3]

include a thorough survey of static data and memory optimization techniques for embedded systems. In [4] and [18], authors have explored a coordinated data and computation reordering for array-based data structures in multimedia applications. They use a linear time algorithm reducing the memory subsystems needs by 50%. Nevertheless, they are not suitable for exploration of complex DDTs employed in modern multimedia applications.

Regarding dynamic embedded software, suitable access methods, power-aware DDT transformations and pruning strategies based on heuristics have started to be proposed for multimedia systems [18][12]. However, these approaches require the development of efficient pruning function costs and fully manual optimizations [5]; otherwise they are not able to capture the evaluation of inter-dependencies of multiple DDTs implementations operating together, as our proposed method using evolutionary computation achieves. Also, in [2] it has already been outline the potential use of MOEAs for dynamic memory optimizations. Nevertheless, this work only performed an initial analysis of one type of MOEA and does not provide a complete analysis of trade-offs between different technologies of MOEAs, as we perform in this paper.

Also, according to the characteristics of certain parts of multimedia applications, several transformations for DDTs and design methodologies [4][18] have been proposed for static data profiling and optimization considering static memory access patterns to physical memories. In this context, the use of MOEA-based optimization has been applied to solve linear and non-linear problems by exploring all regions of the state space in parallel. Thus, it is possible to perform optimizations in non-convex regular functions, and also to select the order of algorithmic transformations in concrete types of source codes [11][12]. However, such techniques are not applicable in DDT implementations, due to the unpredictable nature at compile-time of the stored data.

3. THE DYNAMIC DATA TYPES EXPLORATION PROBLEM

A DDT is a software abstraction by means of which we can manipulate and access data. The implementation of a DDT has two main components. First, it has storage aspects that determine how data memory is allocated and freed at runtime and how this memory is tracked. Second, it includes an access component, which can refer to two different basic access patterns: sequential or iterator-based and random access.

Figure 1 shows an example of a DDTs exploration. The initial code contains two variables, `v1` and `v2`, instantiated as vector and list, respectively. After the exploration process, one candidate solution gives `v1` to be instantiated as *Single Linked List (SLL)* and `v2` as *Double Linked List of Arrays (DLLAR)*. Such instantiation policy tries to minimize memory accesses, memory usage and energy consumption of the final application. It is important to stress that it is unmanageable for the designer to get a totally complete exploration of all the possible DDT implementation combinations using the traditional way for real-life complex applications. For example, in the case of the 3D Physics engine, we optimized memory accesses, memory usage and power consumption for more than 3000 variables.

In general terms, the application to optimize contains a set

```
vector<T1>* v1 = new vector<T1>();
// …
list<T2>* v2 = new list<T2>();
list<T2>::iterator itr = v2->begin();
for(; itr!=v2->end(); ++itr)
  cout << *itr;
// …
```
Initial application code

```
SLL<T1>* v1 = new SLL<T1>();
// …
DLLAR<T2>* v2 = new DLLAR<T2>();
DLLAR<T2>::iterator itr = v2->begin();
for(; itr!=v2->end(); ++itr)
  cout << *itr;
// …
```
Optimized application code

Figure 1: Code before and after the exploration of Dynamic Data Types

of m variables ($\{V\}$) which are candidates to be instantiated as a certain DDT from the set of possible implementation of DDTs library ($\{D\}$) presented in [2][5]. Thus, the goal of our optimization flow is to obtain a set of pairs (variable, DDT) or $(\vec{v}, \vec{d}), v_i \in \{V\}, d_i \in \{D\}, 1 \leq i \leq m$, such that minimizes three objectives: memory accesses, memory usage and energy consumption. Additional constraints, such as minimum and maximum values for all three objectives may be defined.

The proposed optimization framework uses three different phases to perform the automatic exploration of DDTs using MOEAs. Figure 2 shows the different phases required to perform the overall DDTs optimization. In the first phase, we generate an initial profiling of the iterator-based access methods to the different DDTs used in the application. In the second phase, using this detailed report of the accesses, we extract all the information needed by the optimization phase. Finally, an exploration of the design space of DDTs implementation is performed using the genetic algorithm selected (NSGA-II, SPEA2 or VEGA).

Next, we describe the three phases of our flow in detail.

3.1 Profiling

Table 1: DDT library

DDT	Description
AR	Array
AR(P)	Array of pointers
SLL	Single linked list
DLL	Double linked list
SLL(O)	Single linked list with roving pointer
DLL(O)	Double linked list with roving pointer
SLL(AR)	Single linked list of arrays
DLL(AR)	Double linked list of arrays
SLL(ARO)	Single linked list of arrays and roving pointer
DLL(ARO)	Double linked list of arrays and roving pointer

In a first pre-characterization phase we define the equations to evaluate the behavior of DDT implementations by means of parameters such as the number of sequential accesses, random accesses , average size , etc. In our case
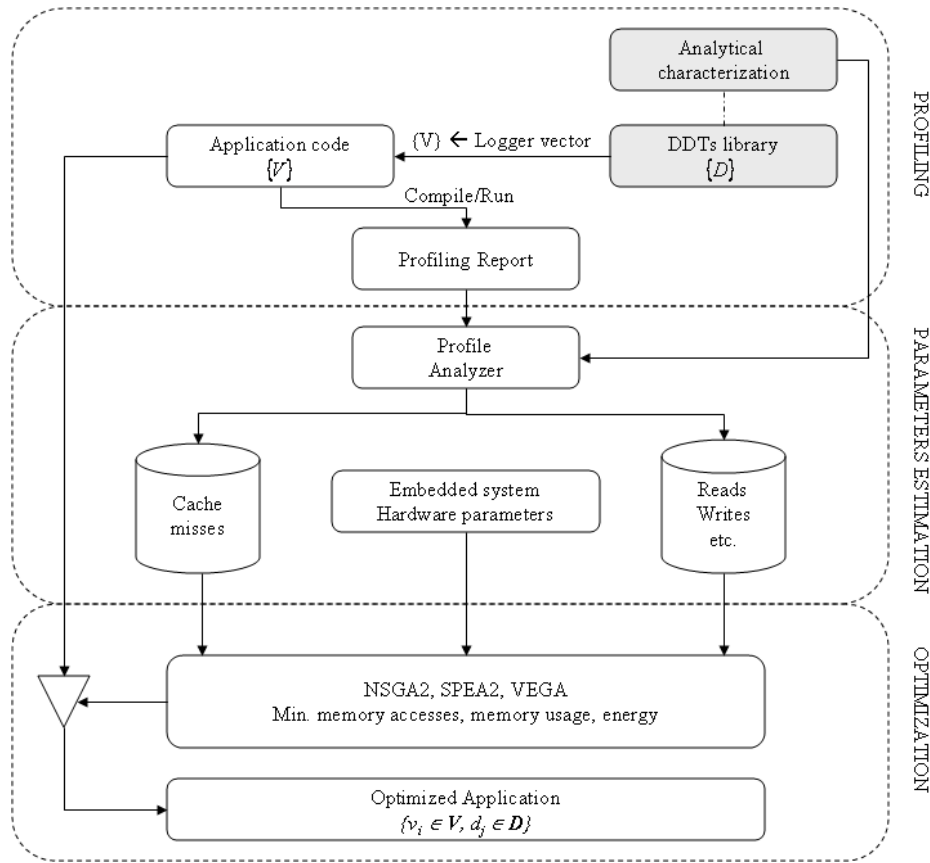
Figure 2: DDTs optimization flow

we have classified the DDT implementations in basic DDT and multi-layer implementations relevant for embedded multimedia applications. Table 1 contains the DDTs implemented [2].

Next, we obtain a profiling report of the application where the following information is logged: accessing of an element, addition of an element, removal of an element, the clearing of the container, iterator operations such as pre-increment or dereference, constructor, destructor, copy constructor and swap operation. To this end, we replace all the candidate variables in the application by our vector DDT implementation, which logs all the needed information. Such replacement is done automatically, where the designer may choose the variables to include in the optimization.

The left side of Figure 3 shows an illustrative example on how the profiling report logs all the variables selected in the original application. The report shows two entries, labeled as LOG_CONSTRUCT_BEGIN, where two variables are instantiated and identified as VariableId 1 and 2. Such entries state that the variables will contain elements of size 12 and 32 Bytes, respectively. Another entry, labeled as LOG_ADD_BEGIN, shows that an element is added to the variable 1, at index 0. The last entry shows that an iterator accesses to a element stored in variable 2, and located at the address 0x85b7290.

3.2 Parameters estimation

In the following phase, we extract all information needed from the profiling report. The purpose is to measure the
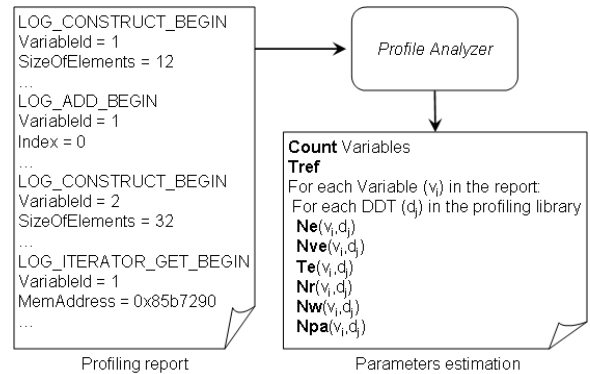

Figure 3: Example of profiling report and parameters estimation

quality of a candidate solution $\vec{v}, \vec{d}$ in the DDT exploration, using parameters such as the number of candidate variables ($Count$ in Figure 3), number of elements stored in the DDT in the worst case ($Ne(\vec{v}, \vec{d})$), average of the number of elements stored ($Nve(\vec{v}, \vec{d})$), size of the elements ($Te(\vec{v}, \vec{d})$, in bytes), size of the pointers ($Tref$, in bytes), number of read accesses ($Nr(\vec{v}, \vec{d})$), number of write accesses ($Nw(\vec{v}, \vec{d})$) and cache misses ($Npa(\vec{v}, \vec{d})$). All the parameters except the last one are obtained with just the information logged in the profiling report and the analytical characterization. Cache

misses are obtained by means of simulation, generating memory traces from the profiling report and the DDT library, using them as input for the *Dinero IV* cache simulator [8] for the particular memory configuration of the target embedded system. We have developed a tool called *Profile Analyzer*, which automates all the process. The right side of Figure 3 depicts an illustrative example on how the Profile Analyzer saves such parameters in an external file, used by the optimization phase.

3.3 Optimization

The last phase is the optimization process. It takes as input the parameters obtained in the previous phase and minimizes three objectives: memory accesses (MA), memory usage (MU) and energy (E), defined by equation 1, where Hw represents the effect that hardware parameters (memory architecture, CPU power, line sizes, memory access time, etc.) have on the optimization, and $(\vec{v}, \vec{d})$ represents one candidate solution.

$$
\begin{aligned}
MA(\vec{v}, \vec{d}) &= f_{MA}(Ne, Nve, Nr, Nw) \\
MU(\vec{v}, \vec{d}) &= f_{MU}(Te, Tref, Ne, Nve) \\
E(\vec{v}, \vec{d}) &= f_E(Nr, Nw, Npa, Hw)
\end{aligned}
\tag{1}
$$

In equation 1, f_{MA} and f_{MU} were taken from [2]. Energy equation of the system is given by the following equation:

$$
\begin{aligned}
f_E &= t_{ex} \times CPU_{pow} + \\
&\quad (N_r + N_w) \times (1 - N_{pa}) \times C_{accE} + \\
&\quad (N_r + N_w) \times N_{pa} \times C_{accE} \times C_{lineS} + \\
&\quad (N_r + N_w) \times N_{pa} \times DRAM_{accP} \times \\
&\quad \left(DRAM_{accT} + \frac{C_{lineS}}{DRAM_{bandW}} \right)
\end{aligned}
\tag{2}
$$

where,

- t_{ex} is the system's total execution time.

- CPU_{pow} is the total processor power excluding the cache power.

- C_{accE} is the cache access energy.

- C_{lineS} is the cache line size.

- $DRAM_{accP}$ is the active power consumed by the DRAM.

- $DRAM_{accT}$ is the DRAM latency time.

- $DRAM_{bandW}$ is the bandwidth of the DRAM.

There exist four components in the energy equation 2. The first term $t_{ex} \times CPU_{pow}$ calculates the processor energy given that execution time takes t_{ex} amount of time. The second term, $(N_r + N_w) \times (1 - N_{pa}) \times C_{accE}$ calculates the amount of energy consumed by the cache. The third term, $(N_r + N_w) \times N_{pa} \times C_{accE} \times C_{lineS}$ calculates the energy cost of writing to cache for each cache miss. The last term, calculates the energy cost of the DRAM to service all the cache misses.

The equation for calculating the system's total execution time t_{ex} is given by:

$$
\begin{aligned}
t_{ex} &= (N_r + N_w) \times (1 - N_{pa}) \times C_{accT} + \\
&\quad (N_r + N_w) \times N_{pa} \times DRAM_{accT} + \\
&\quad (N_r + N_w) \times N_{pa} \times \frac{C_{lineS}}{DRAM_{bandW}} + \\
&\quad T_{bus}
\end{aligned}
\tag{3}
$$

where C_{accT} is the access time of the cache.

There exist four components in the system's execution time shown in equation 3. The first term $(N_r + N_w) \times (1 - N_{pa}) \times C_{accT}$ is for calculating the amount of time taken for the processor to access the cache. The second term $(N_r + N_w) \times N_{pa} \times DRAM_{accT}$ calculates the amount of time required for the DRAM to respond to each cache miss. The third term calculates the amount of time taken to fill a cache line on each cache miss. The bus communication time cost is supposed to be constant (T_{bus}). As the bus communication time is expected to be similar to other systems, such decision will not adversely affect the final results.

Units for time variables in the equations are in seconds, bandwidth is in Bytes/sec., cache line size is in Bytes, power variable is in Watts, and energy unit is in Joules.

To explore this multi-objective optimization problem, we have implemented three popular MOEAs, called VEGA, SPEA2 and NSGA-II. VEGA belongs to the so-called first generation multi-objective evolutionary algorithms and SPEA2 and NSGA-II belong to the second generation. They are discussed next.

When the optimization process ends, it gives the DDT instantiation policy, i.e., which variable should be instantiated by which DDT. We also obtain the gain on memory accesses, memory usage and energy consumption.

4. SELECTION OF MULTI-OBJECTIVE EVOLUTIONARY ALGORITHMS

Multi-objective optimization aims at simultaneously optimizing several contradictory objectives. For such kind of problems, there does not exist a single optimal solution, and some compromises have to be made. We can assume the following N-objective minimization problem:

$$
\begin{aligned}
Minimize \quad &\vec{z} = (f_1(\vec{x}), f_2(\vec{x}), \dots f_N(\vec{x})) \\
subject\ to \quad &\vec{x} \in S
\end{aligned}
\tag{4}
$$

where $\vec{z}$ is the objective vector with N objectives to be minimized, $\vec{x}$ is the decision vector, and $S \subset \Re^m$ is the feasible region in the decision space. A solution $\vec{x} \in S$ is said to dominate another solution $\vec{y} \in S$ (denoted as $\vec{x} \prec \vec{y}$) if the following two conditions are satisfied.

$$
\begin{aligned}
&f_i(\vec{x}) \leq f_i(\vec{y}), \forall i \in \{1, 2, \dots, N\} \\
&f_i(\vec{x}) < f_i(\vec{y}), \exists i \in \{1, 2, \dots, N\}
\end{aligned}
\tag{5}
$$

A decision vector $\vec{x} \in S$ is called *Pareto-optimal* if there does not exist another $\vec{y} \in S$ that dominates it. An objective vector is called Pareto-optimal if the corresponding decision vector is Pareto-optimal. The non-dominated set of the entire feasible search space S is the *Pareto-optimal set (POS)*. The Pareto-optimal set in the objective space is called *Pareto-optimal front (POF)* of the multi-objective problem at hand: it represents the best possible compromises with respect to the contradictory objectives.

A Multi-objective optimization problem is solved, when all its POS is found. Classical optimization methods generally find one of the Pareto optimal solutions by making the initial optimization problem single-objective. Unlike classical methods, MOEAs directly search for the whole POF, allowing decision makers to choose one of the Pareto solutions with more complete information [6].

Up to date, many MOEAs have been developed. Generally speaking, they can be classified into two broad categories: non-elitism and elitism, also called first and second generation evolutionary algorithms. With the elitism approach, MOEAs store in an external set the best solutions of each generation. This set will then be a part of the next generation. Thus, the best individuals in each generation are always preserved, and this helps the algorithm to get close to its POF. Algorithms such as PEAS II, MOMGA II, NSGA-II and SPEA2 are examples of this category. In contrast, the non-elitism approach does not guarantee preserving the set of best individuals for the next generation [20]. Examples of this category include MOGA, HLGA, NPGA and VEGA.

For the purpose of this research, three representative algorithms have been selected to cover the different technologies of evolutionary computation when comparisons are performed. VEGA [20] was selected as a very optimized example of non-elitist MOEAs, where the the concept of POF is not directly incorporated into the selection mechanism of this algorithm. SPEA2 and NSGA-II were selected as two representative elitist-based MOEA [19][7].

5. EXPERIMENTAL METHODOLOGY

In this section we describe the complete method applied to compare the different type of MOEAs while optimizing two real-life dynamic embedded applications.
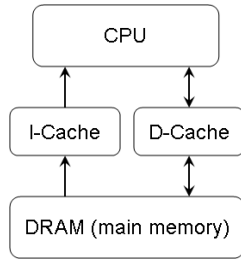
5.1 Embedded System HW/SW Specification

Figure 4: System architecture

The model of the embedded system architecture consisted of a processor with an instruction cache, a data cache, and embedded DRAM as main memory. The data cache uses a write-through strategy. The system architecture is illustrated in Figure 4.

Table 2: System specification

Processor Energy	168mW, 100MHz
Embedded DRAM	100MHz
Energy	19.5 mW
Latency	19.5 ns
Bandwidth	50MB/s

To analyze the effect of MOEAs on embedded system's memory accesses, memory usage and energy consumption, we utilized processor energy from [4], and the access time and energy values for caches of 32KB and embedded 16MB DRAM main memory from [16] and [9], respectively. The processor and memory specification is described in Table 2.

In this paper, we apply MOEAs to two multimedia embedded applications. The first benchmark is VDrift, which is a driving simulation game. The game includes 19 tracks, 28 cars, AI players, networked multiplayer mode, etc. [17]. We logged 49 variables in its source code. The second benchmark is a 3D Physics Engine for elastic and deformable bodies [10], which is a 3D engine that displays the interaction of non-rigid bodies. It includes 3128 dynamic variables in its source code for which we select the optimal DDT implementation.

5.2 Performance Metrics

To compare the performance of different MOEAs, we need to evaluate the obtained set of non-dominated solutions considering: (1) Convergence to POF. (2) Diversity on POF. Since the size of possible DDT implementations is large and it is not possible to cover the exact set of the POF, we compare the obtained *Pareto Front (PF)* with each other. In this direction, we select the following metrics to evaluate the performance of our approach.

Coverage (C): We use the coverage metric [20] to measure convergence. Let PF', PF'' be two sets of non-dominated solutions. The coverage metric can be defined as follows:

$$C(PF', PF'') = \frac{|p'' \in PF''; \exists p' \in PF' : p' \preceq p''|}{|PF''|} \quad (6)$$

The value $C(PF', PF'') = 1$ means that all points in PF'' are dominated by or equal to points in PF'. On the other hand, $C(PF', PF'') = 0$ means that no solutions in PF'' are covered by the set PF'. Both $C(PF', PF'')$ and $C(PF'', PF')$, have to be considered, since $C(PF', PF'')$ is not necessary equal to $C(PF'', PF')$. If $C(PF', PF'') > C(PF'', PF')$, the rate of dominated solutions in PF'' is higher than in PF'.

Spread (D): A spread metric (D) determines in each objective space the maximum range represented by the non-dominated solutions. It was introduced by Ranjithan [13]. A higher value of the spread metric indicates a better performance. It is defined as:

$$D = \sqrt{\sum_{i=1}^{N} \left(\max_{j=1}^{|PF|} f_i(x_j) - \min_{j=1}^{|PF|} f_i(x_j) \right)^2}$$
$$x_j \in PF, j = 1, 2, \ldots, |PF| \quad (7)$$

where N is the number of objectives.

Spacing (S): Schott proposed a metric which allows to measure the distribution of vectors throughout PF [15]. It is defined as:

$$S = \sqrt{\frac{1}{|PF|} \sum_{j=1}^{|PF|} \left(d_j - \bar{d} \right)^2}$$
$$d_j = \min_{x_k \in PF \wedge k \neq j} \sum_{i=1}^{N} |f_i(x_j) - f_i(x_k)| \quad (8)$$

where N is the number of objectives, and $\bar{d}$ is the mean of all d_j. A zero value for this metric means that all members of PF are equidistantly spaced.

We compare the obtained sets of non-dominated solutions by means of the above three criteria.

5.3 Coding a solution

In order to apply a MOEA correctly we need to define a genetic representation of the design space of all possible DDT implementations alternatives. Moreover, to be able to cover all possible inter-dependencies of DDT implementations for different dynamic variables of an application, we must guarantee that all the individuals represent real and feasible solutions to the problem and ensure that the search space is covered in a continuous and optimal way [6].

Table 3: Example of an individual

d_1	d_2	d_3	$\ldots$	d_m	$d_j \in \{D\}$
v_1	v_2	v_3	$\ldots$	v_m	$v_j \in \{V\}$

Table 3 shows the representation of a chromosome. Genes are represented in the first row (gray shaded cells). Each of the chromosomes represents the set of DDT that should be used to instantiate all the corresponding variables in the application from Table 1. For example, the second variable $v_2 \in \{V\}$ will be instantiated by $d_2 \in \{D\}$. A chromosome contains m genes, where m is the number of the variables logged in the application, $m = size(\{V\})$. We may use an integer to represent the values of a gene, and the constraint a gene must satisfy is:

$$1 \leq d_j \leq size(\{D\}) \tag{9}$$

Consequently, if an application contains m variables, each individual (chromosome) has to be constituted by m integer fields (i.e., m genes). Our current implementation of the exploration framework optimizes up to 3128 variables using variations of the 10 possible DDTs contained in Table 1 for each of them; Thus, it can cover large real-life dynamic embedded applications.

To compare the performance of three algorithms, all parameters are set equally. After different tests, we have fixed them to the values indicated in Table 4. The number of non-dominated solutions to preserve between generations is set to the initial population size in the elitism algorithms.

Table 4: Parameters for all three algorithms.

Parameter	VDrift	Physics
Population size	100	200
Number of generations	2000	4000
Probability of crossover	0.80	0.80
Probability of mutation	0.01	0.01

6. EXPERIMENTAL RESULTS

We have explored DDTs for VDrift and Physics with each of the three algorithms proposed (VEGA, SPEA2 and NSGA-II). The coverage, spread and the spacing values are calculated by averaging results of 10 trials.

Figure 5 depicts the Pareto-front (non-dominated individuals) obtained in the best population of both applications. With respect to VDrift, this figure shows that both NSGA-II

and SPEA2 offer the same number of non-dominated individuals and 61% more than VEGA. In the case of Physics, NSGA-II and SPEA2 offer the same optimal solutions. In addition, both NSGA-II and SPEA2 offer 92% more optimal solutions than VEGA. Thus, elitism algorithms cover better the set of possible optimal solutions than non-elitism ones.

Table 5: Coverage metric for VDrift/Physics.

VDrift	VEGA	SPEA2	NSGA-II
VEGA	–	0.005	0.034
SPEA2	0.023	–	0.215
NSGA-II	0.132	0.369	–
Physics	VEGA	SPEA2	NSGA-II
VEGA	–	0.000	0.000
SPEA2	1.000	–	0.002
NSGA-II	1.000	0.462	–

Regarding convergence comparisons, Table 5 shows that the coverage values of NSGA-II are better than VEGA or SPEA2 in both cases (e.g., for Physics $C(NSGA2, VEGA) > C(VEGA, NSGA2)$ is $1 > 0$, and $C(NSGA2, SPEA2) > C(SPEA2, NSGA2)$ is $0.462 > 0.002$). Thus, NSGA-II offers more optimal alternatives to the system designer for the implementation of the final embedded application.

Table 6: Spread and spacing for the three algorithms and two applications.

	VDrift		Physics	
	Spread	δ^2	*Spread*	δ^2
VEGA	1.47E-01	2.23E-03	8.32E-02	1.85E-03
SPEA2	5.42E-01	7.73E-04	2.21E-01	3.35E-04
NSGA-II	7.12E-01	1.25E-04	1.26E-01	2.95E-04
	Spacing	δ^2	*Spacing*	δ^2
VEGA	1.72E-02	9.57E-05	1.11E-02	7.58E-05
SPEA2	1.60E-02	2.47E-06	2.82E-03	2.34E-07
NSGA-II	2.18E-02	4.94E-05	1.91E-03	2.40E-07

The situation is different for spread and spacing metrics, where NSGA-II does not obtain the best performance in all cases. Table 6 shows the spread and spacing for the three algorithms and both VDrift and Physics applications. Regarding VDrift the larger spread is found by NSGA-II (52% better, on average), and the best spacing by SPEA2 (22% better), whereas in the case of Physics SPEA2 obtains the larger spread (53% better) and the lower spacing is found by NSGA-II (64% better).

It has been well established in the MOEA literature that although SPEA2 produces a better distribution compared to NSGA-II, the computational time needed to run SPEA2 is larger. We may conclude from our experiments that NSGA-II offers the larger coverage, whereas SPEA2 offers the better spread for large-scale explorations.

For comparison reasons we present Figure 6 to illustrate the optimization process that our methodology performs. In this test, the set of DDTs was successively implemented using SLL, DLL, etc. All the three objectives have been normalized. Thus, in the end, compared to the combination proposed by our framework.
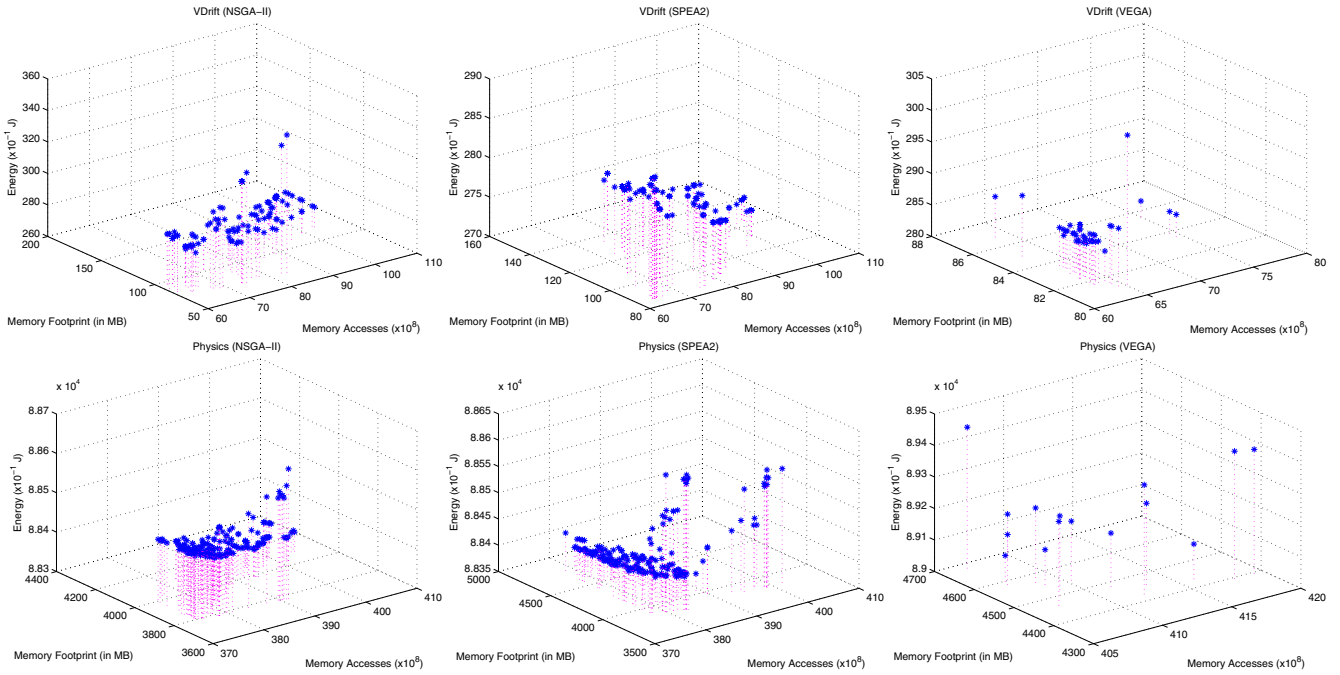
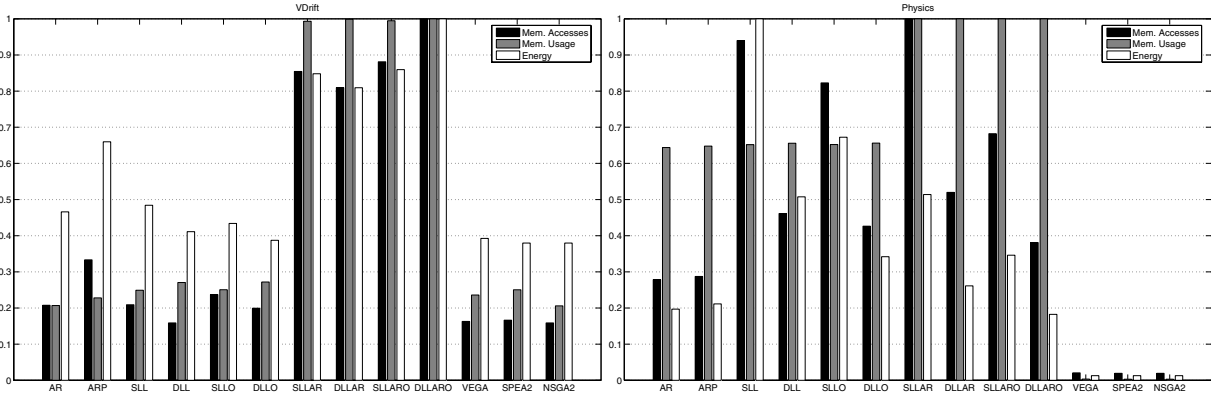Figure 5: 3D Pareto-fronts obtained for VDrift and Physics using NSGA-II, SPEA2 and VEGA.

Figure 6: Overall results for different design metrics coming from various sets of DDTs for both applications.

Figure 6 shows clearly the achieved level of optimization and final gains after applying the proposed optimization flow.

In a second set of experiments we have tested the exploration speed in comparison to different alternative methods. The results obtained for both applications for the different tested exploration methods are shown in Table 7. First, we have compared with an exhaustive exploration. Second, as Table 7 depicts, we have also compared our algorithm with state-of-the-art pruning and optimization methods for DDT implementations presented in [18][5]. In these cases several function costs, and breath-first, deep-first and branch&bound exploration heuristics are used to minimize overall memory accesses, memory usage and energy consumption in embedded multimedia applications. All experiments were executed in an AMD Sempron 3600+ 2GHz, with 1GB DDR memory.

The results in Table 7 outline that the exploration pro-

cess with our method is much faster than using directly the implementations of DDTs or other heuristics, namely, more than $12000\times$ faster (50 seconds *vs* 7 days) for VDrift and $3800\times$ (12 minutes *vs* 32 days) for Physics, due to the larger spread of solutions found in each generation in MOEAs. Note that MOEAs are faster because of the combinatorial complexity involved. There are 10^{49} and 10^{3128} feasible solutions (10 DDTs for 49 and 3128 variables) in the case of VDrift and Physics, respectively.

7. CONCLUSIONS

New multimedia embedded applications are increasingly dynamic, and rely on DDTs to store their data. The selection of optimal DDT implementations for each variable in a particular target embedded system is a very time-consuming process due to the large design space of possible DDTs implementations. In this paper we have analyzed the behavior of MOEAs for the exploration of DDT implementations in

Table 7: Time to explore DDT implementations for VDrift and Physics using different exhaustive and heuristic-based methods versus VEGA, SPEA2 and NSGA-II.

DDTs exploration method	VDrift	Physics
Breadth-First exploration	22 hours	13 days
Deep-First exploration	8 minutes	5 days
Branch&Bound exploration	62 seconds	3 hours
Other VEGA impl. [2]	3 minutes	1 hour
VEGA	50 seconds	12 minutes
NSGA-II	64 seconds	17 minutes
SPEA2	159 seconds	25 minutes

multimedia embedded applications. We have evaluated representative algorithms of elitism and non-elitism MOEAs (VEGA, NSGA-II and SPEA2) and compare them with other heuristics to optimize the DDT implementations on two real-life multimedia embedded applications. The obtained results show that VEGA is 30% quicker than the other two, but SPEA2 and NSGA-II achieve better results (75% on average) regarding covering of the design space of solutions. Moreover, MOEAs can prune the design space in a more effective way than other heuristics, since MOEAs directly search for the whole Pareto Optimal Front, and generate faster a complete range of multi-objective optimal solutions according to each concrete user-defined constraints (i.e., memory accesses, memory usage and/or power consumption). We have also automated all the optimization process, where the designer must select only the set of variables to optimize in the application.

8. ACKNOWLEDGMENTS

9. REFERENCES

[1] J. L. Antonakos and K. C. Mansfield. *Practical Data Structures using C/C++*. Prentice Hall, 1999.

[2] D. Atienza et al. Optimization of dynamic data structures in multimedia embedded systems using evolutionary computation. In *SCOPES '07: Proceedingsof the 10th international workshop on Software & compilers for embedded systems*, pages 31–40, New York, NY, USA, 2007. ACM.

[3] L. Benini and G. de Micheli. System-level power optimization: techniques and tools. *ACM Trans. Des. Autom. Electron. Syst.*, 5(2):115–192, 2000.

[4] F. Catthoor et al. *Data access and storage management for embedded programmable processors*. Kluwer Academic Publishers, 2002.

[5] E. G. Daylight et al. Memory-access-aware data structure transformations for embedded software with dynamic data accesses. *IEEE Transactions on VLSI Systems*, 12:269–280, 2004.

[6] K. Deb. *Multiobjective Optimization using Evolutionary Algorithms*. John Wiley and Son Ltd., 2001.

[7] K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.

[8] J. Edler. Dinero IV trace-driven uniprocessor cache simulator. Available at http://pages.cs.wisc.edu/~markhill/DineroIV, 2007.

[9] K. Hardee et al. A 0.6v 205MHz 19.5ns tRC 16Mb embedded DRAM. In *IEEE International Solid-State Circuits Conference (ISSCC)*, 2004.

[10] L. Kharevych and R. Khan. 3D physics engine for elastic and deformable bodies. Available at http://www.cs.umd.edu/Honors/reports/kharevych.html, 2002. University of Maryland, College Park.

[11] Z. Michalewicz. *Genetic Algorithms + data structures = Evolution Programs*. Springer-Verlag, 1996.

[12] P. R. Panda et al. Data and memory optimization techniques for embedded systems. *ACM Trans. Des. Autom. Electron. Syst.*, 6(2):149–206, 2001.

[13] S. R. Ranjithan, S. K. Chetan, and H. K. Dakshima. Constraint Method-Based Evolutionary Algorithm (CMEA) for Multiobjective Optimization. In E. Zitzler, K. Deb, L. Thiele, C. A. C. Coello, and D. Corne, editors, *First International Conference on Evolutionary Multi-Criterion Optimization*, pages 299–313. Springer-Verlag. Lecture Notes in Computer Science No. 1993, 2001.

[14] J. D. Schaffer. Multiple objective optimization with vector evaluated genetic algorithms. In *Genetic Algorithms and their Applications: Proceedings of the First International Conference on Genetic Algorithms*, pages 93–100, Hillsdale, New Jersey, 1985.

[15] J. R. Schott. *Fault Tolerant Design Using Single and Multicriteria Genetic Algorithm Optimization*. PhD thesis, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, Cambridge, Massachusetts, 1995.

[16] P. Shivakumar and N. P. Jouppi. Cacti 3.0: An integrated cache timing, power, and area model. Technical Report 2001/2, Compaq Computer Corporation, 2001.

[17] Sourceforge. Vdrift racing simulator. Available at http://sourceforge.net/projects/vdrift.

[18] S. Wuytack, F. Catthoor, and H. De Man. Transforming set data types to power optimal data structures. *IEEE Transactions on Computer-Aided Design*, 15(6):619–629, 1996.

[19] E. Zitzler, M. Laumanns, and L. Thiele. SPEA2: Improving the strength pareto evolutionary algorithm for multiobjective optimization. In *Proceedings of the Evolutionary Methods for Design, Optimization and Control with Application to Industrial Problems*, pages 95–100, Barcelona, Spain, 2002.

[20] E. Zitzler and L. Thiele. Multiobjective evolutionary algorithms: a comparative case study and the strength pareto approach. *IEEE Transactions on Evolutionary Computing*, 3(4):257–271, 1998.