

A Hierarchical Coordination Language for Interacting Real-Time Tasks*

Arkadeb Ghosal
UC Berkeley
arkadeb@eecs.berkeley.edu

Thomas A. Henzinger
EPFL
tah@epfl.ch

Daniel Iercan
"Politehnica" U. of Timisoara
daniel.iercan@aut.upt.ro

Christoph M. Kirsch
University of Salzburg
ck@cs.uni-salzburg.at

Alberto
Sangiovanni-Vincentelli
UC Berkeley
alberto@eecs.berkeley.edu

ABSTRACT

We designed and implemented a new programming language called Hierarchical Timing Language (HTL) for hard real-time systems. Critical timing constraints are specified within the language, and ensured by the compiler. Programs in HTL are extensible in two dimensions without changing their timing behavior: new program modules can be added, and individual program tasks can be refined. The mechanism supporting time invariance under parallel composition is that different program modules communicate at specified instances of time. Time invariance under refinement is achieved by conservative scheduling of the top level. HTL is a coordination language, in that individual tasks can be implemented in “foreign” languages. As a case study, we present a distributed HTL implementation of an automotive steer-by-wire controller.

Categories and Subject Descriptors:
C.3 [Special-Purpose and Application-Based Systems]:
Real-time and embedded systems.

General Terms: Languages.

Keywords: Real Time; Automotive Systems.

1. INTRODUCTION

Much current real-time programming proceeds by trial and error: if during a program test some task misses its deadline, then the task priorities are reassigned, and new tests are performed. In rare cases can the timing of a pro-

*This work was supported in part by the GSRC grant 2003-DT-660, by the NSF grant CCR-0208875, and by the Center for Hybrid and Embedded Software Systems (CHESS) at UC Berkeley, which receives support from the National Science Foundation (NSF award #CCR-0225610), the State of California Micro Program, and the following companies: Agilent, DGIST, General Motors, Hewlett Packard, Infineon, Microsoft, and Toyota.

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

EMSOFT'06, October 22–25, 2006, Seoul, Korea.
Copyright 2006 ACM 1-59593-542-8/06/0010 ...\$5.00.

gram be *proved* correct using scheduling theory and/or formal verification. Scheduling analysis becomes difficult when the program structure is irregular, with branches, exceptions, and dynamic task creation. Formal techniques are difficult due to state space explosion.

Part of the problem is that design practice deals with timing in an indirect way, often through low-level constructs such as priorities. In this paper, we present a new high-level coordination language for interacting hard real-time tasks called Hierarchical Timing Language (HTL). Like Giotto [1], our language refers directly to real-time instances, but it is more general than Giotto, in that it offers hierarchical layers of abstraction. Besides adding program structure, a main benefit of the abstraction hierarchy is that feasible schedules for lower layers can be efficiently constructed from feasible schedules for higher layers.

HTL permits composition and refinement of programs without changing their real-time behavior. Parallel program modules communicate with each other and with the environment through so-called *communicators*, of which sensors and actuators are special cases. A communicator defines a sequence of real-time instances of a static variable. Task *reads* and *writes* specify communicator instances. As the read and written time instances of communicators are fixed by a program, they remain unchanged when the context of the program is modified. In other words, the communicator instances specify a *logical execution time* (LET) [1] slot for each task; the actual physical execution of the task must fall within this slot. As long as physical task execution falls within the LET interval, the functional and timing semantics of a program is deterministic, independent of the actual task schedule. In particular, individual program modules can be reused in different contexts without changing their timing behavior, or upgraded without affecting the timing of rest of the system.

In HTL, tasks can be refined iteratively by groups of tasks with precedence relations. Each task refinement is constrained in such a way that if the task is schedulable, then the more detailed replacement group of tasks is schedulable as well. As a consequence, schedulability needs to be checked only for the top level of an HTL program. This property avoids a combinatorial explosion, and permits scheduling to be performed by the HTL compiler. The compiler rejects a program if it cannot guarantee that its timing specification is satisfied on a given platform (which is specified through worst-case execution times for all tasks).

In addition to module composition (concurrency) and task

refinement (hierarchy), HTL supports the collection of tasks into *modes*, which can be composed sequentially. We provide an operational semantics and a compiler for HTL. The compiler performs a schedulability analysis and translates HTL programs into code for the Embedded Machine (E code) [2]. The compiler can generate E code for a distributed HTL implementation; we assume that the Embedded Machine has been implemented on each host over which the HTL implementation is distributed. The semantics of an HTL program remains independent of the number of hosts, but the analysis and code generation takes into account the distribution.

To demonstrate the expressive power and the feature of HTL, we present an automotive steer-by-wire controller as a case study. A steer-by-wire system removes the mechanical linkage between steering wheel and car with a set of sensors, actuators, and a controller distributed over several processors. Typically the sensors and actuators are spread over four processors for each of the wheels; the controller implementation (along with different functionalities like fault detection, supervisory control, and power coordination) is distributed on more than one processor. The example shows the use of communicators and task refinement, while illustrating the need for horizontal and vertical extensions of the software. Horizontally, parallel modules can be appended to the implementation without changing the timing behavior of the implementation. Vertically, the refinement concept can be used to provide (temporal) space for future extensions.

2. LANGUAGE OVERVIEW

Tasks. The computational units of HTL are LET *tasks*. The LET model (Fig. 1) decouples the times when a task reads input and writes output from the time when the task executes. A LET task is sequential code with its own memory space (which cannot be accessed by other tasks) and without internal synchronization points. *Release* and *termination* events, which are triggered by clock ticks or sensor interrupts, determine the LET of the task; the task is *active* between these two events. The input of the task is written into its memory when the release event occurs, not when the task actually *starts* executing. Similarly, the output is made available to other tasks or actuators at the termination event, even if the task *completes* physical execution earlier. Between the start and completion of execution the task may be preempted and resumed any number of times. A LET task is *time-safe* on some given hardware if the task completes execution on that hardware before the termination event occurs. Time-safe LET tasks are time and value deterministic, portable and composable [2].




Figure 1: A LET (logical execution time) task

Communicators. The communication model for HTL is centered around *communicators*. A communicator is a typed variable that can be accessed (read from or written to) only at specific time instances. These time instances are periodic and specified through a communicator period. Sensors and actuators are communicators, but communicators can also be used to exchange data between tasks. A task reads from certain instances of some communicators, computes a function, and updates certain instances of the same or other

communicators. Fig. 2 shows the interaction between four communicators (c_1 , c_2 , c_3 , and c_4 , with periods 2, 3, 4, and 3, respectively) and two tasks (t_1 and t_2). Task t_1 reads the second instances of c_1 and c_4 and updates the fourth instance of c_2 . Task t_2 reads the second instance of c_3 and updates the sixth instance of c_1 and the fifth instance of c_2 . The *read* and *write* instances specify the LET for the tasks: the latest *read* instance determines the release time, and the earliest *write* instance determines the termination time. Thus, the LET of t_1 spans from time 3 to time 9, and the LET of t_2 from 4 to 10.




Figure 2: Communicators and tasks

Communicators are the key to compose HTL programs, because they exhibit deterministic behavior: given sufficient CPU speed for time-safety, the real-time behavior of a program is determined by the input (i.e., the values of all sensor instances), independent of the CPU speed and utilization. The determinism is ensured by prohibiting update races on communicators (different tasks cannot write to the same instance of a communicator), and by ensuring that every communicator instance is updated before it is read.

Modes. HTL generalizes the LET model from tasks to task groups. While tasks with different frequencies can communicate only through communicators, HTL allows direct communication between tasks with identical frequencies. A set of interacting tasks with the same frequency form a *mode* with a specified mode period. The communication flow between tasks within a mode determines an acyclic *precedence* relation on the tasks in the mode.




Figure 3: A mode with task precedences

Fig. 3 shows three tasks t_1 , t_2 , and t_3 within a mode. Task t_1 reads the second instances of c_1 and c_4 , and task t_2 reads the second instance of c_3 . No termination events are specified for t_1 and t_2 ; instead, task t_3 reads the outputs of t_1 and t_2 , and updates the fifth instance of c_2 . The communication between t_1 (and t_2) and t_3 occurs through *ports*. A port is a typed variable, but unlike a communicator, it is not bound to time instances (i.e., as soon as t_1 completes execution, t_3 can read the output port of t_1 and start execution). Thus, within the mode, the two tasks t_1 and t_2 precede the third task t_3 .

The tasks within a mode interact through ports and communicators; tasks from different modes interact only through

communicators. Fig. 4 shows two modes m_1 and m_2 with periods 6 and 12, respectively.




Figure 4: Modes

Modules. In real-time applications, a group of tasks may have to be replaced by an alternate group depending on some specific condition (e.g., a certain sensor reading). HTL accommodates this by allowing *mode switches* at the end of mode periods, which are triggered by conditions on communicator and port values. A network of modes and mode switches is called a *module* (Fig. 5). An HTL *program* is a set of modules and a set of communicators. While the modes within a module are composed sequentially (i.e., at any time, the tasks of at most one mode of a module are active), the modes from different modules are composed in parallel and may interact through communicators. Communicators thus serve several purposes: to exchange data between tasks from different parallel modules, and to exchange data from one task within a module to a later task within the same module (but possibly in a different mode). One mode in each module is specified as the *start mode*.




Figure 5: Modules

Refinement. Specifying all behaviors through mode switching is cumbersome. For a structured and concise specification, we introduce the concept of *mode refinement*: each mode in a program can be specified by an HTL program. This does not add expressiveness to the model; in fact, an HTL program with multiple levels of refinement can be translated into an equivalent “flat” program without refinement. Refinement is useful not only for compact representation, but also for conservatively simplifying the program analysis: we check schedulability only for the top-level program, and constrain refinements so that they preserve schedulability. Fig. 6 shows a mode m , and a mode m' from a program which refines m . HTL imposes certain restrictions on m' to preserve schedulability. First, the period of mode m' is identical to that of m . This ensures that when m switches (which is only possible at the end of its period), then all tasks in the modes refining m have terminated execution. Second, every task in m' refines a unique task in m (shown by dashed arrows); e.g., t'_5 (child) refines t_5 (parent). HTL considers t_5 as a placeholder (an *abstract task*) for t'_5 (the *concrete task*): the abstract task t_5 does not execute at run-time but ensures that t'_5 is accounted for during the schedulability analysis of the top-level program. Therefore, (1) the latest (resp. earliest) communicator read (resp. write) of t'_5 must be equal to or earlier (resp. later) than that of t_5 ; (2) every task that precedes t'_5 must refine a task that precedes t_5 ; and (3) the WCET of t'_5 must be less than or equal to the WCET of t_5 . These three constraints ensure that if t_5

can be scheduled in the top-level program, then t'_5 can be scheduled in the complete program.




Figure 6: Refinement

Refinement can represent both *choice* and *change* of behavior. Choice is expressed when an abstract task t in a mode m is the parent of different (concrete or abstract) tasks in several modes of a program that refines m . Change is expressed by having a (concrete or abstract) task that refines t reading from and writing to different communicators than t does (within the constraints listed above). Fig. 7 shows a diagrammatic view of an HTL program with refinement. The modes contain tasks with precedences.




Figure 7: An HTL program

Distribution. Many embedded applications are distributed: the tasks are distributed on several hosts and interact with each other through communication channels. In HTL, distribution is specified through a mapping of top-level modules to hosts (all refinements are bound to the same hosts). The distribution is implemented by replicating shared communicators on all hosts, and then have the tasks that write to shared communicators broadcast the outputs. The semantics (i.e., the real-time behavior) of an HTL program is independent of the number of hosts, but code generation and schedulability analysis must take the distribution into account. For this purpose, the LET model (Fig. 8) is extended to include both WCETs as well as worst-case output transmission times (WCTTs).




Figure 8: LET model with transmission times

3. STEER-BY-WIRE

A *steer-by-wire* (SBW) control system replaces the mechanical linkage between steering wheel and car wheels by a set of electric motors that control the wheel angle, and a controller that computes the required wheel motor actuation. To maintain a realistic road condition feel for the driver, a

force feedback actuator is placed on the steering wheel. The specific architecture that has been used here is a simplified SBW model used by General Motors for their prototype hydrogen fuel-cell car FX-3. This example is an imitation of the concerns and requirements of automotive design, and does not represent a real set of control algorithms for an actual product or prototype.

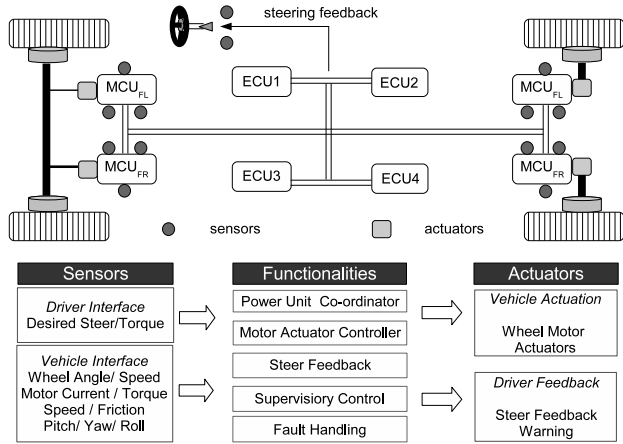


Figure 9: Steer-by-Wire: schematic view

An SBW architecture (Fig. 9) consists of eight hosts (CPUs): four motor control units (MCUs) and four electronic control units (ECUs). The MCUs are placed near the wheels; they detect sensor values related to wheels (wheel angle, motor current, speed, friction, power, pitch, and yaw) and send signals to motor actuators. The ECUs implement functionalities such as computation of wheel motor actuation, steer feedback, supervisory control, fault handling, and power coordination. Supervisory control coordinates between steering, braking, and suspension. The supervisor typically runs in triple-redundant mode (three copies are executed in three different hosts). The fault handling system detects, isolates, and mitigates faults, and warns the driver in case of a fault. The power coordinator handles the coordination of motor current computed by the controller with rest of the power grid. All hosts are connected through a communication link that allows broadcast from any host.




Figure 10: SBW modules

Steer-by-Wire in HTL. The HTL implementation of the SBW functionality consists of fifteen modules (Fig. 10¹): sensor and actuator units for each of the four wheels, and computational units for control, steer feedback, fault diagnosis, power, and supervision (three copies). Fig. 11 shows the modal structure for each of the modules, and the periods of the modes (in milliseconds). The modes capture changes in the computation policy under different environment condition; e.g., wheel actuation needs to be done faster above a critical speed. A specific scenario (Fig. 12) of inter-mode communication is discussed next; refer to <http://htl.cs.unisalzburg.at> for a full specification.

¹In the figures modules and modes are denoted by rectangles and ellipses, respectively.




Figure 11: SBW modes

The angle of the rear-left wheel (while the car is moving at a high speed) is measured by three sensors. The sensor values are read by a task that mediates on the value to be considered and then passes it on to the control task. This necessitates communication between mode **high** (module **sensor_RL**), mode **high** (module **steer feedback**), and mode **high** (module **control**). The modes and communicators have periods of 4,000 and 5,00 microseconds, respectively. Communicators instances are referred relative to the mode period (the 0-th instance corresponds to the start of the mode). Tasks **tA3**, **tA4**, and **tA5** (mode **high** of module **sensor_RL**) read three sensors **sA3**, **sA4**, and **sA5** respectively at the start of the period and update the first instances of the three communicators **cA3**, **cA4** and **cA5**. The task **MEDAngleRL** (mode **high** of module **steer feedback**) reads the first instance of the above communicators and writes to the second instance of **cAngleRL**. The task **cntrlFUN** reads the second instance of **cAngleRL** (among other communicators) and computes the value of the wheel actuation signal. The task **RackPinAct** reads the output of **cntrlFUN** and computes the power requirement for the motors.




Figure 12: Communication in SBW implementation

Fig. 13 shows the refinement of some of the modes discussed above. Refinement allows a concise representation of changes in computation policies. For example, mode **low** (module **control**) is refined by a program that has one module with two modes, **idle** and **motion**. Without hierarchical refinement, module **control** would have 6 modes and 17 mode switches; this is not only inefficient but also error prone. The modules of the root program are distributed over eight hosts: the sensor and actuator modules for each wheel share one host; the modules **control**, **steer feedback**, and **fault diagnosis** are distributed over three hosts along with one supervisor module on each host; the module **power** is assigned its own host.

Implementation. We implemented the SBW controller on eight AMD Duron 1.4GHz machines with 256MB RAM connected by a 100Mbps Ethernet network. The case study is written in 873 lines of HTL code and compiled to around 1,800 E machine instructions per host. The E machine [2] is written in C and executes the generated code with an overhead between 60 and 300 microseconds per time instant for which it is invoked. The tasks are written in C. Our implementation simulates the case study in real time but at a frequency of 2Hz, which is 1,000 times slower than the actual system.




Figure 13: SBW program with all refinements

4. ABSTRACT SYNTAX

We provide the main components of an HTL program in an abstract way. A concrete syntax [3] can be defined for this abstract syntax. An *HTL program* P consists of the following:

1 A set of *communicator declarations* commdecl . A communicator declaration $(c, \text{type}, \text{init}, \pi_c)$ consists of a communicator name c , a data type type , an initial value init , and a communicator period $\pi_c \in \mathbb{N}_{>0}$. The communicator names are unique. The set of declared communicator names for a program P is $\text{comms}(P)$. Given a communicator $c \in \text{comms}(P)$, we write $\text{type}[c]$ for the range of values that c can evaluate to, and $\text{init}[c]$ for the initial value of c .

2 A set of *module declarations* moduledcl . A module declaration $(M, \text{portdecl}, \text{taskdecl}, \text{modedcl}, \text{start})$ consists of a module name M , a set of port declarations portdecl , a set of task declarations taskdecl , a set of mode declarations modedcl , and a start mode start . The module names are unique. The set of declared module names for a program P is $\text{modules}(P)$. Given a module $M \in \text{modules}(P)$, we write $\text{start}[M]$ for the start mode.

2.1 A *port declaration* $(p, \text{type}, \text{init})$ consists of a port name p , a data type type , and an initial value init . The port names are unique for each module. The set of declared port names for a module M is $\text{ports}(M)$. Given a port $p \in \text{ports}(M)$, we write $\text{type}[p]$ for the range of values that p can evaluate to, and $\text{init}[p]$ for the initial value of p .

2.2 A *task declaration* $(t, \text{filist}, \text{folist}, \text{fn})$ consists of a task name t , a list of formal input parameters filist , a list of formal output parameters folist , and an optional task function fn . An element of the lists of formal input and output parameters is a data type type . If $v[\text{type}]$ denotes the range of values for type , the length of filist is m , and the length of folist is n , then fn is a function from $\prod_{1 \leq i \leq m} v[\text{filist}(i)]$ to $\prod_{1 \leq i \leq n} v[\text{folist}(i)]$. The function fn is absent for *abstract* tasks; if fn is present, then t is *concrete*. The task names are unique for each module. The set of declared task names for a module M is $\text{taskset}(M)$.

2.3 A *mode declaration* $(m, \pi_m, \text{invocs}, \text{switches}, \text{refp})$ consists of a mode name m , a mode period $\pi_m \in \mathbb{N}_{>0}$, a set of task invocations invocs , a set of mode switches switches , and an optional program name refp . The program refp is absent for modes that are not refined further; otherwise refp defines the *refining program*. The mode names are unique for each module. The set of declared mode names for a module M is $\text{modes}(M)$, and $\text{start}[M] \in \text{modes}(M)$.

2.3.1 A *task invocation* $(t, \text{alist}, \text{aolist}, \text{ptask})$ consists of a task name $t \in \text{taskset}(M)$, a list of actual input parameters alist , a list of actual output parameters aolist , and an optional task name ptask . An element of the lists of actual input and output parameters is either a port $p \in \text{ports}(M)$, or a pair (c, i) consisting of a communicator name $c \in \text{comms}(P)$ and a communicator instance number $i \in \mathbb{N}$, where $0 \leq i \leq \frac{\pi_m}{\pi_c}$. If P refines a mode of another

program P' , then the communicator c may be declared in P' ; see below. The task name ptask is absent if M is a top-level module; otherwise ptask defines the *parent task* and the invocation associated with the parent task defines the *parent task invocation*. The task names of the invocations are unique for each mode.

2.3.2 A *mode switch* (cnd, m') consists of a condition cnd which is expressed as a predicate on the ports in $\text{ports}(M)$ and communicators in $\text{comms}(P)$ (and communicators declared in super-programs of P ; see below), and a destination mode name $m' \in \text{modes}(M)$. The mode switches are deterministic for each mode, i.e., for port and communicator values, at most one of the mode switch conditions is true.

Hierarchical program structure. A module M_n is a *sub-module* of a module M_1 if $M_n = M_1$, or there exists n modules M_1, M_2, \dots, M_n such that for every pair M_j, M_{j+1} there exists a mode declaration $(m, \cdot, \cdot, \cdot, P) \in \text{modedcl}(M_j)$ with $M_{j+1} \in \text{modules}(P)$ for $1 \leq j < n$. If $M' \in \text{modules}(P')$, $M \in \text{modules}(P)$, and M' is a sub-module of M , then P' is a *sub-program* of P , and P is a *super-program* of P' . A program P is both a sub-program and a super-program of itself. A *top (leaf) level program* is one with no super (sub) programs other than itself. A *flat program* is one which is both a top-level and a leaf-level program. The *abstract program* $\text{abs}(P)$ for program P is the top-level program with all refining programs removed from the mode declarations of P .

Given a mode declaration $(m, \cdot, \cdot, \cdot, P)$, mode m is *parent* of mode m' if m' is a mode in some module of P ; m' is *child* of m . The set of *ancestors* of a mode m , denoted $\text{ancestors}(m)$, is the smallest set of modes that includes the parent of m and the ancestors of the parent. The *start set* of m , denoted $\text{starts}(m)$, is the smallest set that includes m and contains the start sets of all start modes of modules in P .

The *accessible communicator set* for module M in program P is the set of communicators declared by the super-programs of P . The *write-set* for M is the set of communicators occurring as actual output parameters of task invocations of modes in M . The *hierarchical write-set* for M is the set of communicators that belong to the accessible communicator set of M and to the write-sets of all sub-modules of M .

A communicator not written to by any task invocation is called an *input communicator*. Input communicators are written by the physical environment (sensors), or by other programs. *Input communicator set* $\text{icset}(P)$ for program P is the set of input communicators for all sub-programs of P .

Task precedences. A task invocation inv *precedes* another invocation inv' if inv' reads any port written by inv ; in this case, inv' *follows* inv . The *preceding invocation set* of a task invocation inv in mode m , denoted $\text{prec}(\text{inv}, m)$, is the set of task invocations that precede inv . The *read time* (resp. *write time*), $\text{rtime}(\text{inv}, m)$ (resp. $\text{wtime}(\text{inv}, m)$) of a task invocation inv in mode m is the largest (resp. smallest) interval from the start of mode m at which a communicator instance is read (resp. written) by inv or any of its preceding (resp. following) invocations.

Each task invocation has an input (resp. output) port associated with each actual input (resp. output) parameter; we will refer to these ports as *task ports*. The set of input (resp. output) task ports for task invocation inv is $\text{tips}(\text{inv})$ (resp. $\text{tops}(\text{inv})$). A task port has the same type and initial value as the communicator (or module port) corresponding to the actual parameter denoted by the task port. The input (resp. output) task port that reads a port p of a module is $\text{tip}_{\text{inv}}^p$ (resp. $\text{top}_{\text{inv}}^p$). The input (resp. output) task port that reads from (resp. writes to) i -th instance of a communicator c is $\text{tip}_{\text{inv}}^{c,i}$ (resp. $\text{top}_{\text{inv}}^{c,i}$). The set of task ports (both

input and output) for a module M is $\text{tpset}(M)$. The set of (module) ports read (resp. written) by a task invocation inv is $\text{rdset}(\text{inv})$ (resp. $\text{wrset}(\text{inv})$).

Well-formed programs. A program is *well-formed* if it conforms to the following syntactic restrictions (refer [3] for formal definitions):

Constraints on programs: (1) There is only one top-level program; and (2) each mode (other than those of the top-level program) has a unique parent mode.

Constraints on communicators: (1) if a communicator is declared in program P , then it is not redeclared in any other sub-program of P ; (2) if a communicator c is accessed (read or written) by a task invocation or a switch in a mode of module M in program P , then c is declared in one of the super-programs of P ; and (3) if a communicator c belongs to the hierarchical write-set of a module M , then c does not belong to hierarchical write-set of any sibling module of M .

Constraints on task invocations: (1) for every task invocation, the read time is earlier than the write time; (2) the precedence relation on task invocations is acyclic; (3) two task invocations in a mode cannot write to the same port or to the same instance of a communicator; (4) if a task invocation reads from or writes to a communicator c (resp. port p), then the type of c (resp. p) complies to the type of the corresponding formal parameter in task declaration; (5) if a task invocation in a mode m reads from or writes to a communicator c , then period π_m is a multiple of the communicator period π_c ; and (6) every task invocation reads from a communicator or is transitively preceded by a task invocation that reads from a communicator, and writes to a communicator or is transitively followed by a task invocation that writes to a communicator.

Constraints on refinement: (1) if program P refines a mode m , then the period of all modes in P is equal to π_m (this ensures that when there is a mode switch, there is no unsafe termination of tasks in lower-level modes); (2) every task invocation of a mode m that does not belong to the top-level program has an abstract parent task invocation in the parent of m (this ensures that the parent task acts as a placeholder for its children during schedulability analysis); (3) any two distinct task invocations in two modes of two (possibly identical) sibling modules have distinct parent task invocations (this ensures that all tasks that can potentially execute in parallel have unique parents); and (4) if inv' is the parent task invocation of inv , then the read time of inv is not later than that of inv' , the write time of inv is not earlier than that of inv' , and every invocation that precedes inv refines a task invocation that precedes inv' (this ensures that the parent invocation is more constrained in time than the child task invocation, which is used in the schedulability analysis).

Well-timed programs. The *host set* hset for a program P is the set of hosts over which the program is distributed. A *host mapping* hmap is a function that assigns each module in top-level program of P to a host $h \in \text{hset}$. For a given host set hset and host mapping hmap , the *worst-case execution time* (WCET) and the *worst-case transmission time* (WCTT) maps, wemap and wtmap , assign two natural numbers to each task of P : the WCET $\text{wemap}(t)$ bounds the execution time of task t on host $\text{hmap}(t)$; the WCTT $\text{wtmap}(t)$ bounds the broadcasting of the outputs of task t to all hosts. Given a well-formed HTL program P together with an execution platform specified in terms of hset , hmap , wemap , and wtmap , the program P is *well-timed* if for all task invocations $\text{inv}_1 = (t_1, \cdot, \cdot, \cdot)$, $\text{inv}_2 = (t_2, \cdot, \cdot, \cdot)$: if inv_1 is the parent invocation of inv_2 , then $\text{wcet}(t_1) \geq \text{wcet}(t_2)$ and $\text{wctt}(t_1) \geq \text{wctt}(t_2)$. Note that while well-formedness is independent of the execution platform, well-timedness is not.

5. OPERATIONAL SEMANTICS

The semantics (i.e., set of traces) of an HTL program is independent of the execution platform, which will be taken into account later, for code generation and schedulability analysis. The execution of an HTL program yields a (possibly infinite) sequence of configurations, called trace. Each configuration consists of values for all program variables (ports and communicators), a set of triggers, and a set of released (but not yet completed) tasks. A trigger defines an action to be taken at a future event, which is specified by an integer n and a set cmpts of tasks: the trigger becomes enabled as soon as n time ticks have passed *and* all tasks in cmpts have completed execution. In practice, time ticks and task completion event are raised by the execution platform through interrupts; the time unit is required to be a harmonic fraction of all communicator and mode periods. When a trigger becomes enabled, the associated action is carried out; this may be a communicator write (handled by write triggers), mode switch (handled by switch triggers), communicator read (handled by read triggers), or task release (handled by release triggers). Enabled write, switch, read, and release triggers are handled in this order. If no trigger is enabled, then the next time tick is awaited, and any number of released tasks may complete their execution.

Formally, a *trace* of an HTL program is a sequence of configurations u_0, u_1, \dots , where u_0 is the initial configuration, and for all $i > 0$, configuration u_i is a time event, write, switch, read, or release successor of u_{i-1} . Each *configuration* u is a triple $(\text{state}, \text{trgs}, \text{tasks})$, where state is variable state, trgs is a set of triggers, and tasks is a set of tasks. The *variable state* is a valuation of all communicators, module ports, and task ports. Without loss of generality, we assume that all communicator names and port names (both module and task ports) across the refinement hierarchy are unique; we furthermore assume that each task name uniquely identifies a particular task invocation (this can be achieved by duplication and renaming of task declarations). At a configuration u , c_u (resp. p_u) denotes the value of a communicator c (resp. port p) and cnd_u denotes the boolean value of a mode switch condition cnd .

A *trigger* g is a triple (τ, e, a) , where $\tau \in \{\text{w}, \text{s}, \text{d}, \text{r}\}$ is a tag that identifies write, switch, read, and release triggers, respectively; e is an event instance; and a is action to be carried out when the trigger is handled. An *event instance* is a pair (n, cmpts) , where $n \in \mathbb{N}_{\geq 0}$ and cmpts is a set of task names. The trigger is *enabled* if $n = 0$ and $\text{cmpts} = \emptyset$. A configuration is *waiting* if none of its triggers is enabled. For write triggers ($\tau = \text{w}$), the action $a = (c, i, t)$ specifies that the i -th instance of communicator c will be updated to the value of the corresponding output task port of task t . For switch triggers ($\tau = \text{s}$), the action $a = (\text{sw}, m)$ specifies that the condition of mode switch sw of mode m will be checked, causing a possible mode switch. For read trigger ($\tau = \text{d}$), the action $a = (t, c, i)$ specifies that the i -th instance of communicator c will be read into the corresponding input task port of task t . For release triggers ($\tau = \text{r}$), the action $a = t$ specifies that the (unique) invocation of task t will be released.

The five successor relations on configurations are defined formally in Fig. 14. Configuration u has a *time-event successor* if u is waiting; in this case, a time tick event and possibly some task completion events occur. The output task ports of the completed tasks are updated. The completed tasks are removed from the set of released tasks, and from all task sets of the triggers in u . The positive time tick counts of

| | | |
|------------------------|--|--|
| successor [†] | the following conditions hold on $u = (\text{state}, \text{trgs}, \text{tasks})$ | the following conditions hold on $u' = (\text{state}', \text{trgs}', \text{tasks}')$ |
| time event | no enabled trigger in trgs | $\text{tasks}' \subseteq \text{tasks}$, if $(\cdot, (n, \text{cmps}), \cdot) \in \text{trgs}$ then $(\cdot, (n \ominus 1, \text{cmps} \setminus (\text{tasks} \setminus \text{tasks}')), \cdot) \in \text{trgs}'$, $\forall t \in \text{tasks} \setminus \text{tasks}'$: $\forall p \in \text{tops}(t) : p_{u'} = \text{fn}[\prod_{p' \in \text{tips}(t)} p_{u'}]$, $\forall p \in \text{wrset}(t) : p_{u'} = \text{tops}_{t,u'}^p$, $\forall c \in \text{icset}(P) : c_{u'} = \vartheta(\text{type}_c)$ [★] |
| write | $\exists g = (w, (0, \emptyset), (c, i, t)) \in \text{trgs}$ | $c_{u'} = \text{top}_{t,u'}^{c,i}, \text{trgs}' = \text{trgs} \setminus \{g\}, \text{tasks}' = \text{tasks}$ |
| switch | no enabled write trigger in trgs , $\exists g = (s, (0, \emptyset), (sw, m)) \in \text{trgs}$: $(sw = (\text{cnd}, m_1) \wedge$ $\forall (s, (0, \emptyset), (\cdot, m_2)) \in \text{trgs} \setminus \{g\} : m_2 \notin \text{ancestors}(m))$ | if $\neg \text{cnd}_u$ and exists other enabled switch trigger for m in trgs : $\text{trgs}' = \text{trgs} \setminus \{g\}, \text{tasks}' = \text{tasks}$ if $\neg \text{cnd}_u$ and no other enabled switch trigger for m in trgs : $\text{trgs}' = \text{gset}_i(m) \cup \text{trgs} \setminus \{g\}, \text{tasks}' = \text{tasks}$ if cnd_u : $\text{trgs}' = \bigcup_{m_3 \in \text{starts}(m_1)} \text{gset}_i(m_3) \cup \text{trgs} \setminus \text{gset}_r(m, u)$, $\text{tasks}' = \text{tasks}$ |
| read | no enabled write or switch triggers in trgs , $\exists g = (d, (0, \emptyset), (t, c, i)) \in \text{trgs}$ | $\text{tip}_{t,u'}^{c,i} = c_{u'}, \text{trgs}' = \text{trgs} \setminus \{g\}, \text{tasks}' = \text{tasks}$ |
| release | no enabled write or switch or read triggers in trgs , $\exists g = (r, (0, \emptyset), t) \in \text{trgs}$ | $\forall p \in \text{rdset}(t) : \text{tip}_{t,u'}^p = p_{u'}$, $\text{trgs}' = \text{trgs} \setminus \{g\}, \text{tasks}' = \text{tasks} \cup \{t\}$ |

[†] Values of variables remain unchanged from u to u' unless noted.

[★] ϑ non-deterministically assigns a value from type type_c of communicator c .

$$\begin{aligned}
 n \ominus 1 &= n - 1, \text{ if } n > 0 & \text{gset}_i(m) &= \text{Procedure_Invoke_Mode}(m) \\
 &= n, \text{ otherwise} & \text{gset}_r(m, u) &= \{(s, (0, \emptyset), (\cdot, m_4)) \in \text{trgs}(u) : (m_4 = m) \vee (m \in \text{ancestors}(m_4))\}
 \end{aligned}$$

Figure 14: Definition of successor configurations

all triggers in u are reduced by one. All input communicators are non-deterministically assigned a value from their type; for simplicity we assume all input communicators have period one.

Configuration u has a *write successor* if u contains an enabled write trigger. The write trigger is removed and a communicator is updated from an output task port. Configuration u has a *switch successor* if in u no write trigger is enabled, a switch trigger (say, for mode m) is enabled, and there are no enabled switch triggers for any ancestors of m . There are three different scenarios: (1) if the switch condition evaluates to false and there exists another enabled switch trigger from m , then the switch trigger is removed; (2) the switch condition evaluates to false and there exists no other enabled switch trigger from m , then the switch trigger is removed and m is invoked; and (3) if the switch condition (say destination mode is m') evaluates to true, then all enabled switch triggers of m and of the descendants of m are removed, and all modes in $\text{starts}(m')$ are invoked. An *invocation* of mode m (see Alg. 1) involves the following steps: (1) for each task invocation in m , a read trigger is added for each communicator input; (2) for each task invocation in m , a write trigger is added for each communicator output; (3) for each task invocation, a release trigger is added; and (4) for each mode switch, a switch trigger is added.

Configuration u has a *read successor* if in u no write or switch trigger is enabled, but some read trigger is enabled. The read trigger is removed and a communicator is read into an input task port. Configuration u has a *release successor* if in u no write or switch or read trigger is enabled, but a release trigger is enabled. The release trigger is removed, the input ports of a task are loaded with output port values from preceding tasks, and the task is added to the set of

Algorithm 1 Procedure_Invoke_Mode(m)

```

gset =  $\emptyset$ ;
 $\forall \text{inv} = (t, \text{aalist}, \text{aolist}, \cdot) \in \text{invoCs}(m)$ 
   $\forall k \in \mathbb{N}$  s.t.  $\text{aalist}[k] = (c, i)$ 
    add trigger  $(d, (i \cdot \pi_c, \emptyset), (t, c, i))$  to gset;
   $\forall j \in \mathbb{N}$  s.t.  $\text{aolist}[j] = (c, i)$ 
    add trigger  $(w, (i \cdot \pi_c, \emptyset), (c, i, t))$  to gset;
  add trigger  $(r, (n, \text{cmps}), t)$  to gset for
     $n = \text{rtime}(t)$  and  $\text{cmps} = \{t' : (t', \cdot, \cdot) \in \text{prec}(t, m)\}$ ;
 $\forall \text{sw} \in \text{switches}[m]$ 
  add trigger  $(s, (\pi[m], \emptyset), (sw, m))$  to gset;
return gset.

```

released tasks.

The *initial* configuration of a program is defined as follows: the variable state assigns the initial values to all module ports, task input ports and communicators and default values (specified by the types) to task output ports; the trigger set consists of triggers by invoking modes in start set of $\text{start}(M)$ for each module M in top-level program; and an empty task set.

6. HTL COMPILER

We have designed and implemented a compiler for full, distributed HTL in Java. The compiler checks well-formedness, well-timedness, and schedulability of a given HTL program, flattens the program into a semantically equivalent HTL program with only top-level modules, and then generates so-called *E code* for the flattened program targeting the *E(mbedded) Machine* [2]. In our experiments, we have used an existing implementation of the E machine written in C running on Linux. E code specifies the exact real-time instants when port and communicator values are exchanged,

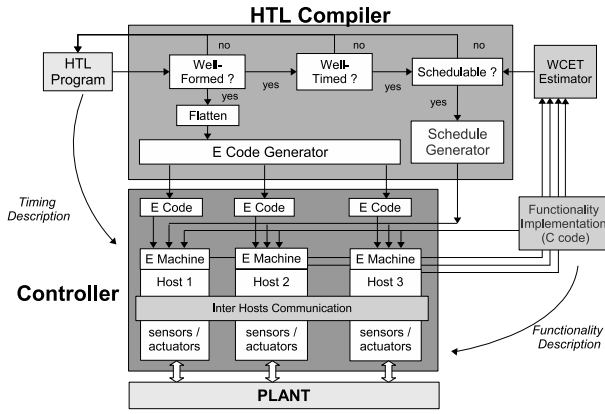


Figure 15: Structure of compiler and target

and when tasks are released and terminated. E code neither implements the actual tasks' functionality nor specifies when released tasks actually execute. Task functionality must be implemented in some other programming language and compiled separately using an appropriate compiler. Here, we have chosen C for implementing tasks, since our version of the E machine is implemented in C. Released tasks are dispatched for execution by an EDF scheduler that is external to the E machine and also implemented in C.

Fig. 15 depicts the structure of the compiler and the target architecture. Each host runs its own E machine and maintains its own copies of all task ports and communicators of an HTL program, even if some task ports and communicators are never accessed by tasks on that host. The compiler generates E code for each host separately. The idea is to compile repeatedly the whole HTL program for each host and to generate E code that implements the whole program on each host, except that the tasks of the modules not mapped to a host are not invoked on that host. Thus the generated E code is identical across all hosts except for the instructions that invoke tasks. Each task invocation involves broadcasting the task's output port values and storing the values in the respective task output ports on all receiving hosts. As a result, each host maintains a complete image of all port and communicator values of an HTL program. Note that all host-to-host transmission is done by the tasks, not by the E code. In our prototype implementation, tasks broadcast via sockets and standard Ethernet. Memory consumption as well as transmission load may be minimized if necessary using, e.g., data-flow analysis, which is, however, future work. In the following, we explain each phase of the compiler.

Checking well-formedness, well-timedness and schedulability. Before flattening the input program, the compiler checks the well-formedness and well-timedness of the program. In particular, the compiler verifies that any concrete task indeed refines its parent task to make sure that the subsequent top-level scheduling test guarantees overall schedulability. The compiler performs an EDF-scheduling test on the abstract, top-level portion of the input program only. If the test succeeds, according to our result in Section 7, the whole input program is schedulable. This result also applies to distributed HTL programs as long as the WCTT for broadcasting the output port values of each task is added to the WCET of the task, and the WCTT includes the time it takes to resolve any collisions even when all hosts try to broadcast at the same time. In our current implemen-

tation, output port values of a task are always broadcast as soon as an invocation of the task completes. Transmission and scheduling techniques that may utilize the network more efficiently, e.g. [4], can also be used but have not been implemented.

Flattening. The HTL compiler flattens a given well-formed and well-timed HTL program into a semantically equivalent flat HTL program. A module of a flat HTL program may only contain modes that do not contain a refining program. Flattening works by essentially computing the product of all modes in the refinement of each top-level module of the original program. This is easy because these modes all have the same period. Only modes in different top-level modules may have different periods. In order to maintain semantical equivalence, flattening needs to prioritize mode switch checking, i.e., mode switches in more abstract modules need to be checked before mode switches in more concrete modules [3]. Flattening an HTL program may in theory result in generated code that is exponentially larger than the size of the input program. However, execution of the generated code is very efficient and is readily supported by existing versions of the E machine. An HTL compiler that may shift the trade-off between code size and execution efficiency more towards smaller code size by generating code directly from the unflattened input program is future work. Note that for such a compiler the design of the E machine may have to be modified as well.

Target machine. The HTL compiler generates E code, which has a semantics that is designed to simplify code generation and can be executed very efficiently [5]. Besides releasing tasks, E code also controls when port and communicator values are copied (or initialized) using so-called *drivers*, which are implemented in C, one for each data type. E code consists of the following instructions: a *call(d)* instruction executes the driver *d*; a *release(t)* instruction releases the task *t* for execution by the EDF scheduler; a *future(g, a)* instruction marks the E code at the address *a* for future execution when the predicate *g* evaluates to true, i.e., when *g* is *enabled*. We call *g* a *trigger*, which observes events such as time ticks and the completion of tasks. Here, we only use triggers that are enabled when all observed events have occurred. The E machine maintains a FIFO queue of trigger-address pairs. If multiple triggers in the queue are enabled at the same instant, the corresponding E code is executed in FIFO order, i.e., in the order in which the *future* instructions that created the triggers were executed. An *if(cnd, a)* instruction branches to the E code at address *a* if predicate *cnd* evaluates to true. We call *cnd* a *condition*, which observes port or communicator states. A *jump(a)* instruction is an absolute jump to E code address *a*, and a *return* instruction completes the execution of an E code sequence.

E code generation. Given a well-formed, well-timed, schedulable, and flattened HTL program and a mapping of its top-level modules to hosts, the HTL compiler generates E code for the program and mapping by invoking Algorithm 2 for each host, which invokes Algorithm 3 to generate E code for each module of the program, which finally invokes Algorithm 4 to generate E code for each mode of each module. The compiler conceptually divides each mode into uniform temporal segments called units. The *unit* of a mode is the smallest time interval at which any two consecutive communicator instances are accessed in that mode. Given a mode *m*, we denote the duration of its unit by $\gamma[m]$, which is the gcd of all access periods of all communicators accessed in *m*. The total number of units of *m* is $\pi[m]/\gamma[m]$, where $\pi[m]$ is the period of *m*. The compiler generates sep-

arate E code blocks for each unit of a mode. The address of an E code block corresponding to unit i of a mode m is denoted by $unit_address[m, i]$. This is a symbolic address to which instructions may forward reference and therefore may need fix up during compilation. We use similar notation for other symbolic addresses.

We also use the following auxiliary operators. The driver $init(x)$ initializes the communicator or task port x . The set $readDrivers(m, u)$ contains the drivers that load the tasks in mode m with values of the communicators that are read by these tasks at unit u . The set $writeDrivers(m, u)$ contains the drivers that load the communicators with the output of the tasks in mode m that write to these communicators at unit u . The set $portDrivers(t)$ contains the drivers that load input ports of task t with the values of module ports on which t depends. The set $complete(t)$ contains the events that signal the completion of the tasks on which task t depends, and that signal the read time of the task t . The set $releasedTasks(m, u)$ contains the tasks in mode m with no precedences that are released at unit u . The set $precedenceTasks(m)$ contains the tasks in mode m with precedences.

Algorithm 2 generates instructions to initialize all communicators and modules. Here, we use instructions of the form $future(0, module_address[M])$, which effectively execute the E code at the address $module_address[M]$ similar to a jump-to-subroutine instruction. However, the actual mechanism is more complicated: for a $future(0, a)$ instruction the E machine appends the already enabled trigger-address pair $(true, a)$ to the trigger queue and then proceeds to the next instruction. Only when the E machine reaches a $return$ instruction, the machine checks the trigger queue again and eventually removes the pair $(true, a)$ from the trigger queue and executes the E code at the address a but not before it executed the E code of all other enabled trigger-address pairs occurring before $(true, a)$ in the queue.

Algorithm 2 GenerateECodeForProgramOnHost(P, h)

```
// initialize communicators
∀c ∈ comms(P): emit(call(Init(c)))
// initialize and start each module
∀M ∈ modules(P): emit(future(0, module_address[M]))
// end initialization phase
emit(return)
// generate code for each module
∀M ∈ modules(P): GenerateECodeForModuleOnHost(M, h)
```

Algorithm 3 generates instructions to initialize all task ports in a module, and to start the execution of the module by jumping to the E code of the first unit of the start mode of the module. PC denotes the compiler's program counter.

Algorithm 3 GenerateECodeForModuleOnHost(M, h)

```
set module_address[M] to PC and fix up
// initialize task ports
∀p ∈ tpset(M): emit(call(Init(p)))
// jump to the start mode at unit 0
emit(jump, unit_address[start[M], 0])
// generate code for each mode
∀m ∈ modes(M): GenerateECodeForModeOnHost(m, h)
```

Algorithm 4 generates the E code for all units of a mode. Only unit 0 contains instructions to check mode switching because mode switching may only occur at the beginning of a mode. When a mode switch occurs, E code execution continues at the $mode_address[m']$ of the target mode m' , not the $unit_address[m', 0]$, since only at most one mode switch

Algorithm 4 GenerateECodeForModeOnHost(m, h)

```
u := 0
while u < π[m]/γ[m] do
  set unit_address[m, u] to PC and fix up
  // update communicators with task output
  ∀d ∈ writeDrivers(m, u): emit(call(d))
  // continue after other modules updated communicators
  emit(future(0, PC + 2))
  emit(return)
  if (u = 0)
    // check mode switches
    ∀(cnd, m') ∈ switches(m): emit(if(cnd, mode_address[m']))
    set mode_address[m] to PC and fix up
  end if
  if (mode m is contained in a module on host h)
    // read communicators into tasks
    ∀d ∈ readDrivers(m, u): emit(call(d))
    // release tasks with no precedences
    ∀t ∈ releasedTasks(m, u): emit(release(t))
  if (u = 0)
    // release tasks with precedences
    ∀t ∈ precedenceTasks(m):
      // wait for tasks on which t depends to complete
      emit(future(complete(t), PC + 2))
      emit(jump(PC + 5 + |portDrivers(t)|))
      // release t after other modules updated communicators
      emit(future(0, PC + 2))
      emit(return)
      // read ports of tasks on which t depends, then release t
      ∀d ∈ portDrivers(t): emit(call(d))
      emit(release(t))
      emit(return)
  end if
end if
// continue mode after γ[m] time
emit(future(γ[m], unit_address[m, u + 1 mod π[m]/γ[m]]))
emit(return)
u := u + 1
end while
```

per time instant may occur. At each time instant, the generated E code uses $future(0, a)$ instructions to write communicators always before any communicator is read making sure that the latest communicator values are used across all modules. Communicator and port values need not be buffered since tasks are invoked at most once per mode period and communicator-to-port transactions are done as soon as possible while port-to-communicator transactions are done as late as possible. It is therefore sufficient to have a single copy of each communicator and task port on each host. Additional memory is not required.

7. SCHEDULABILITY

The schedulability analysis for HTL is given a program P , a set $hset$ of hosts, a host map $hmap$ (an assignment from modules to hosts), a WCET map $wemap$ (an assignment from tasks to positive integers), and a WCTT map $wtmap$ (also an assignment from tasks to positive integers). For scheduling, we extend program configurations with a map $tmap$ that assigns to every task t in the task set, a nonnegative integer $tmap(t)$ that indicates the number of time units that the currently active invocation of t has been executed. An *extended program trace* is a sequence of such extended configurations. A *scheduler* is a function that maps every finite extended program trace (representing the past execution of the program) which ends in a waiting configuration, to a scheduling decision for every time unit and every host: it may decide to keep the host idle (no task is chosen from the task set), execute a task (from the task set), or transmit the output of a task (from the task set). The scheduler is assumed to

be discrete time; i.e., it makes decisions at a clock tick. The clock tick is synchronized over all hosts and is a harmonic fraction of the program clock (which is the minimum interval between communicator accesses or mode switches). The WCETs and the WCTTs are specified as multiples of clock ticks.

The scheduler is *time- and transmission-safe* if every infinite extended program trace produced by the scheduler has the following two properties: (1) a task writing to a communicator must have completed execution and output transmission before the communicator update; and (2) if a host is transmitting, then all other hosts are listening (i.e., neither executing nor transmitting). The *schedulability problem* asks, given P , $hset$, $hmap$, $wemap$, and $wmap$, if there exists a time- and transmission-safe scheduler. It can be solved in time linear in the number of extended program configurations [6], but this is often too expensive.

For efficiency reasons, we solve the schedulability problem only for flat HTL programs, namely, for the top-level program $abs(P)$. Then, due to the HTL constraints on refinement, we can construct a scheduler for the complete program P from a scheduler for $abs(P)$. This is achieved by iteratively scheduling every task execution during time slots in which the parent task is executed; see [3]. In other words, HTL guarantees that top-level schedulability is a sufficient condition for schedulability. For top-level schedulability on a single host, we use an EDF scheduling algorithm for tasks with precedences [7]. For top-level schedulability on multiple hosts, one needs to account for transmission times, and we can use the techniques of [4].

8. RELATED WORK

Timed languages. HTL builds on the LET concept pioneered by the Giotto language [1]. LET-based languages include TDL [8], which like Giotto is restricted to one level of periodic tasks; Timed Multitasking (TM) [9], which defines LET properties through deadlines; and xGiotto [10], an event-triggered LET language. HTL differs from Giotto in that logical execution times are defined through the reading and writing of communicator instances. This adds considerable flexibility, and naturally supports task precedences and hierarchical refinement. At the other extreme, in fully event-triggered timed languages, such as xGiotto, scheduling quickly becomes intractable.

Synchronous languages. Esterel [11], Lustre [12], and Signal [13] are based on the synchrony assumption that the execution platform is sufficiently fast as to complete execution before the arrival of the next environment event. As with timed languages, the behavior of synchronous programs is deterministic. While the synchronous languages theoretically subsume timed languages, HTL offers explicitly a program structure that supports the refinement of tasks into task groups with precedences.

Other real-time languages. Most real-time modeling languages such as Simulink [14] rely on simulators and code generators (e.g., Real-Time Workshop [15]) to define the semantics. There are also soft real-time languages for specialized domains: nesC [16] is targeted towards sensor networks; Erlang [17] towards telecommunication; Flex [18] offers flexible trade-offs between time, resources, and precision. Program execution in any of these languages is not schedule independent.

9. CONCLUSION

We presented HTL, a hierarchical coordination language for safety critical hard real-time applications. HTL is built

upon the Logical Execution Time model of task execution and allows parallel composition of modules and horizontal refinement of tasks without modifying the timing behavior. The hierarchical layers of abstraction allows efficient and concise specification without overloading program analysis. We introduced the restrictions on a general HTL program to guarantee schedulability of lower levels if higher levels of abstraction are schedulable. We also presented the operational semantics for HTL and the implementation of an HTL compiler. The compiler checks well-formedness, well-timedness, and schedulability of a given HTL program, flattens the program into a semantically equivalent HTL program with only top-level modules, and then generates code for the Embedded Machine. A steer-by-wire system was used to illustrate the use and the features of HTL.

10. REFERENCES

- [1] Henzinger, T.A., Horowitz, B., Kirsch, C.M.: GIOTTO: A time-triggered language for embedded programming. Proceedings of the IEEE **91** (2003) 84–99
- [2] Henzinger, T.A., Kirsch, C.M.: The Embedded Machine: Predictable, portable real-time code. In: Proceedings of Programming Language Design and Implementation, ACM Press (2002) 315–326
- [3] Ghosal, A., Henzinger, T.A., Iercan, D., Kirsch, C., Sangiovanni-Vincentelli, A.: Hierarchical timing language. Technical report, UC Berkeley (2006)
- [4] Tindell, K., Clark, J.: Holistic schedulability for distributed hard real-time systems. Microprocessing and Microprogramming - Euromicro Journal **40** (1994) 117–134
- [5] Kirsch, C., Sanvido, M., Henzinger, T.: A programmable microkernel for real-time systems. In: Proc. ACM/USENIX Conference on Virtual Execution Environments, ACM Press (2005) 35–45
- [6] Henzinger, T.A., Kirsch, C.M., Majumdar, R., Matic, S.: Time-safety checking for embedded programs. In: EMSOFT 02: Embedded Software. Lecture Notes in Computer Science 2491. Springer-Verlag (2002) 76–92
- [7] Buttazzo, G.: Hard Real-Time Computing Systems. Kluwer Academic Publisher (1997)
- [8] Farcas, E., Farcas, C., Pree, W., Templ, J.: Transparent distribution of real-time components based on logical execution time. SIGPLAN Not. **40** (2005) 31–39
- [9] Liu, J., Lee, E.A.: Timed multitasking for real-time embedded software. IEEE Control Systems Magazine **23** (2003) 65–75
- [10] Ghosal, A., Henzinger, T.A., Kirsch, C.M., Sanvido, M.A.A.: Event-driven programming with logical execution times. In: Hybrid Systems Computation and Control. Lecture Notes in Computer Science 2993. Springer-Verlag (2004)
- [11] Boussinot, F., de Simone, R.: The ESTEREL language. Proceedings of the IEEE **79** (1991) 1293–1304
- [12] Halbwachs, N., Caspi, P., Raymond, P., Pilaud, D.: The synchronous data-flow programming language LUSTRE. Proceedings of the IEEE **79** (1991) 1305–1320
- [13] Guernic, P.L., Borgne, M.L., Gauthier, T., Maire, C.L.: Programming real time applications with SIGNAL. Proceedings of the IEEE **79** (1991) 1321–1336
- [14] Simulink: (<http://www.mathworks.com/products/simulink/>)
- [15] Real-Time-Workshop: (<http://www.mathworks.com/products/rtw/>)
- [16] Gay, D., Levis, P., von Behren, R., Welsh, M., Brewer, E., Culler, D.: The nesC language: A holistic approach to networked embedded systems. In: Proceedings of Programming Language Design and Implementation, ACM Press (2003) 1–11
- [17] Armstrong, J., Virding, R., Wikström, C., Williams, M.: Concurrent Programming in Erlang. Second edn. Prentice-Hall (1996)
- [18] Kenny, K., Lin, K.J.: Building flexible real-time systems using the FLEX language. IEEE Computer **24** (1991) 70–78