Guaranteeing Real-Time Requirements with Resource-Based Calibration of Periodic Processes*

Richard Gerber
Department of Computer Science
University of Maryland
College Park, MD 20742
USA
rich@cs.umd.edu

Seongsoo Hong
Silicon Graphics Inc.
2011 N. Shoreline Blvd. (MS: 7L-802)
Mountain View, CA 94039
USA
sshong@engr.sgi.com

Manas Saksena
Department of Computer Science
Concordia University
Montreal, Quebec H3G 1M8
Canada
manas@cs.concordia.ca

In IEEE Transactions on Software Engineering, 21(7), July 1995

Abstract

This paper presents a comprehensive design methodology for guaranteeing end-to-end requirements of real-time systems. Applications are structured as a set of process components connected by asynchronous channels, in which the endpoints are the system's external inputs and outputs. Timing constraints are then postulated between these inputs and outputs; they express properties such as end-to-end propagation delay, temporal input-sampling correlation, and allowable separation times between updated output values.

The automated design method works as follows: First new tasks are created to correlate related inputs, and an optimization algorithm, whose objective is to minimize CPU utilization, transforms the end-to-end requirements into a set of intermediate rate constraints on the tasks. If the algorithm fails, a restructuring tool attempts to eliminate bottlenecks by transforming the application, which is then re-submitted into the assignment algorithm. The final result is a schedulable set of fully periodic tasks, which collaboratively maintain the end-to-end constraints.

*This research is supported in part by ONR grant N00014-94-10228, NSF grant CCR-9209333, and NSF Young Investigator Award CCR-9357850. An earlier version of this paper appeared by the title of "Guaranteeing End-to-End Timing Constraints by Calibrating Intermediate Processes," in preliminary form in the *Proceedings of IEEE Real-Time System Symposium*, (December 1994).

1 Introduction

Most real-time systems possess only a small handful of *inherent* timing constraints which will "make or break" their correctness. These are called *end-to-end* constraints, and they are established on the systems' external inputs and outputs. Two examples are:

(1) *Temperature updates rely on pressure and temperature readings correlated within 10ms.*

(2) *Navigation coordinates are updated at a minimum rate of 40ms, and a maximum rate 80ms.*

But while such end-to-end timing parameters may indeed be few in number, maintaining *functionally correct* end-to-end values may involve a large set of interacting components. Thus, to ensure that the end-to-end constraints are satisfied, each of these components will, in turn, be subject to their own *intermediate* timing constraints. In this manner a small handful of end-to-end constraints may – in even a modest system – yield a great many intermediate constraints.

The task of imposing timing parameters on the functional components is a complex one, and it mandates some careful engineering. Consider example (2) above. In an avionics system, a "navigation update" may require such inputs as "current heading," airspeed, pitch, roll, etc; each sampled within varying degrees of accuracy. Moreover, these attributes are used by other subsystems, each of which imposes its own tolerance to delay, and possesses its own output rate. Further, the navigation unit may itself have other outputs, which may have to be delivered at rates faster than 40ms, or perhaps slower than 80ms. And to top it off, subsystems may share limited computer resources. A good engineer balances such factors, performs extensive trade-off analysis, simulations and sensitivity analysis, and proceeds to assign the constraints.

These intermediate constraints are inevitably on the conservative side, and moreover, they are conveyed to the programmers in terms of constant values. Thus a scenario like the following is often played out: The design engineers mandate that functional units A, B and C execute with periods 65ms, 22ms and 27ms, respectively. The programmers code up the system, and find that C grossly over-utilizes its CPU; further, they discover that most of C's outputs are not being read by the other subsystems. And so, they go back to the engineers and "negotiate" for new periods – for example 60ms, 10ms and 32ms. This process may continue for many iterations, until the system finally gets fabricated.

This scenario is due to a simple fact: the end-to-end requirements allow many possibilities for the intermediate constraints, and engineers make what they consider to be a rational selection. However, the basis for this selection can only include rough notions of software structuring and scheduling policies – after all, many times the hardware is not even fabricated at this point!

Our Approach. In this paper we present an alternative strategy, which maintains the timing constraints in their end-to-end form for as long as possible. Our design method iteratively instantiates the intermediate constraints, all the while taking advantage of the leeway inherent in the end-to-end constraints. If the assignment algorithm fails to produce a full set of intermediate constraints, potential bottlenecks are identified. At this point an application analysis tool takes over, determines potential solutions to the bottleneck, and if possible, restructures the application to avoid it. The

result is then re-submitted into the assignment algorithm.

Domain of Applicability. Due to the complexity of the general problem, in this paper we place the following restrictions on the applications that we handle.

Restriction 1: We assume our applications possess three classes of timing constraints which we call *freshness*, *correlation* and *separation*.

- A *freshness constraint* (sometimes called propagation delay) bounds the time it takes for data to flow through the system. For example, assume that an external output Y is a function of some system input X. Then a freshness relationship between X and Y might be: "If Y is delivered at time t, then the X-value used to compute Y is sampled no earlier than $t - 10$ms." We use the following notation to denote this constraint: "$F(Y|X) = 10$."
- A *correlation constraint* limits the maximum time-skew between several inputs used to produce an output. For example, if X_1 and X_2 are used to produce Y, then a correlation relationship may be "if Y is delivered at time t, then the X_1 and X_2 values used to compute Y are sampled no more than within 2ms of each other." We denote this constraint as "$C(Y|X_1, X_2) = 2$."
- A *separation constraint* constrains the jitter between consecutive values on a single output channel, say Y. For example, "Y is delivered at a minimum rate of 3ms, and a maximum rate of 13ms," denoted as $l(Y) = 3$ and $u(Y) = 13$, respectively.

While this constraint classification is not complete, it is sufficiently powerful to represent many timing properties one finds in a requirements document. (Our initial examples (1) and (2) are correlation and separation constraints, respectively.) Note that a single output Y_1 may – either directly or indirectly – be subject to several interdependent constraints. For example, Y_1 might require tightly correlated inputs, but may abide with relatively lax freshness constraints. However, perhaps Y_1 also requires data from an intermediate subsystem which is, in turn, shared with a very high-rate output Y_2.

Restriction 2: All subsystems execute on a single CPU. Our approach can be extended for use in distributed systems, a topic we revisit in Section 7. For the sake of presenting the intermediate constraint-assignment technique, in this paper we limit ourselves to uniprocessor systems.

Restriction 3: The entity-relationships within a subsystem are already specified. For example, if a high-rate video stream passes through a monolithic, compute-intensive filter task, this situation may easily cause a bottleneck. If our algorithm fails to find a proper intermediate timing constraint for the filter, the restructuring tool will attempt to optimize it as much as possible. In the end, however, it cannot redesign the system.

Finally, we stress that we are not offering a *completely automatic* solution. Even with a fully periodic task model, assigning periods to the intermediate components is a complex, nonlinear optimization problem which – at worst – can become combinatorially expensive. As for software restructuring, the specific tactics used to remove bottlenecks will often require user interaction.

Problem and Solution Strategy. We duly note the above restrictions, and tackle the intermediate constraint-assignment problem, as rendered by the following ingredients:

- A set of external inputs $\{X_1, \ldots, X_n\}$, outputs $\{Y_1, \ldots, Y_m\}$, and the end-to-end constraints between them.
- A set of intermediate component tasks $\{\tau_1, \ldots, \tau_l\}$.
- A task graph, denoting the communication paths from the inputs, through the tasks, and to outputs.

Solving the problem requires setting timing constraints for the intermediate components, so that all end-to-end constraints are met. Moreover, during any interval of time utilization may never exceed 100%.

Our solution employs the following ingredients: (1) A periodic, preemptive tasking model (where it is our algorithm's duty to assign the rates); (2) a buffered, asynchronous communication scheme, allowing us to keep down IPC times; (3) the period-assignment, optimization algorithm, which forms the heart of the approach; and (4) the software-restructuring tool, which takes over when period-assignment fails.

Related Work. This research was, in large part, inspired by the real-time transaction model proposed by Burns *et. al.* in [3]. While the model was formulated to express database applications, it can easily incorporate variants of our *freshness* and *correlation* constraints. In the analogue to freshness, a persistent object has "absolute consistency within t" when it corresponds to real-world samples taken within maximum drift of t. In the analogue to correlation, a set of data objects possesses "relative consistency within t" when all of the set's elements are sampled within an interval of time t.

We believe that in output-driven applications of the variety we address, separation constraints are also necessary. Without postulating a minimum rate requirement, the freshness and correlation constraints can be vacuously satisfied – by never outputting any values! Thus the separation constraints enforce the system's progress over time.

Burns *et. al.* also propose a method for deriving the intermediate constraints; as in the data model, this approach was our departure point. Here the high-level requirements are re-written as a set of constraints on task periods and deadlines, and the transformed constraints can hopefully be solved. There is a big drawback, however: the correlation and freshness constraints can inordinately tighten deadlines. E.g., if a task's inputs must be correlated within a very tight degree of accuracy – say, several nanoseconds – the task's deadline has to be tightened accordingly. Similar problems accrue for freshness constraints. The net result may be an over-constrained system, and a potentially unschedulable one.

Our approach is different. With respect to tightly correlated samples, we put the emphasis on simply getting the data into the system, and then passing through in due time. However, since this in turn causes many different samples flowing through the system at varying rates, we perform "traffic control" via a novel use of "virtual sequence numbering." This results in significantly looser periods, constrained mainly by the freshness and separation requirements. We also present a period assignment problem which is optimal – though quite expensive in the worst case.

This work was also influenced by Jeffay's "real-time producer/consumer model" [10], which possesses a task-graph structure similar to ours. In this model rates are chosen so that all messages "produced" are eventually "consumed." This semantics leads to a tight coupling between the execution of a consumer to that of its producers; thus it seems difficult to accommodate relative constraints such as those based on freshness.

Klein *et. al.* surveys the current engineering practice used in developing industrial real-time systems [11]. As is stressed, the intermediate constraints should be primarily a function of the end-to-end constraints, but should, if possible, take into account sound real-time scheduling techniques. At this point, however, the "state-of-the-art" is the practice of trial and error, as guided by engineering experience. And this is exactly the problem we address in this paper.

The remainder of the paper is organized as follows. In Section 2 we introduce the application model and formally define our problem. In Section 3 we show our method of transforming the end-to-end constraints into intermediate constraints on the tasks. In Section 4 we describe the constraint-solver in detail, and push through a small example. In Section 5 we describe the application transformer, and in Section 6 we show how the executable application is finally built.

2 Problem Description and Overview of Solution

We re-state our problem as follows:

- Given a task graph with end-to-end timing constraints on its inputs and outputs,
- Derive periods, offsets and deadlines for every task,
- Such that the end-to-end requirements are met.

In this section we define these terms, and present an overview of our solution strategy.

2.1 The Asynchronous Task Graph

An application is rendered in an asynchronous task graph (ATG) format, where for a given graph $G(V, E)$:

- $V = P \cup D$, where $P = \{\tau_1, \ldots, \tau_n\}$, i.e., the set of tasks; and $D = \{d_1, \ldots, d_m\}$, a set of asynchronous, buffered channels. We note that the external outputs and inputs are simply typed nodes in D.
- $E \subseteq (P \times D) \cup (D \times P)$ is a set of directed edges, such that if $\tau_i \rightarrow d_j$ and $\tau_l \rightarrow d_j$ are both in E, then $\tau_i = \tau_l$. That is, each channel has a single-writer/multi-reader restriction.
- All $\tau_i \in P$ have the following attributes: a period T_i, an offset $O_i \geq 0$ (denoting the earliest start-time from the start-of-period), a deadline $D_i \leq T_i$ (denoting the latest finish-time relative to the start-of-period), and a maximum execution time e_i. The interval $[O_i, D_i]$ constrains the window W_i of execution, where $W_i = D_i - O_i$.

Note that initially the T_i's, O_i's and D_i's are open variables, and they get instantiated by the constraint-solver.

The semantics of an ATG is as follows. Whenever a task τ_i executes, it reads data from all incoming channels d_j corresponding to the edges $d_j \rightarrow \tau_i$, and writes to all channels d_l corresponding to the edges $\tau_i \rightarrow d_l$. The actual ordering imposed on the reads and writes is inferred by the task τ_i's structure.

All reads and writes on channels are asynchronous and non-blocking. While a writer always inserts a value onto the end of the channel, a reader can (and many times will) read data from any location. For example, perhaps a writer runs at a period of 20ms, with two readers running at 120ms and 40ms, respectively. The first reader may use every sixth value (and neglect the others), whereas the second reader may use every other value.

But this scheme raises a "chicken and egg" issue, one of many that we faced in this work. One of our objectives is to support software reuse, in which functional components may be deployed in different systems – and have their timing parameters automatically calibrated to the physical limitations of each. But this objective would be hindered if a designer had to employ the following tedious method: (1) to *first* run the constraint-solver, which would find the T_i's, and *then*, based on the results; (2) to hand-patch all of the modules with specialized IPC code, ensuring that the intermediate tasks correctly correlate their input samples.

Luckily, the ATG semantics enables us to automatically support this process. Consider the ATG in Figure 1(A), whose node τ_4 is "blown up" in Figure 1(B). As far as the programmer is concerned the task τ_4 has a (yet-to-be-determined) period T_4, and a set of asynchronous channels, accessible via generic operations such as "**Read**" and "**Write**." Moreover, the channels can be treated both as unbounded, and as non-blocking.

After the constraint-assignment algorithm determines the task rates, a post-processing phase determines the actual space required for each channel. Then they are automatically implemented as circular, slotted buffers. This is accomplished by running an "awk" script on each module, which instantiates each "**Read**" and "**Write**" operation to select the correct input value.

This type of scheme allows us to minimize the overhead incurred when blocking communication is used, and to concentrate exclusively on the assignment problem. In fact – as we show in the sequel – communication can be *completely unconditional*, in that we do not even require short locking for consistency. However, we pay a price for avoiding this overhead; namely, that the period assignments must ensure that no writer can overtake a reader currently accessing its slot.

Moreover, we note that our timing constraints define a system driven by time and *output requirements*. This is in contrast to reactive paradigms such as ESTEREL [4], which are input-driven. Analogous to the "conceptually infinite buffering" assumptions, the rate assignment algorithm assumes that the external inputs are always fresh and available. The *derived* input-sampling rates then determine the *true* requirements on input-availability. And since an input X can be connected to another ATG's output Y, these requirements would be imposed on Y's timing constraints.

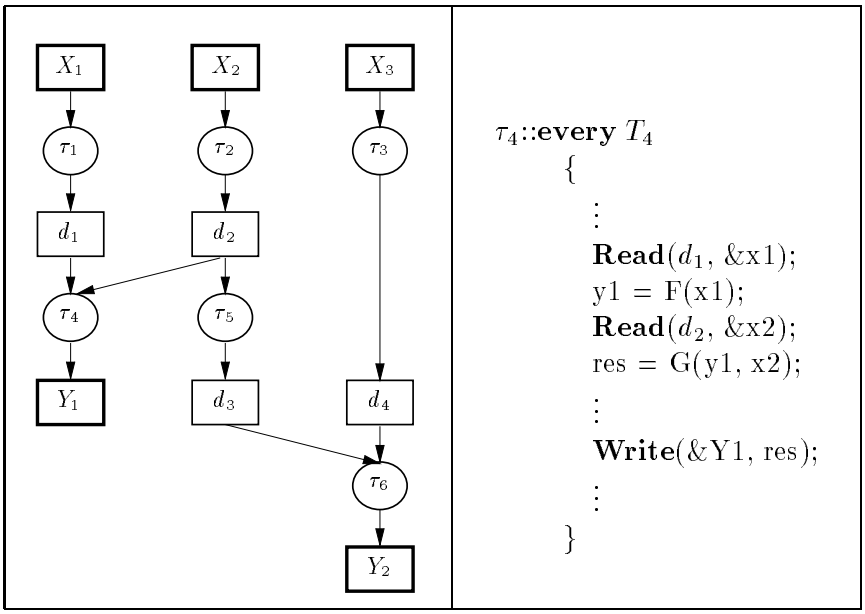

Figure 1: (A) A task graph and (B) code for τ_4.

2.2 A Small Example

As a simple illustration, consider the system whose ATG is shown in Figure 1(A). The system is composed of six interacting tasks with three external inputs and two external outputs. The application's characteristics are as follows:

Freshness	$F(Y_1\|X_1) = 30, \quad F(Y_1\|X_2) = 30$ $F(Y_2\|X_2) = 20, \quad F(Y_2\|X_3) = 15$
Correlation	$C(Y_1\|X_1, X_2) = 3$ $C(Y_2\|X_2, X_3) = 4$
Separation	$l(Y_1) = 21 \quad u(Y_1) = 31$ $l(Y_2) = 29 \quad u(Y_2) = 47$
Max Execution Times:	$e_1 = 6 \quad e_2 = 3 \quad e_3 = 3$ $e_4 = 2 \quad e_5 = 3 \quad e_6 = 2$

While the system is small, it serves to illustrate several facets of the problem: (1) There may be many possible choices of rates for each task; (2) correlation constraints may be tight compared to the allowable end-to-end delay; (3) data streams may be shared by several outputs (in this case that originating at X_2); and (4) outputs with the tightest separation constraints may incur the highest execution-time costs (in this case Y_1, which exclusively requires τ_1).

2.3 Problem Components

Guaranteeing the end-to-end constraints actually poses three sub-problems, which we define as follows.

Correctness: Let $\mathcal{C}$ be the set of derived, intermediate constraints and $\mathcal{E}$ be the set of end-to-end constraints. Then all system behaviors that satisfy $\mathcal{C}$ also satisfy $\mathcal{E}$.

Feasibility: The task executions inferred by $\mathcal{C}$ never demand an interval of time during which utilization exceeds 100%.

Schedulability: There is a scheduling algorithm which can efficiently maintain the intermediate constraints $\mathcal{C}$, and preserve feasibility.

In the problem we address, the three issues cannot be decoupled. Correctness, for example, is often treated as verification problem using a logic such as RTL [9]. Certainly, given the ATG we could formulate $\mathcal{E}$ in RTL and query whether the constraint set is satisfiable. However, a "yes" answer would give us little insight into finding a good choice for $\mathcal{C}$ – which must, after all, be simple enough to schedule. Or, in the case of methods like model-checking ([1], etc.), we could determine whether $\mathcal{C} \Rightarrow \mathcal{E}$ is invariant with respect to the system. But again, this would be an *a posteriori* solution, and assume that we already possess $\mathcal{C}$. On the other hand, a system that is feasible may still not be schedulable under a *known* algorithm; i.e., one that can be efficiently managed by a realistic kernel.

In this paper we put our emphasis on the first two issues. However, we have also imposed a task model for which the greatest number of efficient scheduling algorithms are known: simple, periodic dispatching with offsets and deadlines. In essence, by restricting $\mathcal{C}$'s free variables to the T_i's, O_i's and D_i's, we ensure that feasible solutions to $\mathcal{C}$ can be easily checked for schedulability.

The problem of scheduling a set of periodic real-time tasks on a single CPU has been studied for many years. Such a task set can be dispatched by a calendar-based, non-preemptive schedule (e.g., [16, 17, 18]), or by a preemptive, static-priority scheme (e.g., [5, 12, 13, 15]). For the most part our results are independent of any particular scheduling strategy, and can be used in concert with either non-preemptive or preemptive dispatching.

However, in the sequel we frequently assume an underlying static-priority architecture. This is for two reasons. First, a straightforward priority assignment can often capture most of the ATG's precedence relationships, which obviates the need for superfluous offset and deadline variables. Thus the space of feasible solutions can be simplified, which in turn reduces the constraint-solver's work. Second, priority-based scheduling has recently been shown to support all of the ATG's inherent timing requirements: pre-period deadlines [2], precedence constrained sub-tasks [8], and offsets [14]. A good overview to static priority scheduling may be found in [5].

2.4 Overview of the Solution

Our solution is carried out in a four-step process, as shown in Figure 2. In **Step 1**, the intermediate constraints $\mathcal{C}$ are derived, which postulates the periods, deadlines and offsets as free variables. The challenge here is to balance several factors – correctness, feasibility and simplicity. That is, we

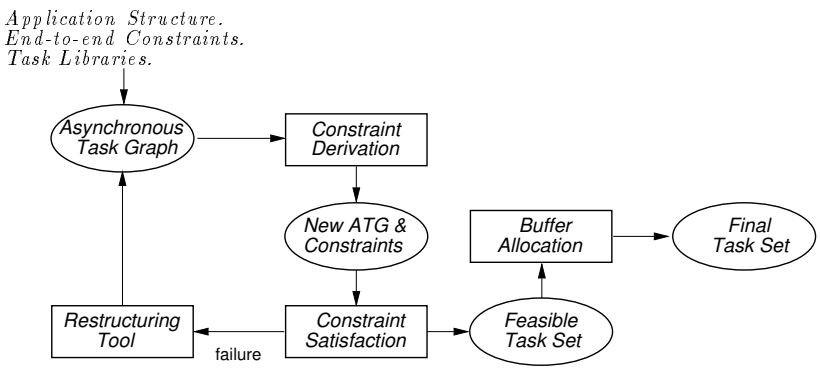

Figure 2: Overview of the approach.

require that any solution to $\mathcal{C}$ will enforce the end-to-end constraints $\mathcal{E}$, and that any solution must also be feasible. At the same time, we want to keep $\mathcal{C}$ as simple as possible, and to ensure that finding a solution is a relatively straightforward venture. This is particularly important since the feasibility criterion – defined by CPU utilization – introduces non-linearities into the constraint set. In balancing our goals we impose additional structure on the application; e.g., by creating new sampler tasks to get tightly correlated inputs into the system.

In **Step 2** the constraint-solver finds a solution to $\mathcal{C}$, which is done in several steps. First $\mathcal{C}$ is solved for the period variables, the T_i's, and then the resulting system is solved for the offsets and deadlines. Throughout this process we use several heuristics, which exploit the ATG's structure.

If a solution to $\mathcal{C}$ cannot be found, the problem often lies in the original design itself. For example, perhaps a single, stateless server handles inputs from multiple clients, all of which run at wildly different rates. **Step 3**'s restructuring tool helps the programmer eliminate such bottlenecks, by automatically replicating strategic parts of the ATG.

In **Step 4**, the derived rates are used to reserve memory for the channels, and to instantiate the "**Read**" and "**Write**" operations. For example, consider τ_4 in Figure 1(B), which reads from channels d_1 and d_2.

Now, assume that the constraint-solver assigns τ_4 and τ_2 periods of 30ms and 10ms, respectively. Then τ_4's **Read** operation on d_2 would be replaced by a macro, which would read every third data item in the buffer – and would skip over the other two.

Harmonicity. The above scheme works only if a producer can always ensure that it is not overtaking its consumers, and if the consumers can always determine which data item is the correct one to read. For example, τ_4's job in managing d_2 is easy – since $T_s = 10$ms and $T_4 = 30$ms, τ_4 will read every third item out of the channel.

But τ_4 has another input channel, d_1; moreover, temporally correlated samples from the two channels have to be used to produce a result. What would happen if the solver assigned τ_1 a period of 30ms, but gave τ_2 a period of 7ms?

If the tasks are scheduled in rate-monotonic order, then d_2 is filled five times during τ_4's first frame, four times during the second frame, etc. In fact since 30 and 7 are relatively prime, τ_4's

selection logic to correlate inputs would be rather complicated. One solution would be to time-stamp each input X_1 and X_2, and then pass these stamps along with all intermediate results. But this would assume access to a precise hardware timer; moreover, time-stamps for multiple inputs would have to be composed in some manner. Worst of all, each small data value (e.g., an integer) would carry a large amount of reference information.

The obvious solution is the one that we adopt: to ensure that every "chain" possesses a common base clock-rate, which is exactly the rate of the task at the head of the chain. In other words, we impose a harmonicity constraint between (producer, consumer) pairs; (i.e., pairs (τ_p, τ_c) where there are edges $\tau_p \rightarrow d$ and $d \rightarrow \tau_c$.)

Definition 2.1 (Harmonicity) *A task τ_2 is harmonic with respect to a task τ_1 if T_2 is exactly divisible by T_1 (represented as $T_2|T_1$[1]).*

Consider Figure 1(A), in which there are three chains imposing harmonic relationships. In this tightly coupled system we have that $T_4|T_1$, $T_4|T_2$, $T_5|T_2$, $T_6|T_5$ and $T_6|T_3$.

3 Step 1: Deriving the Constraints

In this section we show the derivation process of intermediate constraints, and how they (conservatively) guarantee the end-to-end requirements. We start the process by synthesizing the intermediate correlation constraints, and then proceed to treat freshness and separation.

3.1 Synthesizing Correlation Constraints

Recall our example task graph in Figure 3(A), where the three inputs X_1, X_2 and X_3 are sampled by three separate tasks. If we wish to guarantee that τ_1's sampling of X_1 is correctly correlated to τ_2's sampling of X_2, we must pick short periods for both τ_1 and τ_2. Indeed, in many practical real-time systems, the correlation requirements may very well be tight, and way out of proportion with the freshness constraints. This typically results in periods that get tightened exclusively to accommodate correlation, which can easily lead to gross over-utilization. Engineers often call this problem "over-sampling," which is somewhat of a misnomer, since sampling rates may be tuned expressly for coordinating inputs. Instead, the problem arises from poor coupling of the sampling and computational activities.

Thus our approach is to decouple these components as much as possible, and to create specialized samplers for related inputs. For a given ATG, the sampler derivation is performed in the following manner.

[1] $x|y$ iff $\exists \alpha ::\ \alpha y = x$ and $\alpha \geq 1$, where α is an integer.

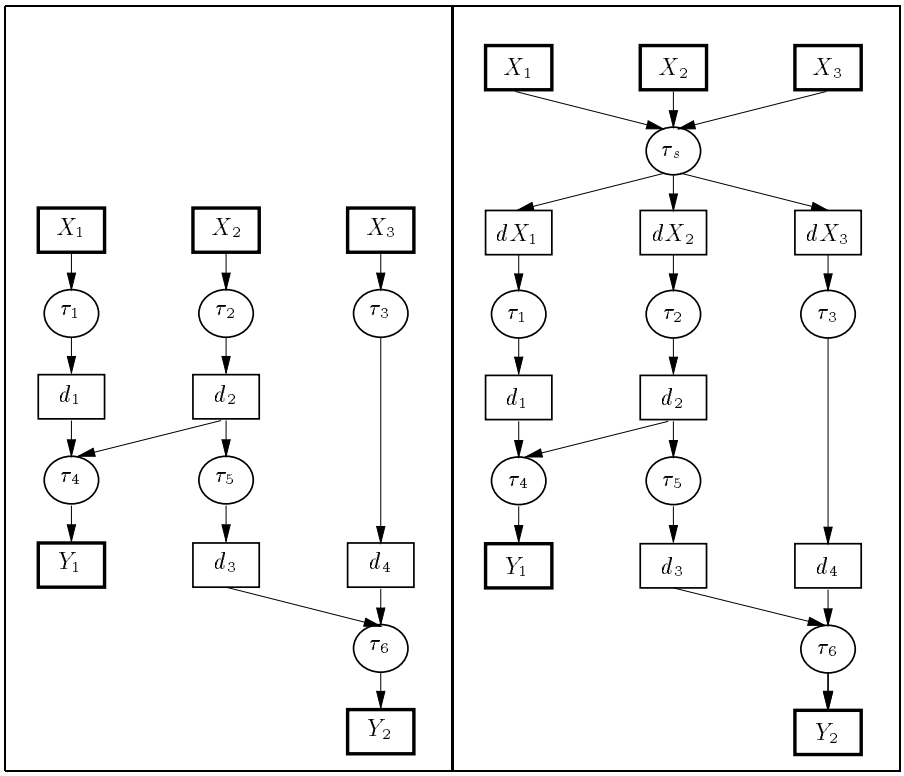

Figure 3: (A) Original task graph and (B) transformed task graph.

foreach Correlation constraint $C_l(Y_k|X_{l_1},\ldots,X_{l_m})$
  Create the set of all input-output pairs associated with C_l, i.e.,
    $T_l := \{(X_{l_i},Y_k)|\, X_{l_i} \in \{X_{l_1},\ldots,X_{l_m}\}\}$
foreach T_l, **foreach** T_k
  If there's a common input X such that there exist outputs Y_i, Y_j
    with $(X,Y_i) \in T_l$, $(X,Y_j) \in T_k$, and
  if chains from X to Y_i and X to Y_j share a common task, then
      Set $T_l := T_l \cup T_k$; $T_k := \emptyset$
foreach T_l, identify all associated sampling tasks, i.e.,
  $S_l := \{\tau|(X,Y) \in T_l \wedge X \rightarrow \tau\}$
  If $|S_l| > 1$, create a periodic sampler τ_{s_l} to take samples for inputs in T_l

Thus the incoming channels from inputs T_l to tasks in S_l are "intercepted" by the new sampler task τ_{s_l}.

Returning to our original example, which we repeat in Figure 3(A). Since both correlated inputs share the center stream, the result is a single group of correlated inputs $\{(X_1,X_2,X_3)\}$. This, in turn, results in the formation of the single sampler τ_s. We assume τ_s has a low execution cost of 1. The new, transformed graph is shown at the right column of Figure 3(B).

As for the deadline-offset requirements, a sampler τ_{s_l} is constrained by the following trivial

relationship

$$D_{s_l} - O_{s_l} \leq t_{cor}$$

where t_{cor} is the maximum allowable time-drift on all correlated inputs read by τ_{s_l}.

The sampler tasks ensure that correlated inputs are read into the system within their appropriate time bounds. This allows us to solve for process rates as a function of both the freshness and separation constraints, which vastly reduces the search space.

However we cannot ignore correlation altogether, since merely sampling the inputs at the same time does not guarantee that they will *remain* correlated as they pass through the system. The input samples may be processed by different streams (running at different rates), and thus they may still reach their join points at different absolute times.

For example, refer back to Figure 3, in which $F(Y_2|X_2) > F(Y_2|X_3)$. This disparity is the result of an under-specified system, and may have to be tightened. The reason is simple: if τ_6's period is derived by using correlation as a dominant metric, the resulting solution may violate the tighter freshness constraints. On the other hand, if freshness is the dominant metric, then the correlation constraints may not be achieved.

We solve this problem by eliminating the "noise" that exists between the different set of requirements. Thus, whenever a fresh output is required, we ensure that there are correlated data sets to produce it. In our example this leads to tightening the original freshness requirement $F(Y_2|X_2)$ to $F(Y_2|X_3)$.

Thus we invoke this technique as a general principle. For an output Y with correlated input sets $X_1, \ldots, X_m$, the associated freshness constraints are adjusted accordingly:

$$F(Y|X_1), \ldots, F(Y|X_m) \ := min\{F(Y|X_1), \ldots, F(Y|X_m)\}$$

3.2 Synthesizing Freshness Constraints

Consider a freshness constraint $F(Y|X) = t_f$, and recall its definition:

> *For every output of Y at some time t, the value of X used to compute Y must have been read no earlier that time $t - t_f$.*

As data flows through a task chain from X to Y, each task τ adds two types of delay overhead to the data's end-to-end response time. One type is *execution time*, i.e., the time required for τ to process the data, produce outputs, etc. In this paper we assume that τ's maximum execution time is fixed, and has already been optimized as much as possible by a good compiler.

The other type of delay is *transmission latency*, which is imposed while τ waits for its correlated inputs to arrive for processing. Transmission time is not fixed; rather, it is largely dependent on our derived process-based constraints. Thus minimizing transmission time is our goal in achieving tight freshness constraints.

Fortunately, the harmonicity relationship between producers and consumers allows us to accomplish this goal. Consider a chain $\tau_1, \tau_2, \ldots, \tau_n$, where τ_n is the output task, and τ_1 is the input

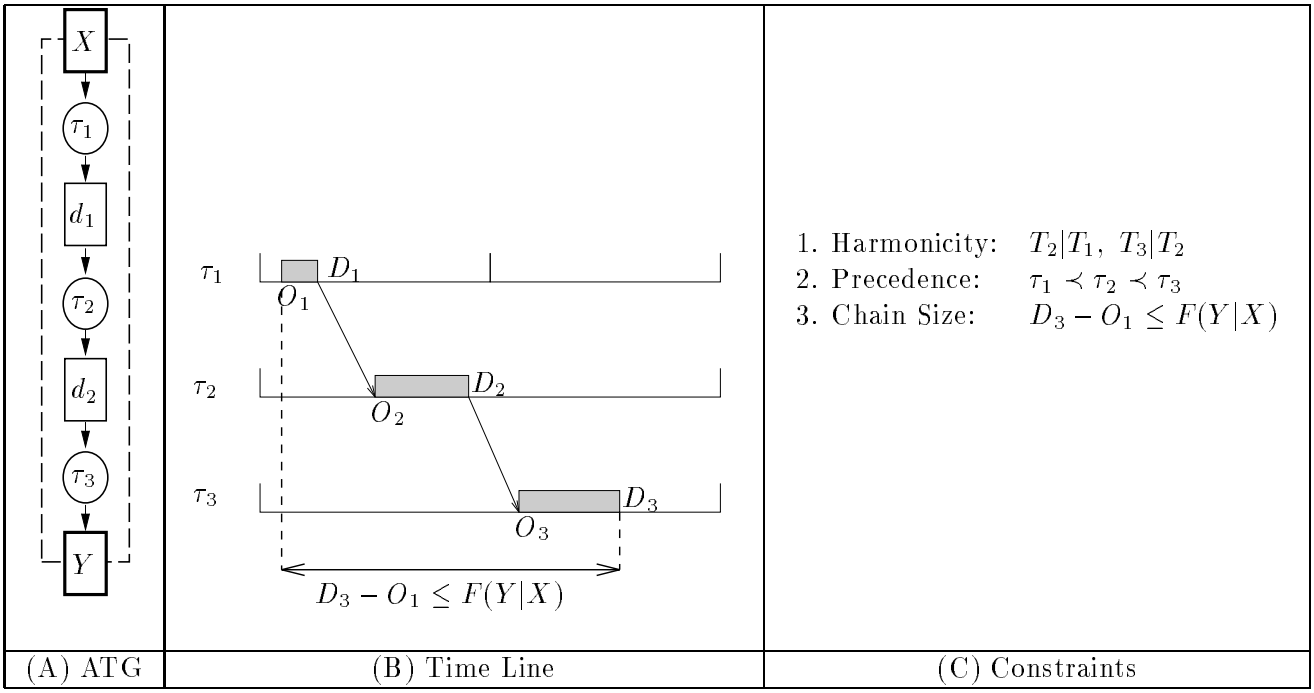

Figure 4: Freshness constraints with coupled tasks.

task. From the harmonicity constraints we get $T_{i+1}|T_i$, for $1 \leq i < n$. Assuming that all tasks are started at time 0, whenever there is an invocation of the output task τ_n, there are simultaneous invocations of every task in the freshness chain.

Consider Figure 4 in which there are three tasks τ_1, τ_2 and τ_3 in a freshness chain. From the harmonicity assumption we have $T_3|T_2$ and $T_2|T_1$.

The other constraints are derived for the entire chain, under the scenario that within each task's minor frame, input data gets read in, it gets processed, and output data is produced. Under these constraints, the worst case end-to-end delay is given by $D_n - O_1$, and the freshness requirement is guaranteed if the following holds:

$$D_n - O_1 \leq t_f$$

Note that we also require a precedence between each producer/consumer task pair. As we show in Figure 4, this can be accomplished via the offset and deadline variables – i.e., by mandating that $D_i \leq O_{i+1}$, for $1 \leq i < n$.

But this approach has the following obvious drawback: *The end-to-end freshness t_f must be divided into fixed portions of slack at each node.* On a global system-wide level, this type of rigid flow control is not the best solution. It is not clear how to distribute the slack between intermediate tasks, without over-constraining the system. More importantly, with a rigid slack distribution, a

$F(Y_1\|X_1)$	$F(Y_1\|X_2)$	$F(Y_2\|X_2)$	$F(Y_2\|X_3)$
$D_4 - O_s \leq 30$	$D_4 - O_s \leq 30$	$D_6 - O_s \leq 15$	$D_6 - O_s \leq 15$
$O_s + e_s + e_1 \leq D_1$	$O_s + e_s + e_2 \leq D_2$	$O_s + e_s + e_2 + e_5 \leq D_5$	$O_s + e_s + e_3 \leq D_3$
$D_1 \leq O_4$	$D_2 \leq O_4$	$D_5 \leq O_6$	$D_3 \leq O_6$
$T_4\|T_1, \quad T_1\|T_s$	$T_4\|T_2, \quad T_2\|T_1$	$T_6\|T_5, \quad T_5\|T_2, \quad T_2\|T_s$	$T_6\|T_3, \quad T_3\|T_s$

Table 1: Constraints due to freshness requirements.

consumer task would not be allowed to execute before its offset, *even if its input data is available.*[2]

Rather, we make a straightforward priority assignment for the tasks in each chain, and let the scheduler enforce the precedence between them. In this manner, we can do away with the intermediate deadline and offset variables. This leads to the following rule of thumb:

> *If the consumer task is not the head or tail of a chain, then its precedence requirement is deferred to the scheduler. Otherwise, the precedence requirement is satisfied through assignment of offsets.*

Example. Consider the freshness constraints for our example in Figure 3(A), $F(Y_1|X_1) = 30$, $F(Y_1|X_2) = 30$, $F(Y_2|X_2) = 15$, and $F(Y_2|X_3) = 15$. The requirement $F(Y_1|X_1) = 30$ specifies a chain window size of $D_4 - O_s \leq 30$. Since τ_1 is an intermediate task we now have the precedence $\tau_s \prec \tau_1$, which will be handled by the scheduler. However, according to our "rule of thumb," we use the offset for τ_4 to handle the precedence $\tau_1 \prec \tau_4$. This leads to the constraints $D_1 \leq O_4$ and $D_s \leq D_1 - e_1$. Similar inequalities are derived for the remaining freshness constraints, the result of which is shown in Table 1.

3.3 Output Separation Constraints

Consider the separation constraints for an output Y, generated by some task τ_i. As shown in Figure 5, the window of execution defined by O_i and D_i constrains the time variability within a period. Consider two frames of τ_i's execution. The widest separation for two successive Y's can occur when the first frame starts as early as possible, and the second starts as late as possible. Conversely, the opposite situation leads to the smallest separation.

Thus, the separation constraints will be satisfied if the following holds true:

$$(T_i + D_i) - O_i \leq u(Y) \quad \text{and} \quad (T_i - D_i) + O_i \geq l(Y)$$

[2]Note that corresponding issues arise in real-time rate-control in high-speed networks.

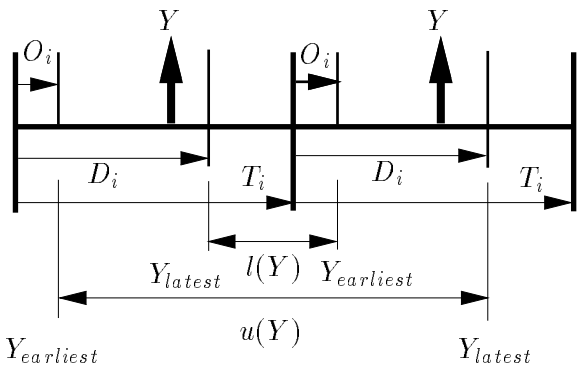

Figure 5: Separation constraints for two frames.

Example. Consider the constraints that arise from output separation requirements, which are induced on the output tasks τ_4 and τ_6. The derived constraints are presented below:

$$(T_4 + D_4) - O_4 \leq u(Y_1) \qquad (T_4 - D_4) + O_4 \geq l(Y_1)$$
$$(T_6 + D_6) - O_6 \leq u(Y_2) \qquad (T_6 - D_6) + O_6 \geq l(Y_2)$$

3.4 Execution Constraints:

Clearly, each task needs sufficient time to execute. This simple fact imposes additional constraints, that ensure that each task's maximum execution time can fit into its window. Recall that (1) we use offset, deadline and period variables for tasks handling external input and output; and (2) we use period variables and precedence constraints for the intermediate constraints.

We can easily preserve these restrictions when dealing with execution time. For each external task τ_i, the following inequalities ensure that window-size is sufficiently large for the CPU demand:

$$O_i + e_i \ \leq \ D_i, \quad O_i \geq 0 \quad D_i \leq T_i$$

On the other hand, the intermediate tasks can be handled by imposing restrictions on their constituent chains. For a single chain, let E denote the chain's execution time from the head to the last intermediate task (i.e., excluding the outputting task, if any). Then the chain-wise execution constraints are:

$$O_h + E \ \leq \ D_m, \quad D_m \leq T_m$$

where O_h is the head's offset, and where D_m and T_m are the last intermediate task's deadline and period, respectively.

Example. Revisiting the example, we have the following execution-time constraints.

$$\begin{array}{llll} O_s + e_s \leq D_s, & O_s \geq 0, & D_s \leq T_s, & \text{sampler task} \\ O_i + e_i \leq D_i, & O_i \geq 0, & D_i \leq T_i, & i = \{4,6\} \\ O_s + e_s + e_i \leq D_i, & D_i \leq T_i & & i = \{1,2,3\} \\ O_s + e_s + e_2 + e_5 \leq D_5, & D_5 \leq T_5 & & \end{array}$$

This completes the set of task-wise constraints $\mathcal{C}$ imposed on our ATG. Thus far we have shown only one part of the problem – how $\mathcal{C}$ can derived from the end-to-end constraints. The end-to-end requirements will be maintained during runtime (1) if a solution to $\mathcal{C}$ is found, and (2) if the scheduler dispatches the tasks according to the solution's periods, offsets and deadlines. Since there are many existing schedulers that can handle problem (2), we now turn our attention to problem (1).

4 Step 2: Constraint Solver

The constraint solver generates instantiations for the periods, deadlines and offsets. In doing so, it addresses the notion of feasibility by using objective functions which (1) minimize the overall system utilization; and (2) maximize the window of execution for each task. Unfortunately, the non-linearities in the optimization criteria – as well as the harmonicity assumptions – lead to a very complex search problem.

We present a solution which decomposes the problem into relatively tractable parts. Our decomposition is motivated by the fact that the non-linear constraints are confined to the period variables, and do not involve deadlines or offsets. This suggests a straightforward approach, which is presented in Figure 6.

1. The entire constraint set $\mathcal{C}$ is projected onto its subspace $\hat{\mathcal{C}}$, constraining only the T_i's.
2. The constraint set $\hat{\mathcal{C}}$ is optimized for minimum utilization.
3. Since we now have values for the T_i's, we can instantiate them in the original constraint set $\mathcal{C}$. This forms a new, reduced set of constraints $\bar{\mathcal{C}}$, all of whose functions are affine in the O_i's and D_i's. Hence solutions can be found via linear optimization.

The back-edge in Figure 6 refers to the case where the nonlinear optimizer finds values for the T_i's, but no corresponding solution exists for the O_i's and D_i's. Hence, a new instantiation for the periods must be obtained – a process that continues until either a solution is found, or all possible values for the T_i's are exhausted.

4.1 Elimination of Offset and Deadline Variables

We use an extension of Fourier variable elimination [6] to simplify our system of constraints. Intuitively, this step may be viewed as the projection of an n dimensional polytope (described by the constraints) onto its lower-dimensional shadow.

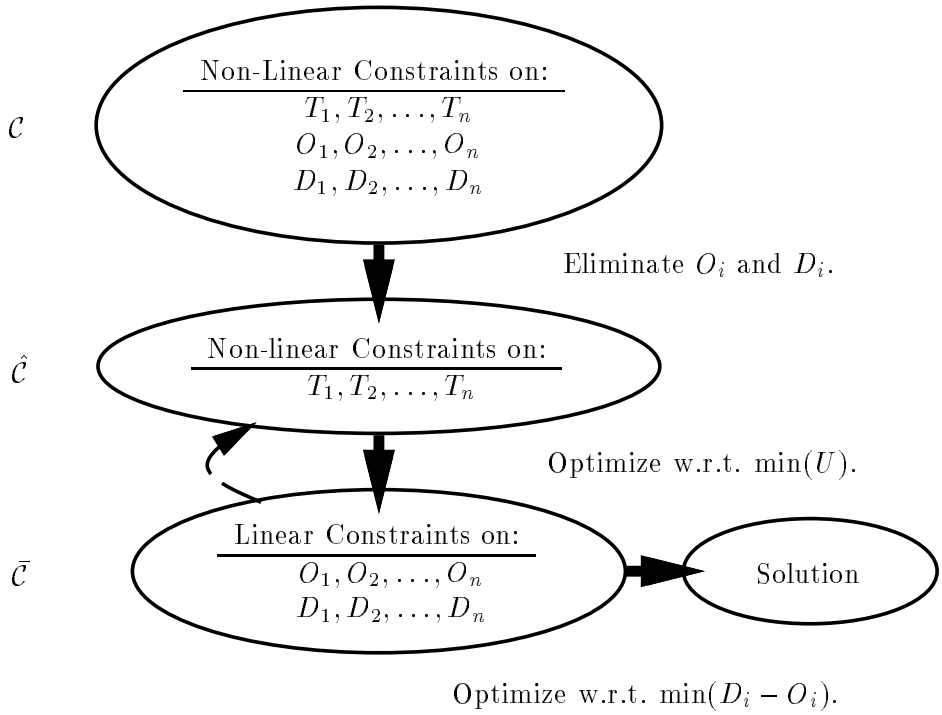

Figure 6: Top level algorithm to obtain task characteristics.

In our case, the n-dimensional polytope is the object described by the initial constraint set $\mathcal{C}$, and where the shadow is the subspace $\hat{\mathcal{C}}$, in which only the T_i's are free. The shadow is derived by eliminating one offset (or deadline) variable at a time, until only period variables remain. At each stage the new set of constraints is checked for inconsistencies (e.g., $0 > 5$). Such a situation means that the original system was over-specified – and the method terminates with failure.

The technique can best be illustrated by a small example. Consider the following two inequalities on $W_4 = D_4 - O_4$:

$$W_4 \geq T_4 + 18 \qquad W_4 \leq 31 - T_4$$

Each constraint defines a line; when W_4 and T_4 are restricted to nonzero solutions, the result is a 2-dimensional polygon. Eliminating the variable W_4 is simple, and is carried out as follows:

$$\begin{aligned} & T_4 + 18 \leq W_4, \quad W_4 \leq 31 - T_4 \\ \Rightarrow \quad & T_4 + 18 \leq 31 - T_4 \\ \Rightarrow \quad & 2T_4 \leq 31 - 18 \\ \Rightarrow \quad & T_4 \leq 6.5 \end{aligned}$$

Since we are searching for integral, nonzero solutions to T_4, any integer in $[0 \ldots 6]$ can be considered a candidate.

When there are multiple constraints on W_4 – perhaps involving many other variables – the same

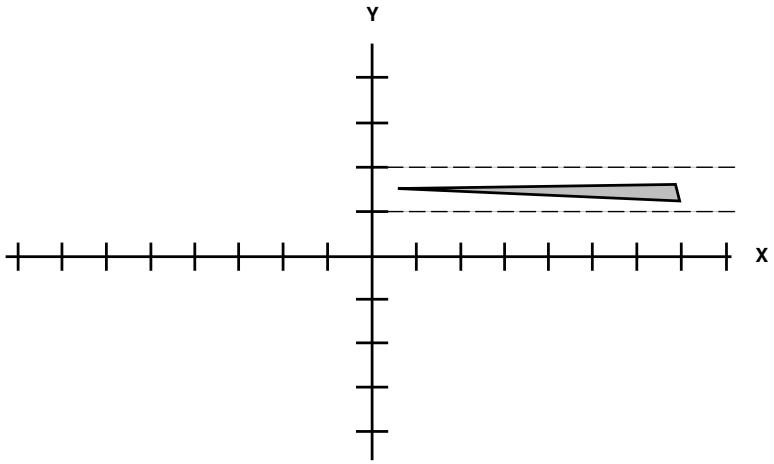

Figure 7: Variable elimination for integer solutions – A deviant case.

process is used. Every constraint "$W_4 \leq \ldots$" is combined with every other constraint "$W_4 \geq \ldots$," until W_4 has been eliminated. The correctness of the method follows simply from the polytope's convexity, i.e., if the original set of constraints has a solution, then the solution is preserved in the shadow.

Unfortunately, the opposite is not true; hence the the requirement for the back-edge in Figure 6. As we have stated, the refined constraint set $\hat{C}$ may possess a solution for the T_i's that do not correspond to any integral-valued O_i's and D_i's. This situation occasionally arises from our quest for integer solutions to the T_i's – which is essential in preserving our harmonicity assumptions.

For example, consider the triangle in Figure 7. The X-axis projection of the triangle has seven integer-solutions. On the other hand, none exist for Y, since all of the corresponding real-valued solutions are "trapped" between 1 and 2.

If, after obtaining a full set of T_i's, we are left without integer values for the O_i's and D_i's, we can resort to two possible alternatives:

1. Search for rational solutions to the offsets and deadlines, and reduce the clock-granularity accordingly, or

2. Try to find new values for the T_i's, which will hopefully lead to a full integer solution.

The Example Application – From $\mathcal{C}$ to $\hat{\mathcal{C}}$. We illustrate the effect of variable elimination on the example application presented earlier. The derived constraints impose lower and upper bounds on task periods, and are shown below. Also remaining are the original harmonicity constraints.

Linear Constraints	τ_s	τ_1	τ_2	τ_3	τ_4	τ_5	τ_6
	$1 \leq T_s$	$7 \leq T_1$	$4 \leq T_2$	$4 \leq T_3$	$23 \leq T_4 \leq 29$	$7 \leq T_5$	$31 \leq T_6 \leq 45$
Harmonicity Constraints	$T_4 \mid T_1$, $T_1 \mid T_s$, $T_4 \mid T_2$, $T_2 \mid T_1$, $T_6 \mid T_5$, $T_5 \mid T_2$, $T_2 \mid T_s$, $T_6 \mid T_3$, $T_3 \mid T_s$						

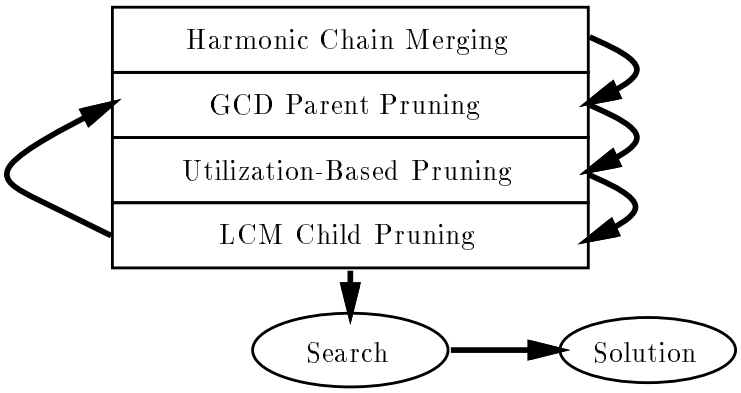

Figure 8: Finding the T's – Pruning.

Here the constraints on the output tasks (τ_4 and τ_6) stem from the separation constraints, which impose upper and lower bounds on the periods.

4.2 From $\hat{\mathcal{C}}$ to $\bar{\mathcal{C}}$: Deriving the Periods

Once the deadlines and offsets have been eliminated, we have a set of constraints involving only the task periods. The objective at this point is to obtain a feasible period assignment which (1) satisfies the derived linear equations; (2) satisfies the harmonicity assumptions; and (3) is subject to a realizable utilization, i.e., $U = \sum \frac{e_i}{T_i} \leq 1$.

As in the example above, the maximum separation constraints will typically mandate that the solution-space for each T_i be bounded from above. Thus we are faced with a decidable problem – albeit a complex one. In fact there are cases which will defeat all known algorithms. In such cases there is no alternative to traversing the entire Cartesian-space

$$[l_1, u_1] \times [l_2, u_2] \times \ldots [l_n, u_n]$$

where there are n tasks, and where each T_i may range within $[l_i, u_i]$.

Fortunately the ATG's structure gives rise to four heuristics, each of which can aggressively prune the search space. The strategy is pictorially rendered in Figure 8.

Let $Pred(i)$ ($Succ(i)$) denote the set of tasks which are predecessors (successors) of task τ_i, i.e., those tasks from (to) which there is a directed path to (from) τ_i. Since the harmonicity relationship is transitive, we have that if $\tau_j \in Succ(\tau_i)$, it follows that $T_j | T_i$. This simple fact leads to three of our four heuristics.

Harmonic Chain Merging extends from following observation: we do not have to solve for each T_i as if it is an arbitrary variable in an arbitrary function. Rather, we can combine chains of processes, and then solve for their base periods. This dramatically reduces the number of free variables.

GCD Parent Pruning is used to ensure that the head of each chain forms a greatest-common-divisor for the entire chain. All tuples which violate this property are deleted from the set of candidate solutions.

Utilization Pruning ensures that candidate solutions maintain a CPU utilization under 100%, a rather desirable constraint in a hard real-time system.

LCM Child Pruning takes an opposite approach to GCD Parent Pruning. It ensures that a task's period is a multiple of its predecessors' combined LCM. Since this is the most expensive pruning measure, it is saved for last.

Harmonic Chain Merging. The first step in the pruning process extends from a simple, but frequently overlooked, observation: that tasks often over-sample for no discernible reason, and that unnecessarily low T_i's can easily steal cycles from the tasks that truly need them. For our purposes, this translates into the following rule:

> If a task τ_i executes with period T_i, and if some $\tau_j \in Pred(\tau_i)$ has the property that $Succ(\tau_j) = \{\tau_i\}$, then τ_j should also execute with period T_i.

In other words, we will never run a task faster than it needs to be run. In designs where the periods are ad-hoc artifacts, tuned to achieve the end-to-end constraints, such an approach would be highly unsafe. Here the rate constraints are *analytically derived* directly from the end-to-end requirements. *We know "how fast" a task needs to be run, and it makes no sense to run it faster.*

This allows us to simplify the ATG by merging nodes, and to limit the number of free variables in the problem. The method is summed up in the following steps:

(1) If $\tau_i \in Pred(\tau_j)$, then $T_j | T_i$ and consequently, $T_i \leq T_j$. The first pruning takes place by propagating this information to tighten the period bounds. Thus, for each task τ_i, the bounds are tightened as follows:

$$l_i = \max\{l_k \mid \tau_k \in Pred(\tau_i)\}$$
$$u_i = \min\{u_k \mid \tau_k \in Succ(\tau_i)\}$$

(2) The second step in the algorithm is to simplify the task graph. Consider a task τ_i, which has an outgoing edge $\tau_i \rightarrow \tau_j$. Suppose $u_i \geq u_j$. Then the maximum value of T_i is constrained only by harmonicity restrictions. The simplification is done by merging τ_i and τ_j, whenever it is safe to set $T_i = T_j$, i.e., the restricted solution space contains the optimal solution. The following two rules give the condition when it safe to perform this simplification.

> **Rule 1:** If a vertex τ_i has a single outgoing edge $\tau_i \rightarrow \tau_j$, then τ_i is merged with τ_j.
>
> **Rule 2:** If $Succ(\tau_i) \subseteq (Succ(\tau_j) \cup \{\tau_j\})$ for some edge $\tau_i \rightarrow \tau_j$, then τ_i is merged with τ_j.

Consider the graph in Figure 9. The parenthesized numbers denote the costs of corresponding nodes. In the graph, the nodes τ_3, τ_5, and τ_1 have a single outgoing edge. Using **Rule 1**, we

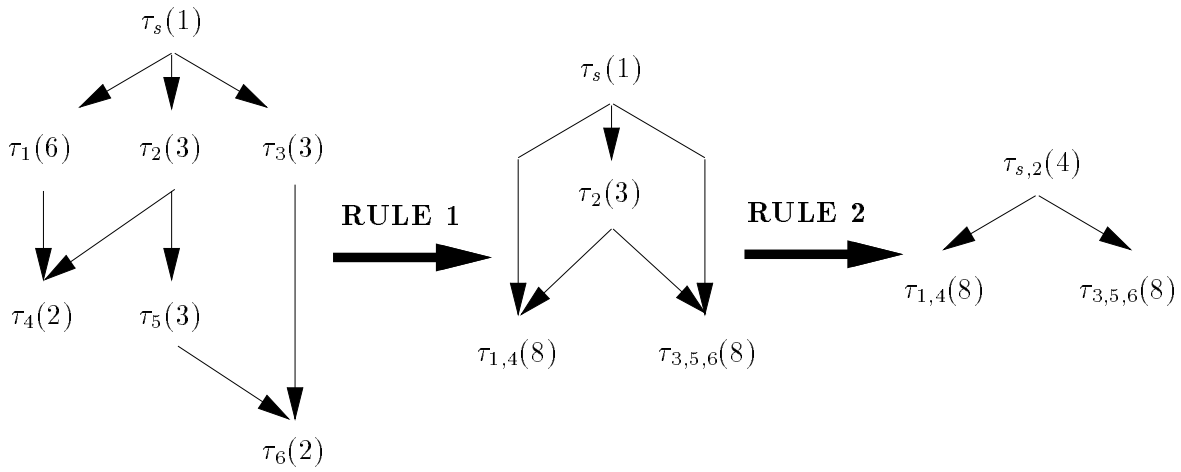

Figure 9: Task graph for harmonicity and its simplification.

merge τ_3 and τ_5 with τ_6, and τ_1 with τ_4. In the simplified graph, $Succ(\tau_s) = \{\tau_4, \tau_6, \tau_2\}$ and $Succ(\tau_2) = \{\tau_4, \tau_6\}$. Thus, we can invoke **Rule 2** to merge τ_s with τ_2.

This scheme manages to reduce our original seven tasks to three sets of tasks, where each set can be represented as a pseudo-task with its own period, and an execution time equal to the sum of its constituent tasks.

At this point we have reduced the structure of the ATG as much as possible, and we turn to examining the search space itself. But even here, we can still use the harmonicity restrictions and utilization bounds as aggressively as possible, with the objective of limiting our search. Let Φ denote the set of feasible solutions for a period T_i, whose initial solution space is denoted as $\Phi_i = \{T_i \mid l_i \leq T_i \leq u_i\}$. The pruning takes place by successively refining and restricting Φ_i for each task.

Algorithm 4.1 combines our three remaining pruning techniques – *GCD Parent Pruning*, *Utilization Pruning* and *LCM Child Pruning*. In the following paragraphs, we explain these steps in detail, and show how they are applied to our example.

GCD Parent Pruning. Consider any particular node τ_i in the task graph. The feasible set of solutions for this node can be reduced by considering the harmonicity relationship with all its successor nodes.

$$\Phi_i := \{T_i \in \Phi_i \mid (\forall \tau_k \in Succ(\tau_i))(\exists T_k \in \Phi_k :: T_k | T_i)\}$$

That is, we restrict Φ_i to values that can provide a base clock-rate for all successor tasks.

Algorithm 4.1 *Prune Feasible Search Space using Harmonicity and Utilization constraints.*

/* U_{max} = maximum allowable utilization. */
/* $T_i^{max} = \max\{T_i \in \Phi_i\}$ */
/* $U_{min} = \sum U_i^{min}$, where $U_i^{min} = e_i/T_i^{max}$ */
Sort the graph in reverse topological order. Let the sorted list be $L = \langle \tau_{i_1}, \tau_{i_2}, \ldots, \tau_{i_n} \rangle$.
for $j := 1$ to n **do** /* traverse the list */
    $\Phi_{i_j} = \{T_{i_j} \in \Phi_{i_j} \mid (\forall \tau_k \in Succ(\tau_{i_j}))(\exists T_k \in \Phi_k :: T_k | T_{i_j})\}$
    /* check utilization condition for each value in Φ_{i_j} */
    foreach $T_{i_j} \in \Phi_{i_j}$ **do**
        foreach $\tau_k \in Succ(\tau_{i_j})$ **do**
            $\hat{U}_k := \frac{e_k}{\hat{T}_k}$, where $\hat{T}_k = \max\{T_k \in \Phi_k \mid T_k | T_{i_j}\}$
            $U_{min} := U_{min} - U_k^{min} + U_k'$
        end
        $U := U - U_{i_j}^{min} + e_i/T_{i_j}$
        if $U > U_{max}$ **then** $\Phi_{i_j} := \Phi_{i_j} - T_{i_j}$
    end
    /* Propagate restricted feasible set to all successors */
    foreach $\tau_k \in Succ(\tau_{i_j})$ **do**
        $\Phi_k := \{T_k \in \Phi_k \mid (\exists T_{i_j} \in \Phi_{i_j} :: T_k | T_{i_j})\}$
    end
end

Figure 10: Pruning feasible space for period derivation.

Within our example, our three merged tasks have the following allowable ranges:

$$\begin{aligned} \Phi_{s,2} &= \{T_{s,2} \mid 4 \leq T_{s,2} \leq 29\} \\ \Phi_{1,4} &= \{T_{1,4} \mid 23 \leq T_{1,4} \leq 29\} \\ \Phi_{3,5,6} &= \{T_{3,5,6} \mid 31 \leq T_{3,5,6} \leq 45\} \end{aligned}$$

The sampler task's period is restricted to values with integral multiples in both $\Phi_{1,4}$ and $\Phi_{3,5,6}$. After deleting members of $\Phi_{s,2}$ that fail to satisfy this property, we are left with the following reduced set:

$$\Phi_{s,2} = \{4, 5, 6, 7, 8, 9, 12, 13, 14\}.$$

Utilization-Based Pruning. Let U_{max} be the upper bound on the utilization that we wish to achieve. At any stage, a lower bound on the utilization for task τ_i is given by:

$$U_i^{min} = \frac{e_i}{T_i^{max}}, \text{where } T_i^{max} = \max\{T_i \in \Phi_i\}$$

If the lower bound on overall utilization $U_{min}(= \sum U_i^{min})$ exceeds U_{max} then there is no solution which satisfies the utilization bound. Now, consider a single task τ_ℓ, and consider a value $\hat{T}_\ell \in \Phi_\ell$. Define $\hat{T}_k$ for all other tasks as follows:

$$\hat{T_k} = \begin{cases} \max\{T_k \in \Phi_k \mid T_k | \hat{T}_\ell\} & \tau_k \in Succ(\tau_l) \\ \max\{T_k \in \Phi_k\} & \tau_k \notin Succ(\tau_l) \end{cases}$$

Then, if $\hat{T}_\ell$ is the period for τ_ℓ, a lower bound on the utilization is given by:

$$U = \sum \frac{e_i}{\hat{T}_i}$$

Clearly, if $U > U_{max}$, then no feasible solution can be obtained with $\hat{T}_\ell$, and hence it may be removed from the feasible set.

Returning to our example, at this point we have the following solution space:

$$\begin{aligned} \Phi_{s,2} &= \{4, 5, 6, 7, 8, 9, 12, 13, 14\} \\ \Phi_{1,4} &= \{T_{1,4} \mid 23 \leq T_{1,4} \leq 29\} \\ \Phi_{3,5,6} &= \{T_{3,5,6} \mid 31 \leq T_{3,5,6} \leq 45\} \end{aligned}$$

Since $\tau_{1,4}$ and $\tau_{3,5,6}$ have no successors, and the utilization bounds are satisfied for all values, no restriction takes place. Now we consider $\tau_{s,2}$, whose period comes from our original sampler task. After testing the possible solutions for utilization, we obtain the reduced set $\Phi_{2,s} = \{9, 12, 13, 14\}$.

LCM Child-Pruning. In general, there may be several chains, each of whose tasks has been restricted by the utilization test. (In our case we only have two chains which share a common source.) In the general case, the reduced feasible set for each task τ_i may be propagated to all successors τ_k. This is done by restricting T_k to integer multiples of T_i.

$$\Phi_k := \{T_k \in \Phi_k \mid (\exists T_i \in \Phi_i :: \ T_k | T_i)\}$$

When we follow this approach in our example, we end up with the following solution space:

$$\begin{aligned} \Phi_{s,2} &= \{9, 12, 13, 14\} \\ \Phi_{1,4} &= \{24, 26, 27, 28\} \\ \Phi_{3,5,6} &= \{36, 39, 42, 45\} \end{aligned}$$

If our objective is to achieve optimality, then examining the remaining candidate solutions is probably unavoidable. In this case, the optimal solution is easily found to be $T_{s,2} = 14, T_{1,4} = 28, T_{3,5,6} = 42$, giving a utilization of 0.7619.

If the remaining solution-space is large, a simple a branch-and-bound heuristic can be employed to control the search. By carefully setting the utilization bound, we can limit the search time required, since the tighter the utilization bound, the greater is the pruning achieved. Thus, by

starting with a low utilization bound, and successively increasing it, we can reduce the amount of search time required to achieve optimally low utilization.

However, if the objective is simply finding a solution – *any solution* – then any of the remaining candidates can be selected.

4.3 Deriving Offsets and Deadlines

Once the task periods are determined, we need to revisit the constraints to find a solution to the deadlines and offsets of the periods. This involves finding a solution which maximizes schedulability.

Variable elimination allows us to select values in the reverse order in which they are eliminated. Suppose we eliminated in following order: $x_1, x_2, \ldots, x_n$. When variable x_i is eliminated, the remaining free variables are $[x_{i+1}, \ldots, x_n]$. Since $[x_{i+1}, \ldots, x_n]$ are already bound to values, the constraints immediately give a lower and an upper bound on x_i.

We use this fact in assigning offsets and deadlines to the tasks. As the variables are assigned values, each variable can be individually optimized. Recall that the feasibility of a task set requires that the task set never demand a utilization greater than one in any time interval. We use a greedy heuristic, which attempts to maximize the window of execution for each task. For tasks which do not have an offset, this is straightforward, and can be achieved by maximizing the deadline. For input/output tasks which have offsets, we also need to fix the position of the window on the time-line. We do this by minimizing the offset for input tasks, and maximizing the deadline for output tasks.

The order in which the variables are assigned is given by the following strategy: First, we assign the windows for each input task, followed by the windows for each output task. Then, we assign the offsets for each task followed by deadline for each output task. Finally, the deadlines for the remaining tasks are assigned in a reverse topological order of the task graph. Thus, an assignment ordering for the example application is given as $\{W_s, W_4, W_6, O_s, D_4, D_6, D_5, D_3, D_1, D_2\}$. The final parameters, derived as a result of this ordering, are shown below.

	τ_s	τ_1	τ_2	τ_3	τ_4	τ_5	τ_6
Period	14	28	14	42	28	42	42
Offset	0	0	0	0	25	0	10
Deadline	3	25	14	10	28	10	15
Exec. Time	1	6	3	3	2	3	2

A feasible schedule for the task set is shown in Figure 11. We note that the feasible schedule can be generated using the fixed priority ordering $\tau_s, \tau_2, \tau_3, \tau_5, \tau_6, \tau_1, \tau_4$.

5 Step 3: Graph Transformation

When the constraint-solver fails, replicating part of a task graph may often prove useful in reducing the system's utilization. This benefit is realized by eliminating some of the tight harmonicity

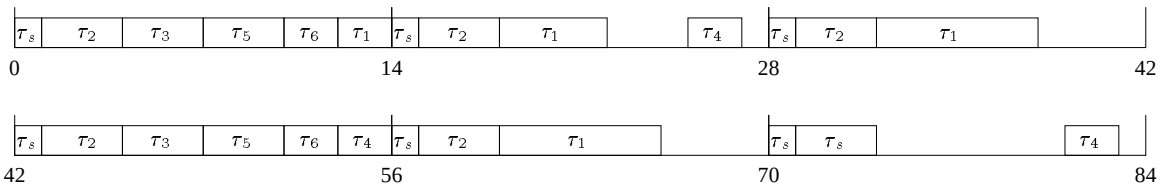

Figure 11: Feasible schedule for example application.

requirements, mainly by decoupling the tasks that possess common producers. As a result, the constraint derivation algorithm has more freedom in choosing looser periods for those tasks.

Recall the example application from Figure 3(B), and the constraints derived in Section 4. In the resulting system, the producer/consumer pair (τ_2, τ_5) has the largest period difference ($T_2 = 14$ and $T_5 = 42$). Note that the constraint solver mandated a tight period for τ_2, due to the coupled harmonicity requirements $T_4|T_2$ and $T_5|T_2$. Thus, we choose to replicate the chain including τ_2 from the sampler (τ_s) to data object d_2. This decouples the data flow to Y_1 from that to Y_2. Figure 12 shows the result of the replication.

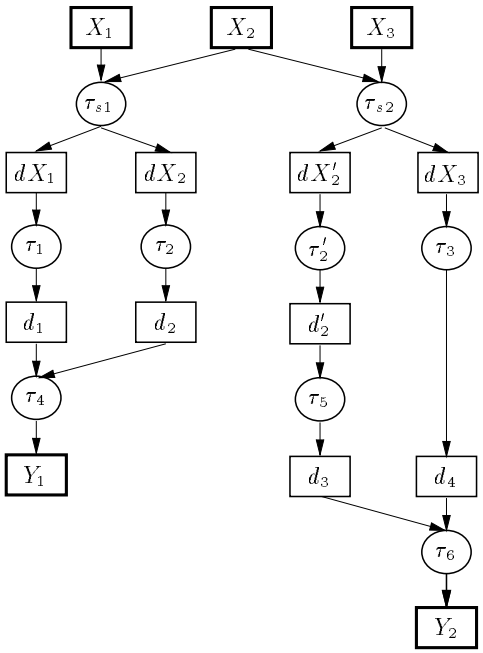

Figure 12: The replicated task graph.

Running the constraint derivation algorithm again with the transformed graph in Figure 12, we obtain the following result. The transformed system has a utilization of 0.6805, which is significantly lower than that of the original task graph (0.7619).

	τ_{s1}	τ_1	τ_2	τ_4	τ_{s2}	τ_2'	τ_3	τ_5	τ_6
Periods	29	29	29	29	45	45	45	45	45
Exec. Time	1	6	3	2	1	3	3	3	2

The subgraph replication technique begins with selecting a producer/consumer pair which requires replication. There exist two criteria in selecting a pair, depending on the desired goal. If the goal is reducing expected utilization, a producer/consumer pair with the maximum period difference is chosen first. On the other hand, if the goal is achieving feasibility, then we rely on the feedback from the constraint solver in determining the point of infeasibility.

After a producer/consumer pair is selected, the algorithm constructs a subgraph using a backward traversal of the task graph from the consumer. In order to avoid excessive replication, the traversal is terminated at the first confluence point. The resulting subgraph is then replicated and attached to the original graph.

The producer task in a replication may, in turn, be further specialized for the output it serves. For example, consider a task graph with two consumers τ_{c1} and τ_{c2} and a common producer τ_p. If we replicate the producer, we have two independent producer/consumer pairs, namely (τ_p, τ_{c1}) and (τ_p', τ_{c2}). Since τ_p' only serves τ_{c2}, we can eliminate all operations that only contribute to the output for τ_{c1}. This is done by *dead code elimination*, a common compiler optimization. The same specialization is done for τ_p.

6 Step 4: Buffer Allocation

Buffer allocation is the final step of our approach, and hence applied to the feasible task graph whose timing characteristics are completely derived. During this step, the compiler tool determines the buffer space required by each data object, and replaces its associated reads and writes with simple macros. The macros ensure that each consumer reads temporally correlated data from several data objects – even when *these* objects are produced at vastly different rates. The reads and writes are nonblocking and asynchronous, and hence we consider each buffer to have a "virtual sequence number."

Combining a set of correlated data at a given confluence point appears to be a nontrivial venture. After all, (1) producers and the consumers may be running at different rates; and (2) the flow delays from a common sampler to the distinct producers may also be different. However, due to the harmonicity assumption the solution strategy is quite simple. Given that there are sufficient buffers for a data object, the following rule is used:

> "Whenever a consumer reads from a channel, it uses the *first* item that was generated within *its* current period."

For example, let τ_p be a producer of a data object d, let $\tau_{c_1}, \ldots, \tau_{c_n}$ be the consumers that read d. Then the communication mechanism is realized by the following techniques (where $L =$

$LCM_{1 \le i \le n}(T_{c_i})$ is the least common multiple of the periods):

(1) The data object d is implemented with $s = L/T_p$ buffers.

(2) The producer τ_p circularly writes into each buffer, one at a time.

(3) The consumer τ_{c_i} reads circularly from slots $(0, T_{c_i}/T_p, \ldots, m \cdot T_{c_i}/T_p)$ where $m = L/T_{c_i} - 1$.

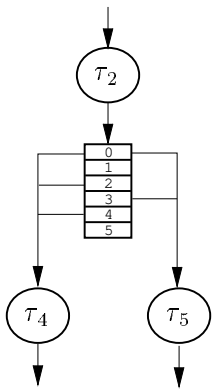

Figure 13: A task graph with buffers.

Consider three tasks τ_2, τ_4 and τ_5 in our example, before we performed graph replication. The two consumer tasks τ_4 and τ_5 run with periods 28 and 42, respectively, while the producer τ_2 runs with period 14. Thus, the data object requires a 6 place buffer ($6 = LCM(28, 42)/14$), and τ_4 reads from slots (0, 2, 4) while τ_5 reads from slots (0, 3). Figure 13 shows the relevant part of the task graph after the buffer allocation.

After the buffer allocation, the compiler tool expands each data object into a multiple place buffer, and replaces each read and write operations with macros that perform proper pointer updates. Figure 14 shows the results of the macro-expansion, after it is applied to τ_4's code from Figure 1(B). Note that τ_1, τ_2 and τ_4 run at periods of 28, 14 and 28, respectively.

7 Conclusion

We have presented a four-step design methodology to help synthesize end-to-end requirements into full-blown real-time systems. Our framework can be used as long as the following ingredients are provided: (1) the entity-relationships, as specified by an asynchronous task graph abstraction; and (2) end-to-end constraints imposed on freshness, input correlation and allowable output separation. This model is sufficiently expressive to capture the temporal requirements – as well as the modular structure – of many interesting systems from the domains of avionics, robotics, control and multimedia computing.

However, the asynchronous, fully periodic model does have its limitations; for example, we cannot support high-level blocking primitives such as RPCs. On the other hand this deficit yields

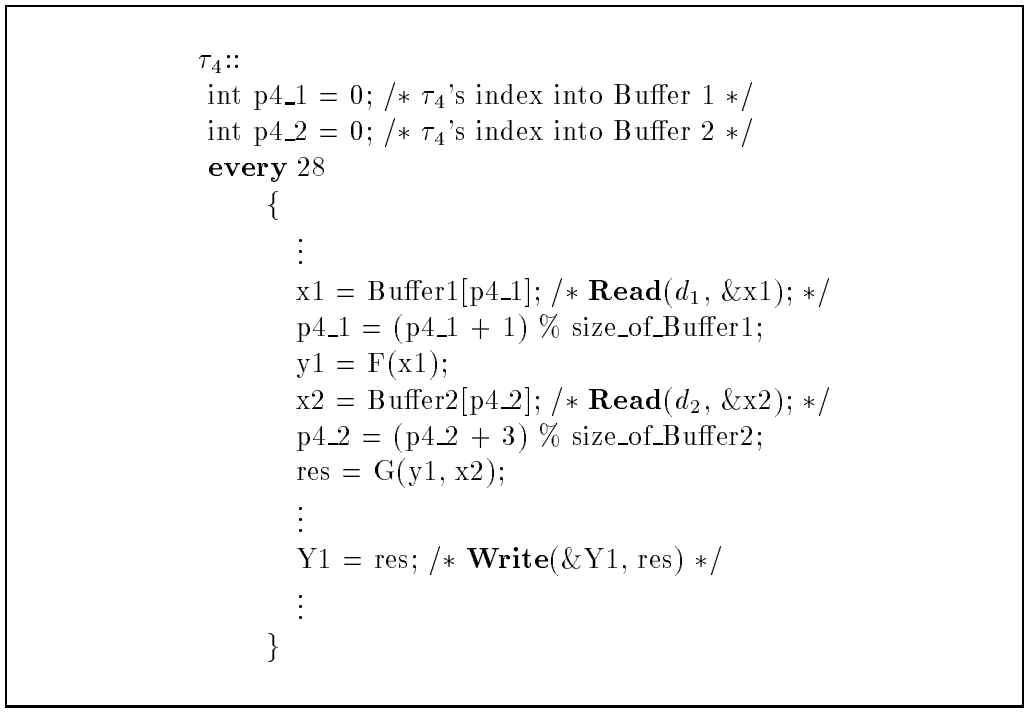

Figure 14: Instantiated code with copy-in/copy-out channels and memory-mapped IO.

significant gains; e.g., handling streamed, tightly correlated data solely via the "virtual sequence numbers" afforded by the rate-assignments.

There is much work to be carried out. First, the constraint derivation algorithm can be extended to take full advantage of a wider spectrum of timing constraints, such as those encountered in input-driven, reactive systems. Also, we can harness finer-grained compiler transformations such as *program slicing* to help transform tasks into read-compute-write-compute phases, which will even further enhance schedulability. We have used this approach in a real-time compiler tool [7], and there is reason to believe that its use would be even more effective here.

We are also streamlining our search algorithm, by incorporating scheduling-specific decisions into the constraint solver. We believe that when used properly, such policy-specific strategies will help significantly in pruning the search space.

But the greatest challenge lies in extending the technique to distributed systems. Certainly a global optimization is impractical, since the search-space is much too large. Rather, we are taking a compositional approach – by finding approximate solutions for each node, and then refining each node's solution-space to accommodate the system's bound on network utilization.

Acknowledgements

The authors gratefully acknowledge Bill Pugh, who was an invaluable resource, critic, and friend throughout the development of this paper. In particular, Bill was our best reference on the topic on nonlinear optimization. We are also grateful for the insightful comments of the TimeWare

group members: Jeff Fischer, Ladan Gharai, Tefvik Bultan and Dong-In Kang (in addition to the authors). In particular, discussions with Ladan and Jeff were great "sounding boards" when we formalized the problem, and they gave valuable advice while we developed a solution.

References

[1] R. Alur, C. Courcoubetis, and D. Dill. Model-checking for real-time systems. In *Proc. of IEEE Symposium on Logic in Computer Science*, 1990.

[2] N. Audsley, A. Burns, M. Richardson, and A. Wellings. Hard real-time scheduling: The deadline-monotonic approach. In *Proceedings of IEEE Workshop on Real-Time Operating Systems and Software*, pages 133–137, May 1991.

[3] N. Audsley, A. Burns, M. Richardson, and A. Wellings. Data consistency in hard real-time systems. Technical Report YCS 203 (1993), Department of Computer Science, University of York, England, June 1993.

[4] G. Berry, S. Moisan, and J. Rigault. ESTEREL: Towards a synchronous and semantically sound high level language for real time applications. In *Proceedings of IEEE Real-Time Systems Symposium*, pages 30–37. IEEE Computer Society Press, December 1983.

[5] A. Burns. Preemptive priority based scheduling: An appropriate engineering approach. In S. Son, editor, *Principles of Real-Time Systems*. Prentice Hall, 1994.

[6] G. Dantzig and B. Eaves. Fourier-Motzkin Elimination and its Dual. *Journal of Combinatorial Theory (A)*, 14:288–297, 1973.

[7] R. Gerber and S. Hong. Semantics-based compiler transformations for enhanced schedulability. In *Proceedings of IEEE Real-Time Systems Symposium*, pages 232–242. IEEE Computer Society Press, December 1993.

[8] M. Harbour, M. Klein, and J. Lehoczky. Fixed Priority Scheduling of Periodic Tasks with Varying Execution Priority. In *Proceedings, IEEE Real-Time Systems Symposium*, pages 116–128, December 1991.

[9] F. Jahanian and A. Mok. Safety analysis of timing properties in real-time systems. *IEEE Transactions on Software Engineering*, 12(9):890–904, September 1986.

[10] K. Jeffay. The real-time producer/consumer paradigm: A paradigm for the construction of efficient, predictable real-time systems. In *ACM/SIGAPP Symposium on Applied Computing*, pages 796–804. ACM Press, February 1983.

[11] M. Klein, J. Lehoczky, and R. Rajkumar. Rate-monotonic analysis for real-time industrial computing. *IEEE Computer*, pages 24–33, January 1994.

[12] C. Liu and J. Layland. Scheduling algorithm for multiprogramming in a hard real-time environment. *Journal of the ACM*, 20(1):46–61, January 1973.

[13] L. Sha, R. Rajkumar, and J. Lehoczky. Priority inheritance protocols: An approach to real-time synchronization. *IEEE Transactions on Software Engineering*, 39:1175–1185, September 1990.

[14] K. Tindell. Using offset information to analyse static priority pre-emptively scheduled task sets. Technical Report YCS 182 (1992), Department of Computer Science, University of York, England, August 1992.

[15] K. Tindell, A. Burns, and A. Wellings. An extendible approach for analysing fixed priority hard real-time tasks. *The Journal of Real-Time Systems*, 6(2):133–152, March 1994.

[16] J. Xu and D. Parnas. Scheduling processes with release times, deadlines, precedence and exclusion relations. *IEEE Transactions on Software Engineering*, 16(3):360–369, March 1990.

[17] X. Yuan, M. Saksena, and A. Agrawala. A Decomposition Approach to Real-Time Scheduling. *Real-Time Systems*, 6(1), 1994.

[18] W. Zhao, K. Ramamritham, and J. Stankovic. Scheduling Tasks with Resource requirements in a Hard Real-Time System. *IEEE Transactions on Software Engineering*, SE-13(5):564–577, May 1987.