

The Embedded Machine: Predictable, Portable Real-Time Code

THOMAS A. HENZINGER and CHRISTOPH M. KIRSCH
University of California, Berkeley

The Embedded Machine is a virtual machine that mediates in real time the interaction between software processes and physical processes. It separates the compilation of embedded programs into two phases. The first phase, the platform-independent compiler phase, generates E code (code executed by the Embedded Machine), which supervises the timing, not the scheduling of, application tasks relative to external events such as clock ticks and sensor interrupts. E code is portable and, given an input behavior, exhibits predictable (i.e., deterministic) timing and output behavior. The second phase, the platform-dependent compiler phase, checks the *time safety* of the E code, that is, whether platform performance (determined by the hardware) and platform utilization (determined by the scheduler of the operating system) enable its timely execution. We have used the Embedded Machine to compile and execute high-performance control applications written in Giotto, such as the flight control system of an autonomous model helicopter.

Categories and Subject Descriptors: D.4.7 [Operating Systems]: Organization and Design—*Real-time systems and embedded systems*

General Terms: Languages

Additional Key Words and Phrases: Real time, virtual machine

ACM Reference Format:

Henzinger, T. A. and Kirsch, C. M. 2007. The embedded machine: Predictable, portable real-time code. *ACM Trans. Program. Lang. Syst.* 29, 6, Article 33 (October 2007), 29 pages. DOI = 10.1145/1286821.1286824 <http://doi.acm.org/10.1145/1286821.1286824>

1. INTRODUCTION

We define a real-time execution model called the Embedded Machine (E machine, for short), which provides a portable target for the compilation of

A preliminary version of this paper appeared in the Proceedings of the 2002 ACM International Conference on Programming Language Design and Implementation, 315–326.

This research was supported in part by the AFOSR MURI grant F49620-00-1-0327 and by the NSF grants CCR-0208875 and CCR-0225610.

Authors' address: T. A. Henzinger, EPFL, Switzerland; C. M. Kirsch, University of Salzburg, Austria.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or direct commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org. © 2007 ACM 0164-0925/2007/10-ART33 \$5.00 DOI 10.1145/1286821.1286824 <http://doi.acm.org/10.1145/1286821.1286824>

ACM Transactions on Programming Languages and Systems, Vol. 29, No. 6, Article 33, Publication date: October 2007.

languages with hard real-time constraints such as Giotto [Henzinger et al. 2003]. E code (the code executed by the E machine) has strong theoretical properties, in particular, its timing and behavior are predictable. These benefits do not come at undue cost in performance. We have demonstrated this by using E code to reimplement the flight control system of an autonomous model helicopter [Kirsch et al. 2002], and to execute synthetic task sets of up to 100 tasks with a runtime overhead of less than 1% of the available CPU time [Kirsch et al. 2005].

1.1 From Platform-Centric to Requirements-Centric Real-Time Programming

In embedded systems, there are two time lines. The interaction of software processes with physical processes (sensors, actuators, clocks) happens in *environment time*. Application requirements are specified in environment time, for example, the actuator is set within 2 clock ticks of a sensor interrupt. On the other hand, the execution of software processes on a specific platform happens in *platform time*. By *platform* we mean the combination of hardware and real-time operating systems (RTOSs). Issues of platform performance, such as worst-case execution times (WCETs), and platform utilization, such as distribution and scheduling, must be addressed in terms of platform time. The art of embedded programming is to reconcile the two time lines.

The E machine proposes a paradigm shift in real-time programming: it permits the programmer to think exclusively in terms of environment time (reactivity), and shifts the burden of reconciliation with platform time to the compiler (schedulability). This paradigm shift is in line with the steady move towards higher-level programming abstractions. In fact, the E machine treats platform time as a resource in the way in which most high-level languages treat memory: the programmer assumes there is enough of it; the compiler makes sure there is enough of it (or fails to compile); the runtime system throws an exception in case the compiler is wrong (usually due to incorrect assumptions about the platform and possible contingencies).

Programming in terms of environment time avoids the two central drawbacks of conventional embedded code, nonportability and unpredictability. The former is immediate: conventional embedded code is intrinsically platform-dependent because it directly or indirectly (say, through priorities) refers to platform time; E code is platform-independent because it refers only to environment time. Predictability, both in timing and functionality, is a less immediate but powerful consequence of programming in terms of environment time. By *task* we mean a software process (or a fragment thereof) without internal synchronization points [Wirth 1996]. Suppose that the inputs of a task t become available at time x (e.g., through a sensor interrupt), and its outputs are required at time y (e.g., as an actuator setting). The programmer and, in turn, the E machine is concerned only with these two times. At environment time x , the inputs are provided to t and the task is turned over to the platform, namely, the RTOS; at environment time y , the outputs of t are read and given to the actuator. The programmer may assume that the task has indeed completed at time y , otherwise the compiler (or, as a last resort, the runtime system) will

complain. However, the programmer cannot know exactly when in the interval $[x, y]$ the task completes. In fact, she cannot even read the outputs of t as soon as they become available as this would introduce an instant of platform time into the program. The strict adherence to environment time allows us to design E code without race conditions. For two concurrent tasks, it does not matter which task completes first, as long as each task completes before its outputs are read.

A computation of the E machine is *time-safe* if each task completes before its outputs are read. Time safety depends, of course, on the platform (performance, distribution, scheduling). A good compiler ensures time safety and, in addition, the runtime system monitors time safety. For E code that refers only to environment time, time safety implies environment determinacy. A computation of the E machine is *environment-determined* if the inputs from the environment processes (e.g., the sensor readings) uniquely determine the outputs of the software processes (e.g., the actuator settings). While time safety ensures timing predictability (the actuators are written at predictable times), environment determinacy also ensures value predictability (the actuators are given predictable values).

Based on these principles, we defined the language Giotto for high-performance control applications [Henzinger et al. 2003]. Giotto supports high-level structuring principles for modern control systems such as periodic task invocation and multiple control modes. In compiling Giotto, we have found it useful to have an intermediate language with none of the high-level concepts but the same platform-independent semantics for mediating between the physical environment (typically, sensors and actuators) and software tasks (typically, control law computations). This intermediate language, which has evolved into E code, offers several benefits. First, it separates the platform-independent from the platform-dependent parts of the Giotto compiler, thus enabling reuse. The platform-independent part of the compiler generates E code from a Giotto program; its main purpose is to specify the timing of all interactions among software tasks and between software tasks and the environment. The platform-dependent part of the compiler checks the time safety of the E code for a given platform with known WCETs and known scheduling scheme (a more ambitious compiler may attempt to synthesize a scheduling scheme that guarantees time safety). Second, E code permits the dynamic implementation of Giotto. Code can be patched at runtime, and, whenever the controller switches mode, new code can be linked at runtime.

While E code has evolved from compiling Giotto, we have found it of considerable independent interest as it illustrates the causalities between the underlying semantic principles and ways to generalize them. One limitation of Giotto for some applications is its time-triggered nature: all Giotto time instants are ticks of an external clock, which, in high-performance control applications, minimizes jitter. By contrast, E code may refer to environment events that are not clock ticks such as sensor interrupts. Also, Giotto cannot refer to the completion time of tasks. For E code, it is a small and natural step to consider inputs not only from the environment processes but inputs from both the environment and the software processes. Then the completion of a software task becomes an

input event which may influence the execution of E code. As this generalization introduces E code references to platform time, environment determinacy is sacrificed and replaced by the weaker but symmetric notion of *input determinacy*: the inputs from the environment processes (e.g., the sensor readings) and from the software processes (e.g., the scheduling scheme) together uniquely determine the outputs of the software processes (e.g., the actuator settings). It is the study of concepts such as time safety and environment/input determinacy that elevates the E machine from an intermediate language for compiling embedded code to a framework for evaluating embedded programming paradigms.

Let us summarize how the programming of control applications using Giotto and the implementation using E code differs from conventional real-time software design. All executions of E code happen at predefined instants of environment time as specified by a control model and, therefore, a Giotto program; for example, a sensor is read every 10ms, the result of a control task is written to an actuator port every 20ms, etc. The compiler, by checking time safety, ensures that the program can be executed, that is, the compiler matches environment time against platform time. If time safety holds, then deterministic timing and output behavior is guaranteed. Otherwise, platform performance (WCETs) or platform utilization (scheduling) must be improved. Recompile supports code reuse on upgraded and different platforms. Conventional real-time software design proceeds in the opposite direction, that is, the programmer's model is the platform (e.g., priority-preemptive scheduling). For example, the actuator ports are typically written whenever the control task completes which is an instant of platform time. Then, code validation is necessary to gain confidence that the application requirements are met (e.g., that the actuator port is updated at least every 20ms) and that the output jitter is acceptable. Code validation, however, is usually difficult, first, because the code exhibits nondeterministic timing and output behavior and, second, because the application requirements are, unlike time safety, nonuniform. If the application requirements are not satisfied, platform performance or utilization needs to be improved. So programming the platform does not necessarily guarantee a better success rate, but it does make platform upgrades and code reuse cumbersome.

1.2 An Overview of the E Machine

The E machine is a mediator between physical processes and software processes. It interpretes E code, which supervises the execution of software processes (written in, e.g., C) in relation to physical events such as clock ticks, and physical states such as sensor values. The E machine has two input interfaces and one output interface.

- Environment inputs*. The physical processes communicate information to the E machine through environment ports such as clocks and sensors.
- Software inputs*. The application software processes, called *tasks*, communicate information to the E machine through task ports.
- Outputs*. The E machine communicates information to the physical processes and to the tasks by calling system processes, called *drivers*, which write to driver ports.

Logically, the E machine does not need to distinguish between environment and task ports since they are both input ports, while driver ports are output ports. A change of value at an input port is called an *input event*. Every input event causes an interrupt that is observed by the E machine and may initiate the execution of E code. E code, in turn, supervises the execution of both tasks and drivers.

—*Tasks*. A task is a piece of application-level code which typically implements a computation activity. When invoked with arguments, a task computes and writes the results to task ports. The execution of a task requires a positive amount of real time, that is, the results cannot be observed until at least one input event happens (e.g., a clock tick or a signal that indicates the completion of the task). A task can be preempted but has no internal synchronization points.

—*Drivers*. A driver is a piece of system-level code which typically facilitates a communication activity. A driver can provide sensor readings as arguments to a task or load task results into actuators or provide task results as arguments to other tasks. The execution of a driver satisfies the *synchrony assumption* [Halbwachs 1993] that it can be performed in logical zero time, that is, before the next input event can be observed. In other words, interrupts that implement input events are disabled during the execution of a driver.

To protect IP, both tasks and drivers may be given as binary executables. E code refers to tasks and drivers only through symbolic references. E code is interpreted system-level code that supervises the execution of a given set of tasks and drivers relative to input events. For this purpose, E code has essentially three instructions.

—*Call driver*. The *call* instruction initiates the execution of a driver. As the driver is synchronous system-level code, the E machine waits until the driver is finished before interpreting the next instruction of E code.

—*Release task*. The *release* instruction hands a task to the operating system. Typically, the task is put into a *ready queue* from which the scheduler of the operating system chooses tasks for execution according to some scheduling scheme. The scheduler is not under control of the E machine. Like the physical environment and the underlying hardware, it is external to the E machine and may or may not be able to satisfy the real-time constraints of E code. Runtime real-time violations are due to a combination of fast physical environment, slow hardware, and inefficient scheduling; they cannot be blamed on any single one of these factors. However, a compiler that checks time safety of the generated E code attempts to rule out such timing violations (and those that remain due to inaccurate assumptions about environment, hardware, and scheduler are handled by runtime exceptions).

—*Future E code*. The *future* instruction marks a block of E code for execution at some future time. It has two parameters: a trigger, which is a predicate that is evaluated with every input event, and the address of a block of E code, which is executed as soon as the trigger evaluates to true. In order to handle

multiple active triggers, the future instruction puts the trigger-address pair into a *trigger queue*. With each input event, all triggers in the trigger queue are evaluated, and the first one to evaluate to true determines the next actions of the E machine.

The E machine is a virtual machine. In an actual implementation of the E machine, E code need not be interpreted but may be compiled into, for instance, C code, or even silicon. The difference between E code and equivalent C code lies in the programming discipline imposed by E code. In particular, the fact that E code relates to time strictly through the trigger queue makes time-safety analysis possible. Moreover, the overhead incurred by E code, rather than optimized C code, is minimal. For code that supervises the timing and interaction of tasks, correctness (i.e., predictability) dwarfs performance as the critical design issue even in high-performance control applications. For example, we have found that in helicopter control, the entire program contains less than 400 instructions of E code.

1.3 A Simple Example with Two Periodic Tasks

We present a highly simplified version of the control program for a model helicopter built at ETH Zürich [Sanvido 1999]. Consider the helicopter in hover mode m . There are two tasks, both given in native code, possibly autogenerated from Matlab/Simulink models, the control task t_1 , and the navigation task t_2 . The navigation task processes GPS input every 10ms and provides the processed data to the control task. The control task reads additional sensor data (not modeled here), computes a control law, and writes the result to actuators (reduced here to a single port). The control task is executed every 20ms. The data communication requires three drivers: a sensor driver d_s , which provides the GPS data to the navigation task; a connection driver d_i , which provides the result of the navigation task to the control task; and an actuator driver d_a , which loads the result of the control task into the actuator. The drivers may process the data in simple ways (such as type conversion) as long as their WCETs are negligible. In general, since E code execution is synchronous and thus cannot be interrupted by other E code, we say that the WCET of an E code block (i.e., the sum of the WCETs of all drivers as well as all E code instructions called in that block) is negligible if it is shorter than the minimal time between any two events that can trigger the execution of E code. In the case of the helicopter software, the WCETs of all E code blocks are at least one order of magnitude shorter than 10ms which is the time between two consecutive invocations of E code in this example.

There are two environment ports, namely, a clock p_c and the GPS sensor p_s ; two task ports, one for the result of each task; and three driver ports—the destinations of the three drivers—including the actuator p_a . Figure 1 shows the topology of the program. We denote ports by bullets, tasks by rectangles, drivers by diamonds, and triggers by circles. Here is a Giotto description of the program timing:




Fig. 1. An example with two periodic tasks.

```

mode  $m()$  period 20 {
  actfreq 1 do  $p_a(d_a)$ ;
  taskfreq 1 do  $t_1(d_i)$ ;
  taskfreq 2 do  $t_2(d_s)$ ;
}

```

The `actfreq 1` statement causes the actuator to be updated once every 20ms; the `taskfreq 2` statement causes the navigation task to be invoked twice every 20ms; etc. Here is the E code generated by the Giotto compiler:

```

 $a_1$ :  call( $d_a$ )            $a_2$ :  call( $d_s$ )
       call( $d_s$ )           release( $t_2$ )
       call( $d_i$ )           future( $g, a_1$ )
       release( $t_1$ )
       release( $t_2$ )
       future( $g, a_2$ )

```

The E code consists of two blocks. The block at address a_1 is executed at the beginning of a period, say, at 0ms. It calls the three drivers, which provide data for the tasks and the actuator, then hands the two tasks to the scheduler, and finally activates a trigger g with address a_2 . When the block finishes, the trigger queue of the E machine contains the trigger g bound to address a_2 , and the ready queue of the scheduler contains two tasks, t_1 and t_2 . Now the E machine relinquishes control only to wake up with the next input event that causes the trigger g to evaluate to true. In the meantime, the scheduler takes over and assigns CPU time to the tasks in the ready queue according to some scheduling scheme. When a task completes, the scheduler removes it from the ready queue.

There are two kinds of input events, one for each environment port, namely, clock ticks and changes in the value of the sensor p_s . The trigger g : $p'_c = p_c + 10$ specifies that the E code at address a_2 will be executed after 10 clock ticks. Logically, the E machine wakes up at every input event to evaluate the trigger, finds it to be false until at 10ms, the trigger is true. An efficient implementation, of course, wakes up the E machine only when necessary, in this case at 10ms. The trigger g is now removed from the trigger queue, and the associated a_2 block is executed. It calls the sensor driver which updates a port read by task t_2 . There are two possible scenarios. The earlier invocation of task t_2 may already have completed and is therefore no longer in the ready queue when the a_2 block is executed. In this case, the E code proceeds to put another invocation of t_2 into

the ready queue, and to trigger the a_1 block in another 10ms, at 20ms. In this way, the entire process repeats every 20ms. The other scenario at 10ms has the earlier invocation of task t_2 still incomplete, that is, in the ready queue. In this case, the attempt by the sensor driver to overwrite a port read by t_2 causes a runtime exception called *time-safety violation*. At 20ms, when ports read by both tasks t_1 and t_2 are updated and ports written by both t_1 and t_2 are read, a time-safety violation occurs unless both tasks have completed, that is, the ready queue must be empty. In other words, an execution of the program is time-safe if the scheduler ensures the following: (1) each invocation of task t_1 at $20n$ ms, for $n \geq 0$ completes by $20n + 20$ ms; (2) each invocation of task t_2 at $20n$ ms completes by $20n + 10$ ms; and (3) each invocation of task t_2 at $20n + 10$ ms completes by $20n + 20$ ms. Therefore, a necessary requirement for time safety is $\delta_1 + 2\delta_2 \leq 20$, where δ_1 is the WCET of task t_1 , and δ_2 is the WCET of t_2 . If this requirement is satisfied, then a scheduler that gives priority to t_2 over t_1 guarantees time safety.

The E code implements the Giotto program correctly only if it is time-safe. During a time-safe execution, the navigation task is executed every 10ms, the control task every 20ms, and the dataflow follows Figure 1. Thus the Giotto compiler needs to ensure time safety when producing E code. In order to ensure this, the compiler needs to know the WCETs of all tasks and drivers (see for example, Ferdinand et al. [2001]) as well as the scheduling scheme used by the operating system. With this information, time safety for E code produced from Giotto can be checked efficiently [Henzinger et al. 2002]. However, for arbitrary E code and platforms, WCET analysis and time-safety checking may be difficult, and the programmer may have to rely on runtime exception handling (see Section 3 for more details). At the other extreme, if the compiler is given control of the system scheduler, it may synthesize a scheduling scheme that ensures time safety [Henzinger et al. 2003].

The time-safe executions of the E code example have an important property. Assuming the two tasks compute deterministic results, for given sensor values that are read at the input port p_s at times 0, 10, 20, . . . ms, the actuator values that are written at the output port p_a at these times are determined, that is, independent of the scheduling scheme. This is because each invocation of the control task t_1 at $20n$ ms operates on an argument provided by the invocation of the navigation task t_2 at $20n - 10$ ms whether or not the subsequent invocation of t_2 at $20n$ ms has completed before the control task obtains the CPU. Time safety, therefore, ensures not only deterministic output *timing*, but also deterministic output *values*; it guarantees predictable, reproducible real-time code. This is made precise in Section 3.

The helicopter may change mode, from hover to descend and, in doing so, apply a different control law. In this case, the control task t_1 needs to be replaced by another task t'_1 . In Section 4, we show how to implement different modes of operation using E code with control-flow instructions and how E code can be changed dynamically, at runtime, while still guaranteeing determinism if no time-safety violations occur. This capability enables the real-time programming of embedded devices that upload code on demand, of code that migrates between hosts, and of code patches.

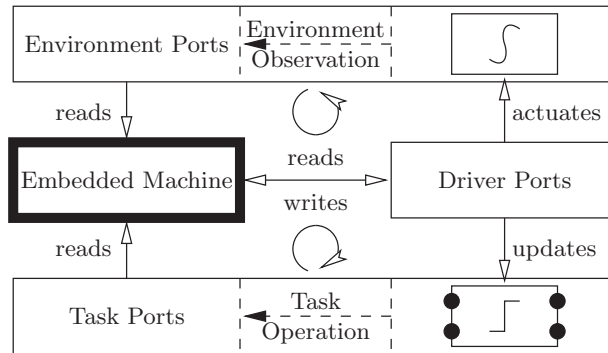


Fig. 2. Ports in the embedded machine.

2. DEFINITION OF THE E MACHINE

The E machine mediates the timing and interaction between environment and software processes. The software processes fall into three categories, drivers, tasks, and triggers. The processes communicate via *ports*. Given a set P of ports, a P state is a function that maps each port in P to a value. The set P is partitioned into three disjoint sets: a set P_E of *environment ports*, a set P_T of *task ports*, and a set P_D of *driver ports*. The read/write access of processes to ports is as follows.

	Environment Ports	Task Ports	Driver Ports
Environment	RW	–	R
Tasks	–	RW	R
Drivers	R	R	RW
Triggers	R	R	R
Input triggers	R	R	–
Environment triggers	R	–	–

The environment, task, and driver ports are updated by the physical environment, by tasks, and by drivers, respectively. All information between the environment and the tasks flows through drivers so environment ports cannot be read by tasks, and task ports cannot be read by the environment. For example, a driver may read an environment port, such as a sensor or a clock, and load the value into a driver port that is read by a task; another driver may read a task port and load the value into a driver port, such as an actuator, which is read by the environment. An *event* is a change of value at a port, for instance, at a sensor p_s , which is observed by the E machine through an interrupt. Such an event interrupt can be characterized by a predicate, namely, $p'_s \neq p_s$, where p'_s refers to the current sensor reading, and p_s refers to the most recent previous sensor reading. Figure 2 shows the E machine as an interface between physical environment and software tasks.

Definition 1. A program of the embedded machine consists of (1) a set P of program ports, (2) a set of drivers, a set of tasks, and a set of triggers, and (3) a set of addresses, and for each address, a finite sequence of instructions.

The instructions, discussed in the following, call drivers, apply scheduling services of the operating system to tasks, and handle interrupts through triggers. We do not prescribe any specific control-flow instructions, but rather view part (3) the E code of a program abstractly as a set of blocks, each with an address and a finite sequence of instructions. We use a function $next(a)$ that returns for an instruction at address a the address of the next instruction. If there is no next instruction, then the function returns bottom. This convention is consistent with any control-flow instructions, structured or unstructured, whose choice is of practical importance but entirely orthogonal to the issues discussed here. The instructions of E code do not manipulate data; all data is handled by drivers, tasks, and triggers which can be implemented in an arbitrary programming language such as C. Abstractly, drivers and tasks are functions from port values to port values, and triggers are boolean functions (i.e., predicates) on port values.

Definition 2. A driver d consists of (1) a set $P[d] \subseteq P_D$ of driver ports, and a set $I[d] \subseteq (P_E \cup P_T)$ of read environment and task ports, and (2) a function $f[d]$ from the $P[d] \cup I[d]$ states to the $P[d]$ states. A task t consists of (1) a set $P[t] \subseteq P_T$ of task ports, and a set $I[t] \subseteq P_D$ of read driver ports, and (2) a function $f[t]$ from the $P[t] \cup I[t]$ states to the $P[t]$ states.

A driver computes on (i.e., reads from and writes to) driver ports and can read from environment and task ports; a task computes on task ports and can read from driver ports. Communication to and from a task, like communication to and from the environment, is only possible through drivers. There is a fundamental difference between drivers and tasks. A driver is nonpreemptable, atomic, single-threaded code, while a task is single-threaded code that is operationally preemptable but logically atomic without internal synchronization points. Logically, a driver is assumed to execute instantaneously in zero time, whereas the execution of a task takes time. Computation in zero time is called *synchronous computation*; computation that takes time is called *scheduled computation*. Operationally, synchronous computation is performed in kernel context with event interrupts disabled. The WCET of synchronous computation (i.e., drivers and E code instructions, including trigger evaluation) as well as the overhead for context switching must be included in the administrative overhead for an accurate schedulability analysis. Scheduled computation happens in user context with event interrupts enabled. In order to validate the real-time behavior of E code through schedulability analysis, it is necessary to know the WCETs of scheduled computation (i.e., tasks).

Definition 3. A trigger g consists of (1) a set $P[g] \subseteq P$ of monitored ports, and (2) a predicate $p[g]$ on pairs of $P[g]$ states, which evaluates to true or false over each pair (s, s') of $P[g]$ states s and s' . We require that $p[g]$ evaluates to false if $s = s'$. The trigger g is an *input trigger* if $P[g] \subseteq (P_E \cup P_T)$ and an *environment trigger* if $P[g] \subseteq P_E$.

In Definition 3, the state s is the state of the ports at the time instant when the trigger is *activated*. The state s' is the state of the ports at the time instant when the trigger is *evaluated*. We assume that all active triggers are evaluated at least at the rate of observed events. An active trigger that evaluates to true may cause a reaction of the E machine. As drivers are executed in logical zero time, a trigger that reads driver ports can evaluate to true at the same logical time instant at which the trigger is activated. Such triggers make possible synchronous reactive communication between drivers [Halbwachs 1993]. Synchronous chain reactions of trigger evaluation and driver execution are not possible with input triggers which read only environment and task ports, nor with environment triggers which read only environment ports. A program is *input-triggered* if all triggers are input triggers and *environment-triggered* if all triggers are environment triggers. While input-triggered programs can react to software events such as the completion of tasks, environment-triggered programs can react only to environment events. An important special case of environment-triggered programs are the *time-triggered* programs whose triggers read only an external clock. For example, every program obtained from a Giotto source is time-triggered.

A program configuration, besides the values of all ports, also tracks the active triggers and the tasks in the ready queue of the operating system. The active triggers are kept in a FIFO queue, called a *trigger queue*, according to their activation order. An active trigger stays in the trigger queue until it evaluates to true, at which point it is removed from the queue. The tasks under OS control are kept in a set, called *task set*, as the organization of the ready queue (e.g., as a priority queue) is unknown to the E machine. A task enters the task set when it is released (i.e., handed over to the OS) and leaves the task set when it completes.

Definition 4. A program configuration consists of (1) a P state s' , called *program state*; (2) a queue of *trigger bindings* (g, a, s) , called *trigger queue*, where g is a trigger, a is an address, and s is a $P[g]$ state; and (3) a set of pairs (t, s) , called *task set*, where t is a task and s is a $P[t] \cup I[t]$ state. A trigger binding (g, a, s) is *enabled* if the trigger predicate $p[g]$ evaluates to true over the pair (s, s') of $P[g]$ states. The configuration c is *input-enabling* if the trigger queue contains no enabled trigger bindings; otherwise, c is *input-disabling*.

An E machine instruction is similar to a machine code instruction. It has a unique opcode and a finite number of arguments, all of which can be represented by integers. Using integers as arguments supports the portability of E code. There are only three fundamental instructions of the E machine. The control-flow effect of these three instructions is trivial. After executing any of the three instructions, the E machine proceeds to the next instruction in the program. Similar to the execution of drivers, E code interpretation happens logically in zero time. E code interpretation is synchronous computation which takes place in kernel context.

Definition 5. An E machine *instruction* is one of the following: $\text{call}(d)$, for a driver d ; $\text{release}(t)$, for a task t ; or $\text{future}(g, a)$, for a trigger g and an address a .

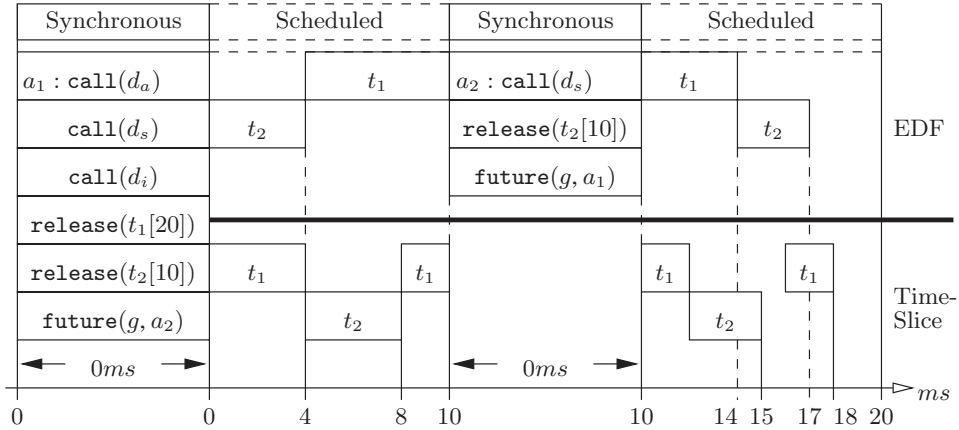


Fig. 3. Earliest-deadline-first (EDF) vs. time-slice scheduling of the tasks from Figure 1.

The $\text{call}(d)$ instruction invokes the binary code of the driver d . The E machine waits until the execution of d is finished, and then proceeds to the next instruction. The $\text{release}(t)$ instruction marks the binary code of the task t for execution by inserting it into the task set. Then the E machine immediately proceeds to the next instruction. The $\text{future}(g, a)$ instruction marks the E code at address a for (possible) execution at the future time when the trigger g next evaluates to true. Operationally, the E machine appends the trigger binding (g, a, s) to the trigger queue, where s is the current program state (this is necessary for evaluating the trigger in the future), and then proceeds to the next instruction. If there is no next instruction, then the E machine gives up control of the CPU and wakes up with each event to evaluate all triggers in the trigger queue. If a trigger binding (g, a, s) is enabled, then it is removed from the trigger queue, and the E code at address a is executed. If several trigger bindings are enabled, then the corresponding blocks of E code are executed consistent with the order of the trigger queue. In other words, the $\text{future}(g, a)$ instruction is similar to binding an interrupt handler to an interrupt where the trigger g defines the interrupt, and the E code at address a defines the interrupt handler.

2.1 The Two-Task Example Revisited

Consider Figure 3 which shows the execution of the E code and tasks from the example in Section 1. The tasks are scheduled either by an earliest-deadline-first (EDF) scheduler (upper half) or a time-slicing scheduler (lower half). The purpose of this example is to demonstrate that the E machine is independent of the scheduling scheme.

We assume that the initial configuration consists of a state s that sets the clock p_c to -10ms , a trigger queue with the single trigger binding (g, a_1, s) , and an empty task set. When the clock p_c reaches 0ms , the trigger g evaluates to true, the enabled trigger binding is removed from the trigger queue, and the E machine starts executing the E code at a_1 . The first three call instructions execute drivers. The $\text{release}(t_1[20])$ instruction releases task t_1 to the scheduler

(either the EDF or the time-slicing scheduler) by inserting it into the task set. The term [20] is an example of an *E machine annotation* which is handed to the scheduler along with the task. E machine annotations describe information that may be provided by the programmer or by the compiler of a high-level language such as Giotto to the system scheduler. For the E machine, annotations have no meaning, but the scheduler may interpret annotations. In particular, an EDF scheduler interpretes the term [20] as the relative deadline of task t_1 with respect to the clock p_c . The subsequent `release(t_2 [10])` instruction inserts task t_2 into the task set. The last instruction, `future(g, a_2)`, ensures that the E code at a_2 will be executed at 10ms. The execution of the E code at a_1 is now finished. The result is a new configuration which consists of a state s' with new values for task and driver ports, the single trigger binding (g, a_2, s') in the trigger queue, and tasks t_1 and t_2 in the task set. Finally, the E machine enables all interrupts and leaves the kernel context.

Now the scheduler, which is independent of the E machine, takes over. The upper half of Figure 3 shows the resulting timeline for the EDF scheduler which invokes task t_2 first at the 0ms instant because t_2 has in this case the earliest relative deadline of 10ms, and therefore a higher priority than task t_1 . Task t_2 executes in user context and completes after, say, 4ms, when it is removed from the task set. Then task t_1 executes. If task t_1 executes, for example, 10ms, it is preempted after 6ms by the trigger g . At 10ms, the E machine disables event interrupts and executes the E code at a_2 by first calling driver d_s to provide a new value from the sensor port p_s to task t_2 . Calling the drivers d_a or d_i at this time would result in a *time-safety violation* because both drivers access ports that are read or written by task t_1 which has not yet completed (see Section 3). The `release(t_2 [10])` instruction inserts task t_2 into the task set. Since the relative deadlines of both tasks in the task set are now equal, an EDF scheduler may assign a priority to t_2 which is lower or higher than the priority of t_1 . In the figure, we assume the priority of t_2 to be lower than the priority of t_1 . The final `future(g, a_1)` instruction marks the E code at a_1 for execution at 20ms. The new E machine configuration has the single trigger binding (g, a_1, s'') , where s'' is the new state, and tasks t_1 and t_2 in the task set. Now the scheduler resumes the execution of task t_1 because it has a higher priority than t_2 . Task t_1 executes for another 4ms and completes. Then task t_2 executes, for instance, this time only for 3ms. Then the system is idle for the next 3ms. At 20ms, the E machine executes again the E code at a_1 , thus closing an infinite loop of periodic invocations of both tasks.

Now consider the timeline of the scheduled computation in the lower half of Figure 3. This shows the task execution according to a time-slicing scheduler instead of an EDF scheduler. The E code is executed exactly the same way as before. The only difference is how the tasks are executed between the synchronous blocks of E code. The time-slicing scheduler ignores the annotations. Each task gets a time slice of 4ms, and tasks are executed in the order of their insertion into the task set in a round-robin fashion. Note that the second slice of task t_1 is preempted by a trigger. Also note that for given sensor readings at 0ms and 10ms, the value loaded by the actuator driver into p_a at 20ms is the same no matter which scheduler is used; that is, the output

behavior of the program is *deterministic*. This is the main property of E code (see Section 3).

2.2 Operational Semantics

We define the semantics of E code operationally using a pseudocode description of the E machine. Algorithm 1 shows the main loop of the machine as it executes a given program. Initially, the trigger queue contains a single trigger binding, and the task set is empty. After entering the main loop, the machine waits for environment and task ports to change their values, that is, for the occurrence of one or more events that may enable input triggers. While the machine waits for input events, scheduled computation may be performed, that is, the scheduler has control of the CPU. Waiting for input events is implemented using event interrupts. The occurrence of an input event wakes up the machine which immediately disables all event interrupts (thus it is still possible for low-level interrupts to preempt the machine as long as they do not interfere with the triggering mechanism of the machine). After the while loop, before the machine loops back to wait for new input events, all event interrupts are enabled again and the scheduler is invoked. The main loop is executed ad infinitum.

The E machine runs through the while loop of Algorithm 1 as long as there are enabled trigger bindings in the trigger queue, each time executing a block of E code that is bound to an enabled trigger. The termination of the while loop is guaranteed for input-triggered programs such as E code generated from Giotto. In an input-triggered program, newly activated triggers are initially disabled and can only be enabled by environment or task activity. In more general programs with arbitrary triggers, it is possible that calling a driver enables a trigger. In this case an explicit termination proof for the while loop is necessary. Synchronous triggering of one E code block by another corresponds to the signaling mechanism of Esterel [Berry 2000] or undelayed data dependency in Lustre [Halbwachs et al. 1991]. For synchronous reactive languages such as Esterel and Lustre, an explicit termination proof of synchronous computation is necessary (typically by ensuring the existence of finite fixed points).

Algorithm 1. The Embedded Machine

```

loop
  wait for input event
  disable trigger-related interrupts
  while there is an enabled trigger in TriggerQueue do
    (g, a, s) := GetFirstEnabledTriggerBinding(TriggerQueue)
    TriggerQueue := RemoveFirstEnabledTriggerBinding(TriggerQueue)
    ProgramCounter := a
    invoke Algorithm 2
  end while
  enable trigger-related interrupts
  invoke system scheduler on TaskSet
end loop

```

Algorithm 2. The E Code Interpreter

```

while ProgramCounter  $\neq \perp$ 
   $i := \text{GetInstruction}(\text{ProgramCounter})$ 
  if call( $d$ ) =  $i$  then
    ProgramState( $P[d]$ ) :=  $f[d](\text{ProgramState}(P[d] \cup I[d]))$ 
  else if release( $t$ ) =  $i$ 
    TaskSet := TaskSet  $\cup \{(t, \text{ProgramState}(P[t] \cup I[t]))\}$ 
  else if future( $g, a$ ) =  $i$ 
    TriggerQueue := TriggerQueue  $\circ (g, a, \text{ProgramState}(P[g]))$ 
  end if
  ProgramCounter := Next(ProgramCounter)
end while

```

Each block of E code is interpreted by Algorithm 2. In the while loop of Algorithm 2, the machine fetches the current instruction from the program, decodes and executes the instruction, and then determines the address of the next instruction. This loop terminates because every block of E code is required to have only a finite number of instructions with sequential control flow. This is true despite the fact that it is often convenient to use control-flow instructions such as absolute or conditional jumps in E code, for example, in the generation of E code from Giotto programs with mode switching (see Section 4). Later, we will also introduce additional instructions for manipulating the task set (such as terminating a task) and trigger queue (such as removing a trigger binding), but these have no fundamental impact on the operation of the E machine.

The execution of an E machine program yields a trace. From one configuration of a trace to the next, there is either an environment event or a software event (i.e., the completion of a task) or the E machine executes one block of E code. The third possibility—E code execution—has precedence as long as there are blocks of E code with enabled triggers.

Definition 6. A *program trace* is an infinite sequence of configurations such that, for all consecutive configurations c and c' , one of the following three conditions holds:

- Environment event.* The configuration c is input-enabling, and c' differs from c in the values of environment ports.
- Task completion.* The configuration c is input-enabling, and c' differs from c in the values of the task ports $P[t]$ for some task t and in the task set. The task set of c contains a pair (t, s) for some $P[t] \cup I[t]$ state s ; the task set of c' results from c by removing the pair (t, s) ; and the values of the task ports $P[t]$ in c' result from applying the function $f[t]$ to state s . In this case, we say that the task t is *completed* at configuration c .
- E code execution.* The configuration c is input-disabling, and c' differs from c in the values of driver ports, in the trigger queue, and in the task set. If (g, a, s) is the first enabled trigger binding in the trigger queue of c , then c' is obtained from c by first removing (g, a, s) from the trigger queue, and then

executing Algorithm 2 with program counter a . In this case, we say that the E code at address a is *executed* at configuration c .

The *input part* of a program trace is the projection of the trace to values for environment and task ports, that is, the trace in which everything has been removed from the program trace except for the values for the environment and task ports; the *output part* is the projection to values for driver ports. The *environment part* is the projection to values for environment ports.

3. PROPERTIES OF E CODE

The benefits of using E code are due to its strong theoretical properties. Time safety is a condition which is satisfied if all real-time requirements are met on a platform, and determinism is a consequence which ensures that the output behavior of E code is predictable.

3.1 Time Safety

A compiler that generates E code is not satisfied by every trace that can result from interpreting the code using Algorithm 2. Rather the compiler expects sufficient performance from the platform such that the computation of a task always completes before drivers access (read or write) ports that are also accessed by the task, and before another invocation of the task is released. A trace that satisfies these conditions is called time-safe because it meets all timing requirements of the source (e.g., Giotto) program.

Definition 7. A program trace is *time-safe* if for every configuration c , if the task set of c contains the pair (t, s) , then E code that is executed at configuration c must obey the following three conditions. (1) For each `call(d)` instruction, $P[d] \cap I[t] = \emptyset$, that is, no driver updates the read driver ports of t . (2) For each `call(d)` instruction, $I[d] \cap P[t] = \emptyset$, that is, no driver reads the task ports of t . (3) For each `release(t')` instruction, $P[t'] \cap P[t] = \emptyset$, that is, no released task accesses the task ports of t .

The E machine throws a runtime exception if any of the previous conditions are violated; we will discuss the exception handling in the following. In order to avoid runtime exceptions, it must be shown that the program is time-safe for the environment and the platform, that is, that all program traces that can occur on the target platform in the target environment are time-safe. Proving time safety requires a schedulability analysis based on the WCETs of all drivers, tasks, and triggers, and, if the program is not time-triggered, also it requires assumptions about the frequency of input events.

As an example, recall the Giotto program from Section 1 which, like all Giotto programs, is time-triggered. For single-CPU platforms with EDF scheduling, the Giotto compiler shows time safety in two steps. First, given the WCETs of all drivers and triggers, the compiler computes the WCETs of all E code blocks. Suppose that the a_1 and a_2 blocks have a WCET of 1ms each, including the overhead for context switching into the E machine. This leaves 18ms of CPU time per 20ms real time for scheduled computation. Second, given the WCETs of all

tasks and having derived the relative deadlines for all task invocations from the Giotto source (relative deadline 20 for each `release(t_1)` instruction, and relative deadline 10 for each `release(t_2)`), the compiler uses an EDF schedulability test to show that all task invocations complete on time. For instance, assuming a WCET of 10ms for t_1 and of 4ms for t_2 , including the overhead for context switching two times from t_2 to t_1 , the EDF schedule is feasible and achieves a theoretical CPU utilization of 100%. Note that, since tasks released by E code can be scheduled according to any scheduling strategy, we may also use rate-monotonic scheduling for the program just presented. In this case, we would achieve the same theoretical CPU utilization of 100% because the task set is harmonic [Liu and Layland 1973]. For E code generated from arbitrary Giotto programs, however, the maximum utilization with rate-monotonic scheduling may only reach 69% [Liu and Layland 1973]. On the other hand, a rate-monotonic scheduler causes less administrative overhead than an EDF scheduler, is easier to implement, and is therefore often preferred in safety-critical applications.

Time-safety checking relies on WCETs. The WCET analysis is itself a difficult problem, but E code somewhat simplifies the analysis of a system by separating functional code such as tasks from kernel code such as drivers and E code instructions. In particular, tasks do not contain any system calls, and thus have a context-free execution time except for the overhead due to context switching. The execution time of each E code instruction is constant, that is, independent of the number of triggers, drivers, tasks, and E code instructions of the program. For input-triggered programs without control-flow instructions (such as for loops), the trigger queue can be evaluated in time linear in the number of future instructions of the program. If the E code is produced from Giotto according to Section 4, then the trigger queue can be evaluated in constant time because the queue never contains more than one trigger at any time during program execution. The overhead through E code interpretation is therefore predictable provided the WCETs of the drivers are known. The WCET analysis of drivers that convert and copy data is easy, whereas device drivers that actually access hardware are difficult to analyze. This, however, is a challenge that is hard to avoid. The WCET analysis of a real-time program using more traditional kernel primitives such as semaphores or signals is more difficult than that of an E code program because of the context-sensitive and thus unpredictable demand for kernel services.

As WCET assumptions may be wrong despite a thorough analysis, the Giotto compiler (see Section 4) generates E code for handling time-safety violations. Algorithm 3 shows the E code interpreter of Algorithm 2 enhanced with exception handling. For exception handling, a second argument is added to the `release` instruction. Suppose that the task t is released by the instruction `release(t, e)`, where e is an E code address. The block of E code at address e is the *exception handler*, and its address is recorded in the task set. If, before t completes, the ports read by t are updated or the ports written by t are read by a driver or another task, then the E machine discards the instruction that causes the exception and jumps to address e . The task t is not terminated, but,

Algorithm 3. The E Code Interpreter with Exceptions

```

while ProgramCounter  $\neq \perp$  do
  i := GetInstruction(ProgramCounter); E :=  $\emptyset$ 
  if call(d) = i then
    E := {e | (t,  $\rightarrow$ , e)  $\in$  TaskSet :  $P[d] \cap I[t] \neq \emptyset \vee I[d] \cap P[t] \neq \emptyset$ }
    if E =  $\emptyset$  then
      ProgramState(P[d]) := f[d](ProgramState( $P[d] \cup I[d]$ ))
    else if release(t, e) = i then
      E := {e' | (t',  $\rightarrow$ , e')  $\in$  TaskSet :  $P[t] \cap P[t'] \neq \emptyset$ }
      if E =  $\emptyset$  then
        TaskSet := TaskSet  $\cup$  {(t, ProgramState( $P[t] \cup I[t]$ ), e)}
    else if future(g, a) = i then
      TriggerQueue := TriggerQueue  $\circ$  (g, a, ProgramState(P[g]))
    end if
    while E  $\neq \emptyset$  do
      (ProgramCounter, E) := ChooseException(E)
      invoke Algorithm 3
    end while
    ProgramCounter := Next(ProgramCounter)
  end while

```

for the case that termination is desired, we add an instruction *terminate*(*t*) to the instruction set of the E machine which removes *t* from the task set. This may be the first instruction of the exception handler. When the exception handler finishes, control flow returns to the instruction that follows the instruction that caused the exception. A single *call* or *release* instruction may cause multiple exceptions, for example, because a driver may read ports from multiple tasks in the task set. We do not specify how simultaneous exceptions are prioritized; this is done by the function *ChooseException*. As runtime exceptions can occur inside an exception handler, Algorithm 3 is invoked recursively, so as to implicitly maintain a stack of return addresses.

In the case of Giotto, the compiler generates an exception handler for each task. If a runtime exception is caused by an instruction that has a conflict with task *t* in the task set, then the associated exception handler terminates *t* and restores the most recent valid values to the task ports of *t*, that is, the values that preceded the terminated invocation of the task. Note that the values in the task ports of a terminated invocation of a task may be inconsistent because updating task ports is not necessarily an atomic operation. We restore the most recent valid values because they are easily available and typically a good approximation of the unknown current values in control applications. This approach requires additional ports and drivers. In our example, the new driver *d_j*, for *j* = 1, 2, stores the result of task *t_j* in a new driver port *p_j*. In case of a runtime exception involving task *t_j*, the exception handler at address *e_j* terminates *t_j* and calls the new driver *d'_j*, which restores the value of *p_j* to the task port of *t_j* before the program proceeds. The generated E code that follows produces exactly the same output behavior for time-safe traces as the original E code from Section 1.

a_1 : <code>call(d_1)</code> <code>call(d_2)</code> <code>call(d_a)</code> <code>call(d_s)</code> <code>call(d_i)</code> <code>release(t_1, e_1)</code> <code>release(t_2, e_2)</code> <code>future(g, a_2)</code>	a_2 : <code>call(d_2)</code> <code>call(d_s)</code> <code>release(t_2, e_2)</code> <code>future(g, a_1)</code>
e_1 : <code>terminate(t_1)</code> <code>call(d'_1)</code>	e_2 : <code>terminate(t_2)</code> <code>call(d'_2)</code>

However, if a time-safety violation occurs, then the new E code handles the associated exception. Suppose that the a_1 block is executed when task t_1 has not yet completed. The attempt to execute d_1 throws a runtime exception that invokes the E code at address e_1 . After t_1 is terminated, the driver d'_1 restores the task port of t_1 and execution proceeds with the `call(d_2)` instruction of the a_1 block.

3.2 Determinism

Figure 3 gave an example where different scheduling schemes (EDF and time slicing) lead to the same output behavior. For either scheduling scheme, the program interacts with the environment at the same constant rate of 10ms and produces the same actuator settings. We now show that this is a general property of E machine programs as long as the platform maintains time safety (i.e., there are no runtime exceptions). Recall that by *input* of the E machine, we refer to both environment and task ports (e.g., input-triggered means triggered by events on these ports); by *environment* we refer to the environment portion of the input; and by *output*, to the driver ports (which includes the actuators).

Definition 8. A program is *input-determined* if, whenever two time-safe program traces agree on the initial configurations and input parts, then they agree on the output parts as well. A program is *environment-determined* if, whenever two time-safe traces agree on the initial configurations and environment parts, then they agree on the output parts.

It is easy to see that all E machine programs are input-determined. However, input determinacy is a rather weak property. Input-determined programs are deterministic only with respect to a given behavior of all environment and software processes. In order to decouple environment and software, a program must be environment-determined, that is, independent of the real-time behavior of the software processes.

FACT. *Every program is input-determined. Every environment-triggered program is environment-determined.*

For example, since the program of Figure 3 is environment-triggered, it is environment-determined, that is, it produces the same actuator settings

independently of the execution order of the tasks as long as the platform maintains time safety. In fact, as Giotto is time-triggered, all E code generated from Giotto sources is environment-determined and so is E code generated from more general sources whose triggers monitor arbitrary environment events. On the other hand, an input-triggered program is in general not environment-determined because it may trigger on the completion of a task execution. Thus, depending on the detailed performance (WCETs) and scheduling of the platform, such a program may exhibit very different output behaviors. In other words, environment determinacy says that the software behavior depends only on the environment behavior but not on the task execution times (as long as the WCETs are met) and the task scheduling; input determinacy says that the software behavior depends on the environment behavior, the task execution times, and the task scheduling, but there are no nondeterministic choices intrinsic to the execution of the E code itself.

Environment determinacy crucially depends on the task model of the E machine. Suppose that tasks were allowed to communicate with each other directly, without going through drivers. In our example, this would mean that task t_1 reads the task port of t_2 instead of calling the connection driver d_i . Now consider the resulting timelines. In the case of EDF, the first invocation of t_1 would read the result of the invocation of t_2 that completes at 4ms. By contrast, in the case of time slicing, t_1 would start first, and thus read the initial value of the task port of t_2 . As a consequence, the result of the first invocation of t_1 , which determines the actuator setting at 20ms, would be different for the two scheduling schemes.

Environment determinacy also depends on the instruction set of the E machine. For example, instead of having a compiler ensure time safety based on schedulability analysis, we could obtain time safety trivially by using the `terminate(t)` instruction to terminate any task t before we interact with t . All traces of a program written in this way are time safe, but the program, even if environment-triggered, would not be environment-determined because the result of a terminated task depends on how long the task has executed. This technique is used in practice with *anytime algorithms* [Dean and Boddy 1988], which give a usable result when terminated at any time (of course, more CPU time usually means a better quality result). To accommodate anytime tasks by the E machine and achieve a form of environment determinacy nonetheless, the concept of time safety can be generalized: instead of checking that a task completes before its results are read, one may check that a task gets enough CPU time to obtain the desired quality of the result.

3.3 Time Liveness

Time safety means that all *finite* trace prefixes satisfy the intended real-time constraints of E code. There still may be undesirable *infinite* traces, however, where the E machine is infinitely faster than the physical environment, for instance, if the E machine executes an infinite loop within an E code block. Yet, even in the absence of infinite loops inside individual blocks of E code, two different blocks of E code may activate and enable each other infinitely

often without intervening input events. A trace without such infinite zero-time behavior is called *time-live*.

Definition 9. A program trace is *time-live* if for every configuration c , if c is input-disabling, then it is followed by some later configuration c' that is input-enabling.

FACT. Every trace of an input-triggered program is *time-live*.

Time liveness, therefore, is not an issue for E code generated from Giotto whose triggers monitor only clocks, nor for E code generated from any other source whose triggers do not look at driver ports. However, E code generation from a synchronous reactive language such as Esterel may, like any Esterel compilation, require nontrivial proofs of time liveness.

The execution of a time-live program, that is, a program whose traces are all time-live, may require a trigger queue of unbounded size which is not desirable in practice. (By contrast, during time-safe program executions, the size of the task set is always bounded by the number of tasks.) For example, two consecutive future instructions to the same block of E code may set off a process that doubles the size of the trigger queue with each input event. A similar problem may arise in the presence of control-flow instructions such as `for` loops. While always finite, the execution of a single block of E code may take more time with each new input event because the block may consist of an ever increasing number of instructions.

Definition 10. A program trace is *bounded time-live* if there exists an integer k such that for every configuration c , if c is input-disabling, then (1) c is followed by at most k input-disabling configurations, and (2) the trigger queue of c contains at most k entries.

Note that bounded time liveness implies time liveness. From Giotto programs, one can always generate E code that is bounded time-live. For E code generation from, for example, Esterel, proving bounded time liveness is necessary to ensure that the delay of executing any synchronous reaction is shorter than the time between any two events that can trigger a synchronous reaction. For a fixed bound k , bounded time liveness can be enforced at runtime by exception handling similar to the handling of time-safety violations. For example, a future instruction may cause a runtime exception if it attempts to create a new trigger binding when the trigger queue already contains k entries.

4. E CODE GENERATION AND LINKING

Generating code from standard high-level programming languages is an optimization problem to reduce the time and space requirements of the code. Generating E code is different in the sense that the time and space requirements of E code are usually negligible compared to the efficiency requirements of the task code. For example, the timing requirements of a set of 100 tasks divided into four groups with different nonharmonic frequencies can be described by around 300 E code instructions, and the execution overhead of the E machine is less than 1% of the total available CPU time [Kirsch et al. 2005]. However,

it is essential to guarantee the correct integration of synchronous and scheduled computation, that is, time safety, and to ensure bounded E code execution, that is, bounded time liveness. When generating E code from Giotto, bounded time liveness is guaranteed by construction but ensuring time safety requires explicit proof. By contrast, when generating E code from synchronous reactive languages such as Esterel or Lustre, we need only the `call` and `future` instructions but not the `release` instruction because these languages do not explicitly support scheduled computation.¹ Therefore, time safety is trivial but achieving bounded time liveness may require proof [Halbwachs 1993]. Also, for code generation from synchronous reactive languages, more flexibility in the manipulation of the trigger queue may be necessary than is offered by the presented minimal definition of the E machine. Useful additions to the instruction set of the E machine include the `cancel(a)` instruction, which removes all trigger bindings with the address *a* from the trigger queue, and the `future(true,a)` instruction, which marks the E code at address *a* for execution after all currently enabled trigger bindings have been processed.

4.1 Compiling Giotto

A Giotto program consists of a functionality part and a timing part. The functionality part contains `port`, `driver`, and `task` declarations, which interface the Giotto program to a functionality implementation typically written in C. The Giotto compiler generates so-called *functionality wrappers*—parameterless procedures—for each driver and task implementation, and stores the wrappers in a table similar to a symbol table. A wrapper calls the corresponding implementation with the proper arguments. The E code, however, refers to the wrapper using only its table index, a portable integer value. From the timing part of the Giotto program, the compiler generates annotated E code where each `release` instruction is annotated with the relative deadline of the released task.

The E code generation from multimode Giotto programs illustrates the use of conditional jumps. Consider the following timing part of a Giotto program with two modes, m_a (representing the helicopter in hover mode) and m_b (descend mode).

```

start  $m_a$  {
  mode  $m_a()$  period 20 {
    actfreq 1 do  $p_a(d_a)$ ;
    exitfreq 1 do  $m_b(c_b)$ ;
    taskfreq 1 do  $t_1(d_i)$ ;
    taskfreq 2 do  $t_2(d_s)$ ; }
  mode  $m_b()$  period 20 {
    actfreq 1 do  $p_a(d_a)$ ;
    exitfreq 1 do  $m_a(c_a)$ ;
    taskfreq 1 do  $t'_1(d_i)$ ;
    taskfreq 2 do  $t_2(d_s)$ ; }}

```

¹An exception is the `exec` statement of Esterel, which resembles scheduled computation in input-determined (but not environment-determined) E machine programs.

The program begins by executing mode m_a , which is equivalent to the (single) mode m of the Giotto program from Section 1 except for the mode switch to mode m_b . A *mode switch* in Giotto has a frequency that determines at which rate an exit condition is evaluated. The exit condition c_b of mode m_a is evaluated once every 20ms. If c_b evaluates to true, then the program switches to mode m_b which is similar to mode m_a except that task t'_1 replaces task t_1 . Task t'_1 computes a different control law on the same ports as t_1 . The mode switch back to m_a evaluates the exit condition c_a also once every 20ms. In order to express mode switching in E code, we use a conditional branch instruction $\text{if}(c, a)$. The first argument c is a *condition*, which is a predicate on some ports. The second argument a is an E code address. The $\text{if}(c, a)$ instruction evaluates the condition c synchronously (i.e., in logical zero time), similar to driver calls, and then either jumps to the E code at address a (if c evaluates to true) or proceeds to the next instruction (if c evaluates to false). The following E code implements the preceding Giotto program.

a_1 : call(d_a) call(d_s) call(d_i) release($t_2[10], e_2$) if(c_b, a'_3)	a_3 : call(d_a) call(d_s) call(d_i) release($t_2[10], e_2$) if(c_a, a'_1)
a'_1 : release($t_1[20], e_1$) future(g, a_2)	a'_3 : release($t'_1[20], e'_1$) future(g, a_4)
a_2 : call(d_s) release($t_2[10], e_2$) future(g, a_1)	a_4 : call(d_s) release($t_2[10], e_2$) future(g, a_3)

The two E code blocks in the left column implement mode m_a ; the two blocks on the right implement m_b . The exception handlers for the three tasks (e_1 , e'_1 , and e_2) are omitted. Note that, no matter which conditional branches are taken, the execution of any block terminates within a finite number of E code instructions. This concludes the first, platform-independent phase of the Giotto compiler.

The second, platform-dependent phase of the Giotto compiler performs a time-safety check for the generated E code and a given platform. For single-CPU platforms with WCET information and an EDF-based scheduling scheme and for the simple code generation strategy illustrated in the example, the time-safety check is straightforward [Henzinger et al. 2002]. For distributed platforms, complex scheduling schemes, or complex code generation strategies, this, of course, may not be the case. The code generation strategy has to find the right balance between E code and E machine annotations. An extreme choice is to generate E code that at all times maintains a singleton task set, which makes the scheduler's job trivial but E code generation difficult. The other extreme is to release tasks as early as possible with precedence annotations that allow the scheduler to order task execution correctly. This moves all control over the timing of software events from the code generator to the scheduler. In other words, the compiler faces a trade-off between static (E machine) scheduling and

dynamic (RTOS) scheduling. Our strategy, which releases tasks and computes deadlines according to the logical semantics of the Giotto source, chooses a compromise that suggests itself for control applications. To achieve controller stability and maximal performance, it is often necessary to minimize the jitter on sensor readings and actuator updates. This is accomplished by generating separate, time-triggered blocks of E code for calling drivers that interact with the physical environment. In this way, the time-sensitive parts of a program are executed separately [Wirth 1977], and for these parts, platform time is statically matched to environment time as closely as possible at the E code level. On the other hand, for the time-insensitive parts of a program, the scheduler is given maximal flexibility.

4.2 Dynamic Linking

Software modularization is an important concept in the non-real-time world for improving software reusability and reducing software complexity. Software modularization requires the use of symbolic references in executable code rather than direct references. Resolving symbolic references at compile time and runtime is called *linking* and *dynamic linking*, respectively. Since E code uses only symbolic references, E code (and tasks, drivers) can be linked statically as well as dynamically.

For an example of static linking, the two columns of E code that are generated for the two modes of the Giotto program above can be compiled separately and then linked to a complete E code executable. For dynamic linking, we leverage the dynamic nature of the trigger queue by maintaining trigger bindings to unloaded E code in the queue. This approach requires control over the queue as provided by the `cancel(a)` instruction which removes all trigger bindings with the address *a* from the trigger queue. Hence a `cancel` instruction may negate the effect of several future instructions. Consider the following program (annotations and exception handling are omitted) which implements the dynamic loading and linking of the two Giotto modes m_a and m_b during helicopter flight.

a_1 : <code>cancel(a_3)</code> <code>call(d_a)</code> <code>call(d_s)</code> <code>call(d_i)</code> <code>release(t_1)</code> <code>release(t_2)</code> <code>future(g, a_2)</code> <code>future(h_b, a_3)</code>	a_3 : <code>cancel(a_1)</code> <code>call(d_a)</code> <code>call(d_s)</code> <code>call(d_i)</code> <code>release(t'_1)</code> <code>release(t_2)</code> <code>future(g, a_4)</code> <code>future(h_a, a_1)</code>
a_2 : <code>call(d_s)</code> <code>release(t_2)</code> <code>future(g, a_1)</code>	a_4 : <code>call(d_s)</code> <code>release(t_2)</code> <code>future(g, a_3)</code>

Suppose that only the E code in the left column is currently loaded in the E machine. We begin by executing the a_1 block with an empty trigger queue

and an empty task set. Thus the initial `cancel(a_3)` instruction has no effect. The final `future(h_b, a_3)` instruction activates a new trigger h_b and binds it to the still unknown address a_3 . The trigger predicate of h_b is $(p'_c = p_c + 20) \wedge c_b$, that is, h_b becomes enabled at 20ms if the condition c_b evaluates to true (h_a is defined analogously using condition c_a). At 10ms, the E machine executes the a_2 block. The `future(g, a_1)` instruction appends the trigger g to the trigger queue after the trigger h_b , which is still in the queue. At 20ms, if c_b is not true, then h_b is not enabled and thus skipped, but g is enabled, which causes the a_1 block to be executed again. Now the first `cancel(a_3)` instruction removes the h_b trigger from the queue. On the other hand, if c_b is true, then the E machine attempts to execute instead the a_3 block. As the a_3 block is not available, the E machine starts the loader and linker to retrieve it. This overhead needs to be taken into account by time-safety analysis. Once the E code is available, the E machine begins by executing the `cancel(a_1)` instruction which removes the g trigger from the queue. The rest of the E code at a_3 and a_4 is analogous to the E code at a_1 and a_2 except that task t'_1 (whose code can also be loaded and linked dynamically) replaces task t_1 .

From the perspective of the RTOS, dynamic linking of E code means changing the interrupt and task schedules at runtime, which is standard practice. Before accepting a new interrupt or task schedule, the scheduler typically applies a schedulability test and rejects the request if the test fails. If the test succeeds, then the new interrupt and task schedules are feasible, but we have no assurances about determinism properties of the system. On the other hand, a successful schedulability test of the task set in dynamically linked E code establishes time safety, and thus provides the same strong determinism properties as static E code.

4.3 E Machine Implementations

There are several implementations of the E machine. The simplest one is written in C for Lego Mindstorms robots using the open-source LegOS operating system. The interesting aspect of this implementation is that it is a kernel patch, which makes the E machine part of the kernel, rather than the highest-priority thread outside the kernel. The E machine implementation for the ETH helicopter [Kirsch et al. 2002] is written in Oberon using a custom-designed RTOS on a StrongARM embedded processor. The interesting feature of this implementation, besides high performance, is that tasks are implemented as subroutines rather than as coroutines [Wirth 1996]. In this case, the E machine is an interrupt handler bound to a real-time clock. Preemption works through reentrant interrupts. We have also used a revised version of this implementation to benchmark the E machine on synthetic task sets of up to 100 tasks [Kirsch et al. 2005].

The third implementation of the E machine is in C under Linux using POSIX threads and semaphores. The E machine and each task runs in its own thread. Each task thread runs at a lower priority than the E machine thread and uses a unique semaphore on which it waits until the E machine signals the semaphore. For example, upon executing a `release(t)` instruction, the

E machine signals the semaphore of the thread that implements task t . When t completes, the thread loops back and waits on the semaphore for the next $\text{release}(t)$ instruction. The goal of the Linux implementation is to port it to the RTOS VxWorks, which also features POSIX threads, however, with real-time guarantees. The Linux version also includes a dynamic loader and linker for the binary format of E code. Moreover, it features a distributed E machine implementation that runs interacting E machines on each host of a distributed system. The hosts communicate using the UDP protocol on BSD sockets. The distributed version is also supported by the Giotto compiler which can generate E code separately for each E machine. The goal is to port the implementation to a network with real-time guarantees such as a time-triggered bus.

5. CONCLUSION

E code is predictable, portable, hard real-time code. E code is hard real-time because it relates to environment (physical) time, rather than platform (CPU) time; this simplifies code validation at the expense of code generation. E code is predictable because the timing and behavior of a program depends only on the external inputs; there are no internal race conditions. E code is portable because it is independent of the platform, in particular, of the scheduler. The combination of these attributes is made possible through the notion of *time safety*. E code is time-safe if its timing requirements are met on the chosen platform. In some cases, such as Giotto source programs, time safety can be guaranteed statically by the compiler [Henzinger et al. 2002]. In other cases, time-safety violations must be handled dynamically by the runtime system.

5.1 Related Work

The problem of real-time programming is to define abstractions that capture the interaction between physical processes and software processes. Two major research communities have approached this problem from different directions: the synchronous reactive language community has studied zero-delay synchronous computation [Halbwachs 1993], and the real-time systems community has focused on the theory of scheduled computation [Buttazzo 1997] and corresponding programming languages [Burns and Wellings 1997]. Both fields have had a big impact on the design of the E machine which attempts to bring together the concepts of synchronous and scheduled computation. Synchronous computation relates software processes to physical processes by modeling software processes as instantaneous reactions to physical stimuli. The main challenge is to prove the analogue of time liveness, that is, the existence of finite reactions [Berry 2000]. However, large blocks of synchronous computation naturally exhibit delayed reactivity which may deteriorate the determinism of the synchronous model. Real-time scheduling theory, on the other hand, relates software processes to physical processes by imposing a constraint system of release times and deadlines on the software processes. The main challenge is to prove the analogue of time safety, that is, the existence of a feasible schedule.

However, the arrangement of computation according to a schedule that depends on the behavior of both physical and software processes leads to an inherently nondeterministic model.

From the E machine perspective, the synchronous reactive language community deals with organizing a time-live trigger queue, and the real-time systems community worries about scheduling a time-safe task set. This can be seen more clearly from various intermediate languages that have been proposed for synchronous and for scheduled computation. Examples from the synchronous domain include the automata-based portable object code [Plaice and Saint 1998] for Esterel and Lustre, and halt point graphs [Weil et al. 2000] for Esterel. Examples from the scheduled domain include the abstract machines JAM and BEAM for code generation from the functional real-time language Erlang [Armstrong 1997]. Each of these formalisms is richer than the E machine in some respects as they permit more general manipulations of the trigger queue or of the task set, or more general control flow and data handling. The E machine attempts to exploit and at the same time restrict the possibilities in both the synchronous and the scheduled domains by identifying a set of primitives that (1) guarantee strong determinism properties of the code and (2) are sufficiently rich to be useful in practice. An important practical question raised by the E machine asks which parts of embedded software are best modeled by synchronous computation and which by scheduled computation. For example, E code generated from Giotto uses synchronous computation to implement the data transport (I/O) from and to the physical system and between software processes, while time-consuming data computation is implemented as scheduled computation. (In a distributed Giotto system, data communication across networks with nonnegligible latency is also implemented as scheduled computation.)

It should be noted that Wirth already emphasized in 1977 the value of separating logical programs from physical platforms to obtain a discipline of real-time programming [Wirth 1977]. He suggested that the time-dependent program parts, that is, the blocks of synchronous computation, are “few, simply structured, and without loops with an unknown number of repetitions” in order to maintain correctness independently from the execution time of the non-real-time code, that is, the units of scheduled computation. In particular, he proposed a “ban on the notion of interrupt” unless “interrupts can be ignored in considerations about a system’s computational state and can be confined to timing considerations only.”

5.2 Recent Work

In this article, the E machine made use of a system scheduler given by the RTOS, which is taken into account by the compiler during time-safety analysis. A more radical approach replaces the system scheduler with a compiler-generated scheduler. Such an approach is pursued in Henzinger et al. [2003] where the system scheduler is replaced by a second virtual machine called the *Scheduling Machine* (or S machine), which interprets scheduling instructions. In this setting, the compiler produces code for both the E machine, which

reacts to input events (such as environment events and task completions) by calling drivers and releasing tasks, and the S machine, which takes over between these reactions by dispatching released tasks to the CPU. This scheme affords maximum flexibility for synthesizing time-safe schedules. A distributed implementation of the E and S machines, together with a compiler for generating distributed code from Giotto programs, is presented in Henzinger et al. [2005].

ACKNOWLEDGMENTS

We thank Marco Sanvido for help with the helicopter implementation, Arkadeb Ghosal and Slobodan Matic for help with the Giotto compiler, and Edward Lee and Jie Liu for many valuable discussions.

REFERENCES

- ARMSTRONG, J. 1997. The development of Erlang. In *Proceedings of the ACM International Conference on Functional Programming*. 196–203.
- BERRY, G. 2000. The foundations of Esterel. In *Proof, Language, and Interaction: Essays in Honour of Robin Milner*, G. Plotkin, C. Stirling, and M. Tofte, Eds. MIT Press.
- BURNS, A. AND WELLINGS, A. 1997. *Real-Time Systems and Programming Languages*. Addison-Wesley.
- BUTTAZZO, G. 1997. *Hard Real-Time Computing Systems: Predictable Scheduling Algorithms and Applications*. Kluwer.
- DEAN, T. AND BODDY, M. 1988. An analysis of time-dependent planning. In *Proceedings of the National Conference on Artificial Intelligence*. 49–54.
- FERDINAND, C., HECKMANN, R., LANGENBACH, M., MARTIN, F., SCHMIDT, M., THEILING, H., THESING, S., AND WILHELM, R. 2001. Reliable and precise WCET determination for a real-life processor. In *Proceedings of the International Workshop on Embedded Software*. Lecture Notes in Computer Science, Vol. 2211, Springer, 469–485.
- HALBWACHS, N. 1993. *Synchronous Programming of Reactive Systems*. Kluwer.
- HALBWACHS, N., CASPI, P., RAYMOND, P., AND PILAUD, D. 1991. The synchronous dataflow programming language Lustre. In *Proceedings of the IEEE 79*, 1305–1320.
- HENZINGER, T., HOROWITZ, B., AND KIRSCH, C. 2003. Giotto: A time-triggered language for embedded programming. In *Proceedings of the IEEE 91*, 84–99.
- HENZINGER, T., KIRSCH, C., MAJUMDAR, R., AND MATIC, S. 2002. Time-safety checking for embedded programs. In *Proceedings of the International Workshop on Embedded Software*. Lecture Notes in Computer Science, Vol. 2491, Springer, 76–92.
- HENZINGER, T., KIRSCH, C., AND MATIC, S. 2003. Schedule-carrying code. In *Proceedings of the International Workshop on Embedded Software*. Lecture Notes in Computer Science, Vol. 2855, Springer, 241–256.
- HENZINGER, T., KIRSCH, C., AND MATIC, S. 2005. Composable code generation for distributed Giotto. In *Proceedings of the ACM Conference on Languages, Compilers, and Tools for Embedded Systems*. 21–30.
- KIRSCH, C., SANVIDO, M., AND HENZINGER, T. 2005. A programmable microkernel for real-time systems. In *Proceedings of the ACM International Conference on Virtual Execution Environments*. 35–45.
- KIRSCH, C., SANVIDO, M., HENZINGER, T., AND PREE, W. 2002. A Giotto-based helicopter control system. In *Proceedings of the International Workshop on Embedded Software*. Lecture Notes in Computer Science, Vol. 2491, Springer, 46–60.
- LIU, C. AND LAYLAND, J. 1973. Scheduling algorithms for multiprogramming in a hard-real-time environment. *J. ACM* 20, 46–61.
- PLAICE, J. AND SAINT, J.-B. 1998. The Lustre-Esterel portable format. Tech. rep. INRIA Sophia Antipolis.

- SANVIDO, M. 1999. A computer system for model helicopter flight control; Technical memo Nr. 3: The software core. Tech. rep. 317, ETH Zurich.
- WEIL, D., BERTIN, V., CLOSSE, E., POIZE, M., VENIER, P., AND PULOU, J. 2000. Efficient compilation of Esterel for real-time embedded systems. In *Proceedings of the ACM International Conference on Compilers, Architectures, and Synthesis for Embedded Systems*. 2–8.
- WIRTH, N. 1977. Toward a discipline of real-time programming. *Comm. ACM* 20, 577–583.
- WIRTH, N. 1996. Tasks versus threads: An alternative multiprocessing paradigm. *Software: Concepts Tools* 17, 6–12.

Received August 2003; revised July 2005; August 2006; accepted May 2007