

A Framework for the Derivation of WCET Analyses for Multi-Core Processors

Michael Jacobs, Sebastian Hahn, Sebastian Hack

Saarland University, Saarbrücken, Germany

Email: jacobs@cs.uni-saarland.de, sebastian.hahn@cs.uni-saarland.de, hack@cs.uni-saarland.de

Abstract—Multi-core processors share common hardware resources between several processor cores. As a consequence, the performance of one processor core is influenced by the programs executed on the concurrent cores. We refer to this phenomenon as *shared-resource interference*. An explicit consideration of all such interference effects is in general combinatorially infeasible. This makes a precise worst-case execution time (WCET) analysis for multi-core processors challenging.

In order to reduce the complexity, WCET analyses for multi-core processors coarsely approximate the behavior of the considered applications. However, current approaches are only applicable to rather restricted classes of hardware platforms. We propose a framework for the derivation of WCET analyses for multi-core processors. It relaxes the restricting assumptions that existing approaches are based on.

The derivation starts from a WCET analysis that makes maximally pessimistic assumptions about the shared-resource interference. More precise interference bounds for the concrete system are subsequently lifted to the approximation of the analysis. The lifted bounds are finally incorporated in the analysis in order to model the interference in a more precise way.

I. INTRODUCTION

For a timing-critical application it is important that the time needed to deliver the results of its calculations does not exceed a deadline dictated by the physical environment. A timing-critical application may consist of several programs that interact. Knowledge about the worst-case execution time (WCET) [1] of each such program allows us to verify the timeliness of the overall application. It is safe to replace the WCET of a program by an upper bound on its execution times (a so-called WCET bound) in this verification step. However, the timeliness of an application can often only be verified if the WCET bounds are relatively tight. WCET analyses are used for the calculation of WCET bounds.

The execution times of a program depend on the possible execution behaviors at the micro-architectural level of the processor that executes the program. Modern processors are too complex to exhaustively simulate or measure the execution times of all possible behaviors. WCET analyses for those processors need to approximate some of the micro-architectural details in order to reduce the inherent complexity [2], [3]. Approximation often comes at the cost of a less tight WCET bound.

The use of multi-core processors can reduce the weight, the energy consumption and the production costs of computer systems. Hence, they are likely to also be used for timing-critical applications in the long run.

However, multi-core processors consist of several processor cores, which share common resources such as buses or caches. The resource sharing has a significant impact on the overall performance of a system [4] because the cores compete for the shared resources. For example, an access request to a shared bus may be blocked for some cycles before it is granted because a concurrent core is granted access first. This effect is commonly referred to as *shared-resource interference*.

The WCET analysis of programs executed on multi-core processors needs to take into account all possible interference effects due to

resource sharing. An exact consideration of all such effects requires in general an exhaustive enumeration of all possible interleavings of accesses to the shared resources by the different processor cores. Such an enumeration is combinatorially infeasible.

Most of the current approaches to WCET analysis for multi-core processors [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], [16] try to find a level of approximation that avoids this complexity without sacrificing precision too much. Unfortunately, they are restricted to Time-Division-Multiple-Access (TDMA) bus arbitration or not sound in the presence of indirect interference effects, which most modern multi-core platforms exhibit.

Contributions

We propose a framework for the derivation of WCET analyses for multi-core processors. An instance of our framework—derived according to the criteria proposed in this paper—is guaranteed to be a *sound* WCET analysis. The derivation starts from a baseline WCET analysis that makes maximally pessimistic assumptions about the shared-resource interference. We can infer more precise interference bounds from the specification of the concrete system. Lifting these bounds to the approximation of the baseline analysis avoids overly pessimistic assumptions about the interference.

Our iterative overapproximation analyzes each processor core on its own and still incorporates cumulative information about the concurrent cores in the lifted interference bounds. In this way, it finds a trade-off between the performance of analyzing each core in isolation and the precision of simultaneously considering all processor cores.

Our framework has been successfully used in the development of a novel analysis [17] that avoids the restrictions of the existing approaches.

II. RELATED WORK

Most approaches [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15] to WCET analysis or response time analysis for multi-core processors rely on compositionality [18] in the sense that they start with a timing analysis that ignores the shared-resource interference. Subsequently, they add bounds on the direct interference effects to their results. In modern micro-processors, however, the overall impact of the interference can exceed the direct interference effects [19]. Thus, these approaches are not applicable to current hardware platforms.

An approach by Chattopadhyay et al. [16] supports complex processor core pipelines. It is restricted to TDMA bus arbitration. Most multi-core processors on the market, however, implement event-driven bus arbitration protocols.

A recent approach by Kelter and Marwedel [20] supports complex multi-core processors equipped with event-driven bus arbitration. However, it relies on the enumeration of all interleavings of accesses to the shared bus by the different cores. Therefore, we expect it to not scale to realistic application scenarios.

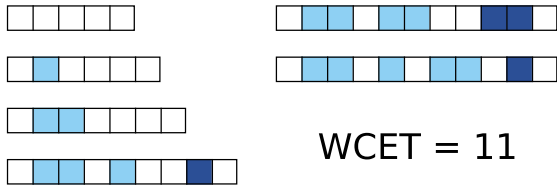


Figure 1: All six behaviors of an example program executed on a hardware platform. The program has a WCET of eleven time units.

A novel analysis developed by our group [17] overcomes the restrictions and simplifying assumptions of previous approaches. To the best of our knowledge, it is the first approach to the calculation of co-runner-sensitive WCET bounds that scales to multi-core processors with out-of-order execution and event-driven bus arbitration.

This paper presents the concepts we applied during the derivation of our novel analysis. They are embedded in a general and formally sound framework for the derivation of WCET analyses for multi-core processors.

III. MOTIVATION

We motivate the key principle of our framework by considering an example program executed on a hardware platform. Figure 1 shows all six possible execution behaviors of the program. Each sequence of boxes represents one execution behavior. White boxes stand for time units of non-interfered execution. The boxes colored in light blue represent *direct interference effects* like cycles blocked at a shared bus or needed to serve a miss in the shared cache. Dark boxes denote the prolonging effects of timing anomalies that are a consequence of earlier interference. A processor core pipeline might, for example, only speculate in a particular situation if it is blocked at the shared bus. If the prediction turns out as false, the execution time is prolonged by more than the blocked cycles. Such *indirect interference effects* can be observed in modern multi-core processors [19]. Sound WCET analyses for such platforms have to take these indirect effects into account.

The example program has a WCET of eleven time units as its longest execution behavior takes this long.

In this example, we assume that no execution behavior of the example program can exhibit more than five direct interference effects. Such assumptions can for instance be inferred from the specification of a bus arbiter if the actual set of execution behaviors is unknown.

A. Classical Compositional Analysis

Existing compositional analyses [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15] first calculate a basic timing bound that does not take into account behaviors exhibiting interference effects. Subsequently, they add an upper bound on the direct interference effects to the result. This principle is depicted in Figure 2. The longest behavior without interference takes five time units. The maximum of five direct interference effects is subsequently added.

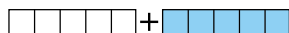


Figure 2: Compositional analyses typically only consider behaviors without interference. Subsequently, they add an upper bound on the direct interference effects. In the presence of indirect effects, this is unsound.

This shows that the common way of decomposing analyses is efficient but unsound in the presence of indirect interference effects. In order to be sound, an upper bound on the indirect effects has to be additionally incorporated.

The indirect interference effects highly depend on the interaction between the shared resources and the processor cores. As this interaction is typically not considered by compositional analyses, they can at best provide very pessimistic bounds on the indirect interference effects. In our example system, the amount of indirect interference effects cannot exceed the amount of direct ones. Note that this is hard or impossible to show for real-world hardware, for example due to domino effects [21]. The analysis results in a high overestimation of the WCET as shown in Figure 3.



Figure 3: In order to be sound, such a decomposition of timing analysis would have to pessimistically assume that each direct effect leads to an indirect effect. However, this results in a high overestimation.

As a consequence, the common way of decomposing timing analysis is not able to provide sound and precise WCET bounds for modern multi-core processors.

B. Key Principle of our Framework

The derivation of a WCET analysis in our framework starts from an overapproximation of all behaviors of the program. This overapproximation is maximally pessimistic with respect to the shared-resource interference. Each access request to a shared bus can, for example, be blocked arbitrarily long before being granted by the arbiter. Similarly, each access to a shared cache can be a hit or a miss. Figure 4 shows such an overapproximation for our example program. It contains two infeasible behaviors that cannot be observed when actually executing the program (dashed box).

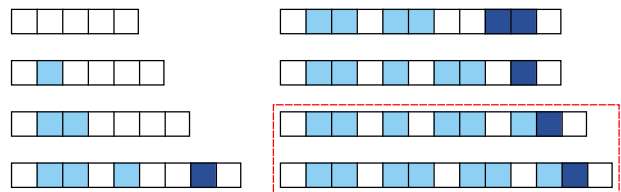


Figure 4: The pessimistic overapproximation of all behaviors of the example program contains two infeasible behaviors (dashed box).

We exploit the upper bound on the direct interference effects to prune the two infeasible behaviors, which exhibit more than five direct interference effects. The remaining execution behaviors result in an exact WCET bound of eleven time units.

In this way, our framework supports the development of timing analyses that explicitly model the impact of the interference on the processor cores and, thus, precisely bound the indirect interference effects.

Note that the pessimistic overapproximation is only a conceptual starting point. The implementation of the analysis derived from our framework will not materialize this overapproximation. Instead, it can directly leave out many of the infeasible behaviors during analysis.

Our example program only exhibits six execution behaviors. Programs executed on real-world hardware platforms, however, exhibit too many execution behaviors to exhaustively enumerate them. Thus, it is common to approximate some of the micro-architectural details [2]. In the next section, we formally show how the principle presented above can be applied to such approximations.

IV. PROPERTY LIFTING

This section describes the concept of property lifting, which is the central part of our framework.

A. Concrete Execution Behavior and Time

We consider a multi-core processor consisting of the set $Cores$ of n processor cores.

$$Cores = \{C_1, \dots, C_n\}$$

For simplicity, we assume that each core only runs one program and that each program may at most be executed once per system run. This restriction is not inherent to our framework but only made to simplify the notation. For a more detailed discussion on how to overcome this restriction, we refer to [22]. In the following, we use the term *system* to refer to the combination of the hardware including the multi-core processor and the software executed on it.

The system may exhibit different execution behaviors depending on its initial state, external input parameters and clock drift effects. Let $Traces$ be the set of all execution behaviors of the system. Its superset $Universe$ additionally contains the spurious behaviors that might be described by imprecise analyses. Spurious behaviors can, for example, be sequences of concrete system states that cannot be observed during any execution of the concrete system.

$$Universe \supseteq Traces$$

The program executed on processor core C_i can be assigned an execution time per execution behavior. This time is given by the function et_{C_i} .

$$et_{C_i} : Universe \rightarrow \mathbb{N} \cup \{\infty\}$$

The WCET of the program executed on core C_i is its maximal execution time over all execution behaviors of the considered system.

$$WCET_{C_i} = \max_{t \in Traces} et_{C_i}(t) \quad (1)$$

B. Approximation by Abstract Traces

Modern processors usually exhibit too many execution behaviors to allow for an exhaustive consideration of all of them. The set $Traces$ is simply too large. Therefore, it is mandatory to introduce some kind of approximation. The goal is to not have to argue separately about each concrete execution behavior.

In our view, an abstract model of the considered system is given by the tuple $(\widehat{Traces}, \gamma_{trace})$. \widehat{Traces} is the set of abstract traces of the model. Depending on the chosen way of approximation, an abstract trace might for example be a sequence of abstract states in an analysis based on abstract interpretation [23] or the combination of a sequence of superblocks [5] and a corresponding sequence of blocking cycle counts. The function γ_{trace} maps those abstract traces to subsets of the universe of execution behaviors. Note that $\mathcal{P}(Universe)$ denotes the power set of this universe of execution behaviors.

$$\gamma_{trace} : \widehat{Traces} \rightarrow \mathcal{P}(Universe)$$

We say that an abstract model $(\widehat{Traces}, \gamma_{trace})$ is an overapproximation of $Traces$ iff:

$$\bigcup_{\hat{t} \in \widehat{Traces}} \gamma_{trace}(\hat{t}) \supseteq Traces \quad (2)$$

We assume that for each core C_i there is an upper bound on its execution times per abstract trace. This bound shall be given by $UB_{et_{C_i}}$.

$$\begin{aligned} & UB_{et_{C_i}} : \widehat{Traces} \rightarrow \mathbb{N} \cup \{\infty\} \\ \forall \hat{t} \in \widehat{Traces} : UB_{et_{C_i}}(\hat{t}) & \geq \max_{t \in \gamma_{trace}(\hat{t})} et_{C_i}(t) \end{aligned} \quad (3)$$

From (2) and (3) it follows that the abstract model provides an upper bound to the WCET as defined in (1) by:

$$\max_{\hat{t} \in \widehat{Traces}} UB_{et_{C_i}}(\hat{t}) \geq WCET_{C_i} \quad (4)$$

From now on we only consider abstract models that are overapproximations of $Traces$.

C. Infeasible Abstract Traces

The method used to obtain the set of abstract traces (e.g. a static analysis exploring abstract states) might introduce imprecision. Therefore, there may be abstract traces that do not describe any execution behavior of the considered system. We call them *infeasible* abstract traces.

$$\widehat{Infeas} = \{\hat{t} \in \widehat{Traces} \mid \gamma_{trace}(\hat{t}) \cap Traces = \emptyset\} \quad (5)$$

Correspondingly, we refer to $\widehat{Traces} \setminus \widehat{Infeas}$ as the set of *feasible* abstract traces. In fact, it follows from (5) that the set of feasible abstract traces is an overapproximation of $Traces$.

$$\bigcup_{\hat{t} \in \widehat{Traces} \setminus \widehat{Infeas}} \gamma_{trace}(\hat{t}) \supseteq Traces \quad (6)$$

Based on an abstract model $(\widehat{Traces}, \gamma_{trace})$, which is an overapproximation of $Traces$, we define a set $Deriv_{(\widehat{Traces}, \gamma_{trace})}$ of further abstract models as follows:

$$\begin{aligned} Deriv_{(\widehat{Traces}, \gamma_{trace})} = \\ \{(\widehat{Traces}', \gamma_{trace}') \mid \widehat{Traces} \supseteq \widehat{Traces}' \supseteq \widehat{Traces} \setminus \widehat{Infeas}\} \end{aligned} \quad (7)$$

Intuitively, each element of $Deriv_{(\widehat{Traces}, \gamma_{trace})}$ is an overapproximation of $Traces$. So we can calculate an upper bound to the WCET based on any member of $Deriv_{(\widehat{Traces}, \gamma_{trace})}$:

$$\begin{aligned} \forall (\widehat{Traces}', \gamma_{trace}') \in Deriv_{(\widehat{Traces}, \gamma_{trace})} : \\ \max_{\hat{t} \in \widehat{Traces}'} UB_{et_{C_i}}(\hat{t}) \geq WCET_{C_i} \end{aligned} \quad (8)$$

As a consequence, we can ignore an arbitrarily chosen set of infeasible abstract traces in an abstract model. A WCET bound based on the remaining abstract traces is still guaranteed to be sound.

The calculation of WCET bounds is based on upper bounds on the execution times per abstract trace (3). If an abstract model makes conservative assumptions about the behavior at the shared resources, some infeasible abstract traces might assume an amount of shared-resource interference that exceeds the maximum possible amount for the concrete system. As upper bounds on the execution times of such infeasible abstract traces are likely to be very pessimistic, ignoring those abstract traces—as in (8)—might improve the tightness of the resulting WCET bound.

However, it depends heavily on the particular abstract model $(\widehat{Traces}, \gamma_{trace})$ and the upper bounds on the execution times per abstract trace whether the WCET bound can be tightened by leaving out some infeasible abstract traces.

We introduced the abstract model to not have to materialize the set $Traces$. The definition of infeasible abstract traces, however, is also based on $Traces$. Therefore, we cannot directly use this definition to detect infeasible abstract traces. The following subsection describes how we can use properties of the system under consideration to detect some infeasible abstract traces.

assume a compound abstract model as baseline. It shall be composed of one abstract model per processor core.

$$\text{Models} = \text{Cores} \quad (i)$$

Further assume that each abstract model can only provide detailed information about the processor core it is specialized on. In particular, this means:

$$\begin{aligned} \forall C_i \in \text{Cores} : \\ \forall \widehat{t}^{C_i} \in \widehat{\text{Traces}}^{C_i} : \\ \forall C_j \in (\text{Cores} \setminus \{C_i\}) : \\ \quad {}^{LB}\#\text{blockedCycles}_{C_j}(\widehat{t}^{C_i}) = 0 \wedge \\ \quad {}^{UB}\#\text{accessCycles}_{C_j}(\widehat{t}^{C_i}) = \infty \end{aligned} \quad (j)$$

In combination with (21) and (22) this implies the following equalities:

$$\begin{aligned} \forall C_i \in \text{Cores} : \\ {}^{LB}\#\text{blockedCycles}_{C_i}(\widehat{t}) = {}^{LB}\#\text{blockedCycles}_{C_i}(\pi_{\text{trace}}^{C_i}(\widehat{t})) \wedge \\ {}^{UB}\#\text{accessCycles}_{C_i}(\widehat{t}) = {}^{UB}\#\text{accessCycles}_{C_i}(\pi_{\text{trace}}^{C_i}(\widehat{t})) \end{aligned} \quad (k)$$

This allows us to rewrite the lifted property \widehat{P}_{wc} as follows:

$$\begin{aligned} \forall \widehat{t} \in \widehat{\text{Traces}} : \\ \widehat{P}_{wc}(\widehat{t}) \\ \stackrel{(g)}{\Leftrightarrow} [{}^{LB}\#\text{blockedCycles}_{C_i}(\widehat{t}) \\ \leq \sum_{C_j \in (\text{Cores} \setminus \{C_i\})} {}^{UB}\#\text{accessCycles}_{C_j}(\widehat{t})] \\ \stackrel{(k)}{\Leftrightarrow} [{}^{LB}\#\text{blockedCycles}_{C_i}(\pi_{\text{trace}}^{C_i}(\widehat{t})) \\ \leq \sum_{C_j \in (\text{Cores} \setminus \{C_i\})} {}^{UB}\#\text{accessCycles}_{C_j}(\pi_{\text{trace}}^{C_j}(\widehat{t}))] \end{aligned} \quad (m)$$

This time, the lifted property \widehat{P}_{wc} is not guaranteed to hold for all abstract traces. Hence, it can effectively detect infeasible abstract traces. [Example end]

However, the cross product in the definition of $\widehat{\text{Traces}}$ already gives a hint that $\widehat{\text{Traces}}$ might become quite large. Thus, the compound consideration of several abstract models is likely too complex in most cases.

C. Projections of the Compound Results

Taking a closer look at the set $\widehat{\text{LessTraces}}$ derived from the compound abstract model, it turns out that we are not really interested in the set of all combinations of abstract traces from the different abstract models. It would be sufficient to know for each $M_a \in \text{Models}$ which members of $\widehat{\text{Traces}}^{M_a}$ are contained in a compound abstract trace of $\widehat{\text{LessTraces}}$. Those subsets can be obtained by projecting the members of $\widehat{\text{LessTraces}}$ to their different components. We define the projections in a general way on arbitrary subsets $\widehat{\text{SomeTraces}}$ of $\widehat{\text{Traces}}$.

$$\begin{aligned} \forall M_a \in \text{Models} : \\ \pi^{M_a}(\widehat{\text{SomeTraces}}) = \{\pi_{\text{trace}}^{M_a}(\widehat{t}) \mid \widehat{t} \in \widehat{\text{SomeTraces}}\} \end{aligned} \quad (23)$$

Obviously, each projection $\pi^{M_a}(\widehat{\text{SomeTraces}})$ is a subset of the set of abstract traces of the corresponding abstract model.

$$\begin{aligned} \forall \widehat{\text{SomeTraces}} \subseteq \widehat{\text{Traces}} : \\ \forall M_a \in \text{Models} : \\ \pi^{M_a}(\widehat{\text{SomeTraces}}) \subseteq \widehat{\text{Traces}}^{M_a} \end{aligned} \quad (24)$$

Please note that $\widehat{\text{SomeTraces}}$ is a subset of the cross product over its projections.

$$\begin{aligned} \forall \widehat{\text{SomeTraces}} \subseteq \widehat{\text{Traces}} : \\ \widehat{\text{SomeTraces}} \\ \subseteq \pi^{M_1}(\widehat{\text{SomeTraces}}) \times \dots \times \pi^{M_m}(\widehat{\text{SomeTraces}}) \end{aligned} \quad (25)$$

Furthermore, it is rather obvious that the projection functions π^{M_a} are monotone.

$$\begin{aligned} \forall \widehat{\text{SomeTraces}}, \widehat{\text{OtherTraces}} \subseteq \widehat{\text{Traces}} : \\ \forall M_a \in \text{Models} : \\ [\widehat{\text{SomeTraces}} \subseteq \widehat{\text{OtherTraces}}] \\ \Rightarrow [\pi^{M_a}(\widehat{\text{SomeTraces}}) \subseteq \pi^{M_a}(\widehat{\text{OtherTraces}})] \end{aligned} \quad (26)$$

Each projection $\pi^{M_a}(\widehat{\text{LessTraces}})$ is a superset of the feasible abstract traces of the corresponding $\widehat{\text{Traces}}^{M_a}$. Consider (27) for a formal proof of this statement. According to (7), (24) and (27), each abstract model $(\pi^{M_a}(\widehat{\text{LessTraces}}), \gamma_{\text{trace}}^{M_a})$ is a member of $\text{Deriv}_{(\widehat{\text{Traces}}^{M_a}, \gamma_{\text{trace}}^{M_a})}$.

$$\begin{aligned} \forall M_a \in \text{Models} : \\ (\pi^{M_a}(\widehat{\text{LessTraces}}), \gamma_{\text{trace}}^{M_a}) \in \text{Deriv}_{(\widehat{\text{Traces}}^{M_a}, \gamma_{\text{trace}}^{M_a})} \end{aligned} \quad (28)$$

Thus, each abstract model $(\pi^{M_a}(\widehat{\text{LessTraces}}), \gamma_{\text{trace}}^{M_a})$ can be used to calculate a WCET bound based on it.

$$\begin{aligned} \forall M_a \in \text{Models} : \\ \max_{\widehat{t}^{M_a} \in \pi^{M_a}(\widehat{\text{LessTraces}})} {}^{UB}et_{C_i}(\widehat{t}^{M_a}) \geq \text{WCET}_{C_i} \end{aligned} \quad (29)$$

But how precise is a WCET bound based on the projection of $\widehat{\text{LessTraces}}$ compared to one that is directly based on $\widehat{\text{LessTraces}}$? In general, we might lose precision by restricting ourselves to the projections $\pi^{M_a}(\widehat{\text{LessTraces}})$.

$$\begin{aligned} \text{WCET}_{C_i} \\ \stackrel{(15)}{\leq} \max_{\widehat{t} \in \widehat{\text{LessTraces}}} {}^{UB}et_{C_i}(\widehat{t}) \\ \stackrel{(22)}{=} \max_{\widehat{t} \in \widehat{\text{LessTraces}}} \min_{M_a \in \text{Models}} {}^{UB}et_{C_i}(\pi_{\text{trace}}^{M_a}(\widehat{t})) \\ \stackrel{(31)}{\leq} \min_{M_a \in \text{Models}} \max_{\widehat{t} \in \widehat{\text{LessTraces}}} {}^{UB}et_{C_i}(\pi_{\text{trace}}^{M_a}(\widehat{t})) \\ \stackrel{(23)}{=} \min_{M_a \in \text{Models}} \max_{\widehat{t}^{M_a} \in \pi^{M_a}(\widehat{\text{LessTraces}})} {}^{UB}et_{C_i}(\widehat{t}^{M_a}) \end{aligned} \quad (30)$$

We can prove the second \leq relation used in (30) by assuming its opposite and deriving a statement from it that contradicts to the definition of the minimum.

$$\begin{aligned} \max_{\widehat{t} \in \widehat{\text{LessTraces}}} \min_{M_a \in \text{Models}} {}^{UB}et_{C_i}(\pi_{\text{trace}}^{M_a}(\widehat{t})) \\ > \min_{M_a \in \text{Models}} \max_{\widehat{t} \in \widehat{\text{LessTraces}}} {}^{UB}et_{C_i}(\pi_{\text{trace}}^{M_a}(\widehat{t})) \\ \Rightarrow \exists \widehat{t}^* \in \widehat{\text{LessTraces}} : \exists M_* \in \text{Models} : \\ \min_{M_a \in \text{Models}} {}^{UB}et_{C_i}(\pi_{\text{trace}}^{M_a}(\widehat{t}^*)) \\ > \max_{\widehat{t} \in \widehat{\text{LessTraces}}} {}^{UB}et_{C_i}(\pi_{\text{trace}}^{M_*}(\widehat{t})) \\ \stackrel{\text{def.}}{\geq} {}^{UB}et_{C_i}(\pi_{\text{trace}}^{M_*}(\widehat{t}^*)) \quad \downarrow \\ \max \end{aligned} \quad (31)$$

However, we additionally assume each abstract model is focused on one processor core.

$$\text{Models} = \text{Cores} \quad (32)$$

The boolean predicate \widehat{P}^{M_a} used in F^{M_a} takes an abstract trace from \widehat{Traces}^{M_a} and the current vector of approximation variables as parameters. It is defined as follows.

$$\begin{aligned} \forall M_a \in Models : \\ \forall \widehat{t}^{M_a} \in \widehat{Traces}^{M_a} : \\ \widehat{P}^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx}) \\ \Leftrightarrow \forall P_k \in Prop : \widehat{P}_k^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx}) \end{aligned} \quad (40)$$

The $\widehat{P}_k^{M_a}$ are properties that overapproximate the \widehat{P}_k lifted to the compound abstract model. They shall fulfill the following criterion with respect to the \widehat{P}_k .

Soundness Criterion (C2):

$$\begin{aligned} \forall \overrightarrow{Approx} \in \mathcal{P}(\widehat{Traces}^{M_1}) \times \dots \times \mathcal{P}(\widehat{Traces}^{M_m}) : \\ \forall \widehat{t}^{M_a} \in \widehat{Traces}^{M_a} : \\ [\exists (t^{M_1}, \dots, t^{M_{a-1}}) \in \widehat{Approx}^{M_1} \times \dots \times \widehat{Approx}^{M_{a-1}} : \\ \exists (t^{M_{a+1}}, \dots, t^{M_m}) \in \widehat{Approx}^{M_{a+1}} \times \dots \times \widehat{Approx}^{M_m} : \\ \widehat{P}_k(t^{M_1}, \dots, t^{M_{a-1}}, \widehat{t}^{M_a}, t^{M_{a+1}}, \dots, t^{M_m})] \\ \Rightarrow \widehat{P}_k^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx}) \end{aligned} \quad (C2)$$

Criterion (C2) allows us to show that each approximation variable is guaranteed to be an overapproximation of the corresponding projection after arbitrary sequences of updates of the approximation variables:

$$\begin{aligned} \forall M_a \in Models : \\ \pi^{M_a}(\widehat{LessTraces}) \subseteq \widehat{Approx}^{M_a} \end{aligned} \quad (H1)$$

Proof: As a consequence of (24), the claim in (H1) trivially holds for the initial values of the approximation variables as specified in (36). For the general case, however, we assume the hypothesis (H1) to hold after a given sequence of approximation variable updates. In an inductive way, we can use this assumption to show that the hypothesis is preserved by an additional simultaneous update of an arbitrarily chosen set of the approximation variables. For the details of this induction step, please refer to (41) and (42). The inequation chain in (41) shows that all sets contained in it are in fact equal. This information is used in (42) to show that the $F^{M_a}(\overrightarrow{Approx})$ are guaranteed to be supersets of the projections $\pi^{M_a}(\widehat{LessTraces})$. According to the equation system given by (37), the approximation variables \widehat{Approx}^{M_a} are updated to the values of the functions $F^{M_a}(\overrightarrow{Approx})$. This proves that the simultaneous update of an arbitrarily chosen set of approximation variables is guaranteed to preserve the hypothesis given by (H1). \square

As a consequence of (42), we can optionally use the alternative definitions of the update functions F^{M_a} given in (43). Their use is equivalent to the use of the definitions in (37). Depending on the implementation details of an instance of our framework, it could be more straightforward to use one style of definition or the other.

$$\begin{aligned} \forall M_a \in Models : \\ F^{M_a}(\overrightarrow{Approx}) \\ := \{ \widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Approx}^{M_a} \wedge \widehat{P}_k^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx}) \} \end{aligned} \quad (43)$$

It follows from hypothesis (H1) that we can bound the content of the sets \widehat{Approx}^{M_a} from above and from below.

$$\begin{aligned} \widehat{Traces}^{M_a} \\ \supseteq \widehat{Approx}^{M_a} \\ \stackrel{(38)}{\supseteq} \pi^{M_a}(\widehat{LessTraces}) \\ \stackrel{(H1)}{\supseteq} \pi^{M_a}(\widehat{LessTraces}) \\ \stackrel{(27)}{\supseteq} \widehat{Traces}^{M_a} \setminus \widehat{Infeas}^{M_a} \end{aligned} \quad (44)$$

As a consequence of (7) and (44), we see that each abstract model $(\widehat{Approx}^{M_a}, \gamma_{trace}^{M_a})$ is a member of $Deriv_{(\widehat{Traces}^{M_a}, \gamma_{trace}^{M_a})}$.

$$\begin{aligned} \forall M_a \in Models : \\ (\widehat{Approx}^{M_a}, \gamma_{trace}^{M_a}) \in Deriv_{(\widehat{Traces}^{M_a}, \gamma_{trace}^{M_a})} \end{aligned} \quad (45)$$

Thus, each abstract model $(\widehat{Approx}^{M_a}, \gamma_{trace}^{M_a})$ can be used to calculate a WCET bound based on it.

$$\begin{aligned} \forall M_a \in Models : \\ \max_{\widehat{t}^{M_a} \in \widehat{Approx}^{M_a}}^{UB} etc_i(\widehat{t}^{M_a}) \geq WCET_{C_i} \end{aligned} \quad (46)$$

In addition, the $\widehat{P}_k^{M_a}$ shall fulfill the following criterion.

Monotonicity Criterion (C3):

$$\begin{aligned} \forall \overrightarrow{Approx}, \overrightarrow{Approx}' \in \mathcal{P}(\widehat{Traces}^{M_1}) \times \dots \times \mathcal{P}(\widehat{Traces}^{M_m}) : \\ \forall \widehat{t}^{M_a} \in \widehat{Traces}^{M_a} : \\ [\forall M_b \in Models : \\ \widehat{Approx}'^{M_b} \subseteq \widehat{Approx}^{M_b}] \\ \Rightarrow [\widehat{P}_k^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx}') \Rightarrow \widehat{P}_k^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx})] \end{aligned} \quad (C3)$$

Let \overrightarrow{Approx}' be the vector of approximation variables after an arbitrary sequence of updates of the approximation variables. Criterion (C3) allows us to show that the following additional hypothesis holds:

$$\begin{aligned} \forall M_a \in Models : \\ F^{M_a}(\overrightarrow{Approx}) \subseteq \widehat{Approx}^{M_a} \end{aligned} \quad (H2)$$

Proof: According to (36) and (37), the hypothesis (H2) trivially holds for a vector \overrightarrow{Approx} just initialized. For the inductive step, assume that hypothesis (H2) holds for a given vector \overrightarrow{Approx} of approximation variables. Let \overrightarrow{Approx}' be the successor of \overrightarrow{Approx} after the simultaneous update of an arbitrarily chosen set of approximation variables:

$$\overrightarrow{Approx}' = (\widehat{Approx}'^{M_1}, \dots, \widehat{Approx}'^{M_m}) \quad (47)$$

$$\begin{aligned} \forall M_a \in Models : \\ \widehat{Approx}'^{M_a} \in \{ \widehat{Approx}^{M_a}, F^{M_a}(\overrightarrow{Approx}) \} \end{aligned} \quad (48)$$

It follows from (H2) and (48) that each component of \overrightarrow{Approx}' is a subset of its corresponding counterpart in \overrightarrow{Approx} .

$$\begin{aligned} \forall M_a \in Models : \\ \widehat{Approx}'^{M_a} \subseteq \widehat{Approx}^{M_a} \end{aligned} \quad (49)$$

$$\begin{aligned}
& \pi^{M_a}(\widehat{LessTraces}) \\
& \stackrel{(23)}{=} \{\pi_{trace}^{M_a}(\hat{t}) \mid \hat{t} \in \widehat{LessTraces}\} \\
& \stackrel{(13)}{=} \{\pi_{trace}^{M_a}(\hat{t}) \mid \hat{t} \in \widehat{LessTraces} \wedge \widehat{P}(\hat{t})\} \\
& \stackrel{(25)}{\subseteq} \{\pi_{trace}^{M_a}(\hat{t}) \mid \hat{t} \in \pi^{M_1}(\widehat{LessTraces}) \times \dots \times \pi^{M_m}(\widehat{LessTraces}) \wedge \widehat{P}(\hat{t})\} \\
& \stackrel{(H1)}{\subseteq} \{\pi_{trace}^{M_a}(\hat{t}) \mid \hat{t} \in \widehat{Approx}^{M_1} \times \dots \times \widehat{Approx}^{M_m} \wedge \widehat{P}(\hat{t})\} \\
& \stackrel{(38)}{\subseteq} \{\pi_{trace}^{M_a}(\hat{t}) \mid \hat{t} \in \widehat{Traces}^{M_1} \times \dots \times \widehat{Traces}^{M_m} \wedge \widehat{P}(\hat{t})\} \\
& \stackrel{(18)}{=} \{\pi_{trace}^{M_a}(\hat{t}) \mid \hat{t} \in \widehat{Traces} \wedge \widehat{P}(\hat{t})\} \\
& \stackrel{(13)}{=} \{\pi_{trace}^{M_a}(\hat{t}) \mid \hat{t} \in \widehat{LessTraces}\}
\end{aligned} \tag{41}$$

$$\begin{aligned}
& \pi^{M_a}(\widehat{LessTraces}) \\
& \stackrel{(41)}{=} \{\pi_{trace}^{M_a}(\hat{t}) \mid \hat{t} \in \widehat{Approx}^{M_1} \times \dots \times \widehat{Approx}^{M_m} \wedge \widehat{P}(\hat{t})\} \\
& \stackrel{(19)}{=} \{\widehat{t}^{M_a} \mid (\widehat{t}^{M_1}, \dots, \widehat{t}^{M_a}, \dots, \widehat{t}^{M_m}) \in \widehat{Approx}^{M_1} \times \dots \times \widehat{Approx}^{M_m} \wedge \widehat{P}(\widehat{t}^{M_1}, \dots, \widehat{t}^{M_a}, \dots, \widehat{t}^{M_m})\} \\
& \stackrel{(11)}{=} \{\widehat{t}^{M_a} \mid (\widehat{t}^{M_1}, \dots, \widehat{t}^{M_a}, \dots, \widehat{t}^{M_m}) \in \widehat{Approx}^{M_1} \times \dots \times \widehat{Approx}^{M_m} \wedge \forall P_k \in Prop : \widehat{P}_k(\widehat{t}^{M_1}, \dots, \widehat{t}^{M_a}, \dots, \widehat{t}^{M_m})\} \\
& \stackrel{(C2)}{\subseteq} \{\widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Approx}^{M_a} \wedge \forall P_k \in Prop : \widehat{P}_k^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx})\} \\
& \stackrel{(40)}{=} \{\widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Approx}^{M_a} \wedge \widehat{P}^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx})\} \\
& \stackrel{(38)}{\subseteq} \{\widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Traces}^{M_a} \wedge \widehat{P}^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx})\} \\
& \stackrel{(37)}{=} F^{M_a}(\overrightarrow{Approx})
\end{aligned} \tag{42}$$

We assume that the update functions F^{M_a} can be applied to $\overrightarrow{Approx'}$ in the same way as to \overrightarrow{Approx} . Equation (50) shows that $F^{M_a}(\overrightarrow{Approx'})$ is a subset of $F^{M_a}(\overrightarrow{Approx})$.

$$\begin{aligned}
& F^{M_a}(\overrightarrow{Approx'}) \\
& \stackrel{(37)}{=} \{\widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Traces}^{M_a} \wedge \widehat{P}^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx'})\} \\
& \stackrel{(40)}{=} \{\widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Traces}^{M_a} \\
& \quad \wedge \forall P_k \in Prop : \widehat{P}_k^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx'})\} \\
& \stackrel{(C3)}{\subseteq} \{\widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Traces}^{M_a} \\
& \quad \wedge \forall P_k \in Prop : \widehat{P}_k^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx})\} \\
& \stackrel{(40)}{=} \{\widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Traces}^{M_a} \wedge \widehat{P}^{M_a}(\widehat{t}^{M_a}, \overrightarrow{Approx})\} \\
& \stackrel{(37)}{=} F^{M_a}(\overrightarrow{Approx})
\end{aligned} \tag{50}$$

Based on those results, it is straightforward to show that the hypothesis also holds for $\overrightarrow{Approx'}$, which concludes the inductive

proof of (H2):

$$\begin{aligned}
& \widehat{Approx'}^{M_a} \\
& \stackrel{(48)}{=} \left\{ \begin{array}{l} \widehat{Approx}^{M_a} \\ F^{M_a}(\overrightarrow{Approx}) \end{array} \right. \\
& \stackrel{(H2)}{\supseteq} F^{M_a}(\overrightarrow{Approx}) \\
& \stackrel{(50)}{\supseteq} F^{M_a}(\overrightarrow{Approx'})
\end{aligned} \tag{51}$$

□

The intuition behind (H2) is that the update of an approximation variable is guaranteed to never increase its set of abstract traces. As the calculation of the WCET bounds is based on the abstract trace sets, we can be sure that the update of some approximation variables can never result in worse WCET bounds.

Example: Coming back to the example of Sections V-A and V-B, we can further lift the property \widehat{P}_{wc} —as defined in equation (m)—in a way that satisfies criteria (C2) and (C3):

$$\begin{aligned}
& \widehat{P}_{wc}^{C_i}(\widehat{t}^{C_i}, (\widehat{Approx}^{C_1}, \dots, \widehat{Approx}^{C_n})) \\
& \Leftrightarrow [{}^{LB}\#blockedCycles_{C_i}(\widehat{t}^{C_i}) \leq \\
& \quad \sum_{C_j \in (Cores \setminus \{C_i\})} \max_{\widehat{t}^{C_j} \in \widehat{Approx}^{C_j}} {}^{UB}\#accessCycles_{C_j}(\widehat{t}^{C_j})]
\end{aligned} \tag{n}$$

Note that the right-hand side of the inequation in property $\widetilde{P}_{wc}^{C_i}$ does not depend on the abstract trace t^{C_i} . It contains cumulative information about the processor cores competing against C_i . Thus, this right-hand side is constant over all evaluations of $\widetilde{P}_{wc}^{C_i}$ during an update of \widetilde{Approx}^{C_i} . Therefore, we can precompute the constant right-hand side based on the other approximation variables before updating \widetilde{Approx}^{C_i} . The constant right-hand side can subsequently be used in an integer linear programming constraint. [17] [Example end]

Moreover, we can use hypothesis (H2) to show that all sets contained in the following cyclic inequation chain are equal.

$$\begin{aligned}
& F^{M_a}(\widetilde{Approx}) \\
& \stackrel{(37)}{=} \{ \widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in F^{M_a}(\widetilde{Approx}) \wedge \widehat{P}^{M_a}(\widehat{t}^{M_a}, \widetilde{Approx}) \} \\
& \stackrel{(H2)}{\subseteq} \{ \widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Approx}^{M_a} \wedge \widehat{P}^{M_a}(\widehat{t}^{M_a}, \widetilde{Approx}) \} \\
& \stackrel{(38)}{\subseteq} \{ \widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Traces}^{M_a} \wedge \widehat{P}^{M_a}(\widehat{t}^{M_a}, \widetilde{Approx}) \} \\
& \stackrel{(37)}{=} F^{M_a}(\widetilde{Approx})
\end{aligned} \tag{52}$$

An interesting consequence of (52) is that, in particular, the following equation holds.

$$\begin{aligned}
& \{ \widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Approx}^{M_a} \wedge \widehat{P}^{M_a}(\widehat{t}^{M_a}, \widetilde{Approx}) \} \\
& \stackrel{(52)}{=} \{ \widehat{t}^{M_a} \mid \widehat{t}^{M_a} \in \widehat{Traces}^{M_a} \wedge \widehat{P}^{M_a}(\widehat{t}^{M_a}, \widetilde{Approx}) \}
\end{aligned} \tag{53}$$

Using (53) in the induction step (42) of the soundness proof instead of (38) allows us to replace the last \subseteq by an equal. With this improvement in place, we see that (42) has only left a single \subseteq . This means, we can exactly point out where the update of an approximation variable may lose precision compared to the projections of the compound consideration of all models at once. The compound consideration of all models only keeps those \widehat{t}^{M_a} that occur in a combination with abstract traces from the other models that fulfills all lifted properties \widehat{P}_k . However, the specially lifted properties $\widehat{P}_k^{M_a}$ may lead to different results. In their presence, an abstract trace \widehat{t}^{M_a} is not pruned as soon as for each \widehat{P}_k there is a particular combination with abstract traces from the other models that fulfills \widehat{P}_k —this is a consequence of (C2). The intersection of those sets of combinations of abstract traces for the different \widehat{P}_k could possibly be empty for a particular \widehat{t}^{M_a} . In that case, this \widehat{t}^{M_a} is not contained in $\pi^{M_a}(\widehat{LessTraces})$, but in its overapproximation \widetilde{Approx}^{M_a} .

Please note that the additional abstract traces introduced in this way by overapproximation are inherent to considering each abstract model on its own. They can occur even if we choose the $\widehat{P}_k^{M_a}$ in a way that the \Rightarrow in criterion (C2) can be shown to be replaceable by a \Leftrightarrow . This inherent amount of overapproximation only depends on the abstract models $(\widehat{Traces}^{M_a}, \gamma_{trace}^{M_a})$ and the way in which the \widehat{P}_k are chosen. Of course, it might lead to further overapproximation if we choose the $\widehat{P}_k^{M_a}$ in a way that we cannot prove the additional equivalence relation.

VI. ADVANTAGES OF THE FRAMEWORK

This section highlights the benefits of using our framework.

Standard derivation procedure: The framework is a common starting point for the derivation of future WCET analyses for multi-core processors. It has been successfully used in the development of a novel analysis that avoids the restrictions of previous approaches (cf. Section II).

Soundness guarantee: We show in this paper that an instance of our framework is a sound WCET analysis. This soundness is a consequence of a sound baseline analysis and the property lifting according to the criteria presented above. A sound baseline analysis is easily obtained by adapting a single-core WCET analysis in a way that makes it maximally pessimistic with respect to the shared-resource interference [17]. Hence, our framework essentially reduces the soundness proofs of its instances to showing the soundness of the property lifting steps involved in their derivations.

Assumptions about the system always explicit: The declarative style of our framework makes it mandatory to explicitly list all properties that a derived analysis assumes about the system under analysis. This makes sure that a derived analysis does not rely on implicit assumptions (except those that its baseline analysis already relies on).

Clean separation between concrete system and approximation: Existing analyses often try to directly incorporate properties of the concrete system in their level of approximation. However, this is mostly based on intuition and, thus, very error-prone. The concept of property lifting, in contrast, provides a clean separation between system properties and their implications on the approximation.

Trade-off between efficiency and precision: The iterative overapproximation (Section V-D) forms a trade-off between the efficiency of analyzing the programs of one processor core in isolation (Section V-A) and the precision of performing a simultaneous consideration of the detailed behaviors of the programs executed on all processor cores (Section V-B).

Not limited to multi-core processors: The principles presented throughout this paper are not limited to the analysis of multi-core processors. Some of the techniques used in single-core WCET analysis can also be seen as instances of our framework. The micro-architectural analysis, for example, typically has no notion of loop bounds. Thus, it pessimistically assumes that each loop body in the program can be executed indefinitely. Loop bounds of the concrete program are subsequently lifted to the level of approximation that the path analysis operates on. The lifted loop bounds are typically implemented as additional constraints in an implicit path enumeration [25].

VII. FRAMEWORK INSTANTIATION WORKFLOW

Figure 5 sketches the typical workflow of deriving WCET analyses as instances of our framework. A derivation that only relies on the concept of property lifting (cf. Section IV) comprises two logical steps. The derivation of an analysis that iteratively overapproximates the results of properties lifted to a compound abstract model (cf. Section V) requires an additional lifting step.

We successfully used our framework for the derivation of two novel WCET analyses for multi-core processors with a shared bus and Round-Robin bus arbitration: a co-runner-insensitive analysis and a co-runner-sensitive one [17]. This section describes—at a high level—how the derivation of each of these analyses follows the instantiation workflow sketched in Figure 5.

The derivation of our *co-runner-insensitive* analysis comprises two steps:

Step 1: The derivation starts from a baseline analysis focusing on one core and assuming that each access request to the shared bus can be blocked indefinitely by the bus arbiter. Furthermore, we consider a system property that bounds the maximum amount of blocked cycles per access to the shared bus under Round-Robin arbitration.

Step 2: The Round-Robin property is lifted to the baseline analysis. The lifted property is subsequently added to the implementation of

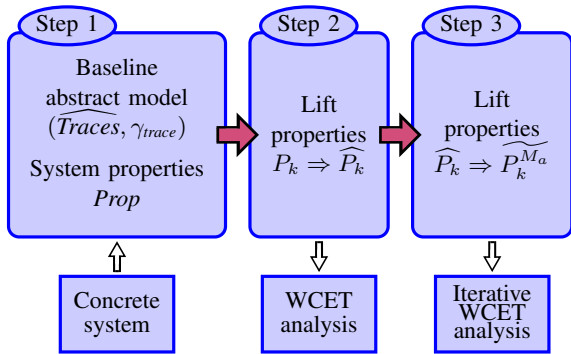


Figure 5: Framework Instantiation Workflow

the baseline analysis in order to prune infeasible behavior at its level of approximation.

The derivation of our *co-runner-sensitive* analysis, in contrast, comprises three steps:

Step 1: The derivation starts from a compound baseline abstract model that consists of one co-runner-insensitive analysis per processor core. Thus, the compound baseline analysis argues about all processor cores. We consider a property that bounds the blocked cycles of the core under analysis based on the access cycles of the concurrent cores assuming a work-conserving bus arbitration (like e.g. Round-Robin).

Step 2: The work-conserving property is lifted to the baseline abstract model. However, it would be unpractical to enumerate all combinations of abstract traces of the analyses for the different cores (since the compound abstract model is defined as cross product over its components).

Step 3: Hence, we further lift the already lifted property to the component analysis only focusing on the core under analysis. To this end, we assume per concurrent core the maximum amount of access cycles possible in any interval no longer than the current WCET bound of the core under analysis. The resulting analysis starts by calculating the co-runner-insensitive WCET bound for the core under analysis. Then, it calculates upper bounds on the concurrent access cycles and subsequently recalculates the WCET bound. This process is repeated until a fixed point is reached.

VIII. EXPERIMENTAL EVALUATION

We evaluate our analysis prototype for multi-core processors with ARM[®] instruction set, a shared memory bus, and Round-Robin bus arbitration. Our experiments consider cores with in-order pipelines (five stages) as well as cores that support out-of-order execution (Tomasulo dynamic scheduling, three functional units, and speculative execution). We also consider two scenarios with respect to the local instruction memories of the cores. First, we assume a local instruction scratchpad that is statically initialized with all programs executed on the core. Secondly, we consider a local instruction cache (1KiB size, least-recently-used [LRU] replacement policy) that is connected to the shared bus. Table II lists the four resulting core configurations. All core configurations assume a local LRU data cache of size 1KiB. The shared bus connects the cores to an SRAM background memory that serves accesses with a fixed latency of ten cycles. Note that we do not precisely model any particular commercial processor.

We consider a dual-core, a quad-core, and an octa-core processor per core configuration. For each resulting hardware configuration, we calculate a co-runner-insensitive WCET bound per benchmark. Our benchmark suite contains 31 benchmarks of the Mälardalen suite [26]

and six larger programs generated from SCADE models¹. We assume non-preemptive task scheduling as providing timing guarantees for preemptive multitasking is an unsolved problem for realistic hardware platforms. Note that we do not perform response time analysis. This work focuses on WCET analysis.

Our experiments take about 107 minutes on a quad-core Intel[®] Core[™] i7 processor clocked at 2.4 GHz and provided 8 GiB of main memory.

We normalize the WCET bound and the analysis runtime per benchmark and considered processor to the corresponding values of an analysis that ignores the shared-bus interference. Table I lists the geometric means of the resulting factors for the considered hardware platforms.

The results underline the strong impact of the shared-bus interference on the WCET bounds: the average deviation factors of the WCET bounds from bounds ignoring the interference reach up to 1.756 (3.267, 6.284) for dual-core (quad-core, octa-core) processors. However, note that the calculated WCET bounds are co-runner-insensitive. Thus, they implicitly assume arbitrarily aggressive bus access behavior of the programs executed on the concurrent cores. As we have shown before [17], a co-runner-sensitive analysis can lead to a significant reduction of the WCET bounds.

In contrast to classical compositional approaches [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], [15], our analysis prototype supports hardware platforms exhibiting indirect interference effects (cf. Section III). We estimate the additional cost of considering indirect interference effects by comparing the analysis runtime to the runtime of an analysis that ignores all interference effects (which is the main part of classical compositional timing analysis). The average increase in analysis runtime is moderate (up to 5.4 percent) for hardware platforms with in-order execution or instruction scratchpads ($Conf_{is}^{io}$, $Conf_{is}^{ooo}$, $Conf_{ic}^{io}$). The combination of out-of-order execution and instruction caches ($Conf_{ic}^{ooo}$), however, leads to a significantly higher—though still bearable—increase in analysis runtime (up to 15.9 percent on average). Intuitively, the complexity of modeling the pipeline features multiplies with the complexity of modeling the shared-bus interference by non-determinism. Note that these runtime results are a significant improvement compared to the numbers we reported in our earlier work. The improvement stems from engineering improvements (which are not in the scope of this paper) of the implementation of our analysis.

The average runtime increase factors for dual-core, quad-core, and octa-core processors with the same core configuration are essentially identical for all our experiments (the small deviations are caused by the heavy use of hash sets in our prototype implementation). Intuitively, for the considered processor core configurations, the core pipelines already converge for each access to the shared bus in a dual-core processor. As a consequence, further cycles blocked at the shared bus do not result in new pipeline states. An optimization (fast-forwarding of converged chains [17]) in our analysis prototype exploits this convergence. For analyses relying on the enumeration of all interleavings of bus access requests by the different processor cores [20], in contrast, each additional core increases the analysis runtime by a factor. Thus, such analyses do not scale to high numbers of processor cores.

IX. FUTURE WORK

Our current prototype implementation only takes into account shared-bus interference. We plan to also consider shared caches and

¹<http://www.esterel-technologies.com/products/scade-suite>

	$Conf_{is}^{io}$			$Conf_{is}^{ooo}$			$Conf_{ic}^{io}$			$Conf_{ic}^{ooo}$		
	2-Core	4-Core	8-Core	2-Core	4-Core	8-Core	2-Core	4-Core	8-Core	2-Core	4-Core	8-Core
WCET bound	1.579	2.728	5.022	1.660	2.978	5.609	1.677	3.024	5.714	1.756	3.267	6.284
analysis runtime	1.033	1.033	1.028	1.054	1.046	1.051	1.050	1.037	1.038	1.159	1.152	1.149

Table I: Average deviation factors of WCET bound and analysis runtime of the calculation of co-runner-insensitive WCET bounds with respect to an analysis assuming no interference

	in-order execution	out-of-order execution
local instruction scratchpad	$Conf_{is}^{io}$	$Conf_{is}^{ooo}$
local instruction cache	$Conf_{ic}^{io}$	$Conf_{ic}^{ooo}$

Table II: Evaluated processor core configurations

cache coherence in a future version of our tool. A long-term goal is the modeling of commercially available multi-core processors with our techniques.

Furthermore, we plan to study the impact of complex processor core features like store buffers and speculation on the performance of our analysis approach. In this context, we will investigate performance improvements of our tool in order to further reduce the performance overhead due to the consideration of shared-resource interference. As a result of our studies, we will give recommendations for the design of future multi-core hardware platforms to enable their use in timing-critical embedded system.

X. SUMMARY

We present a framework for the derivation of WCET analyses for multi-core processors. It centers around the concept of property lifting. Instances of the framework are sound WCET analyses. The framework has been successfully used in the development of a novel analysis that avoids the restrictions of existing approaches.

ACKNOWLEDGMENTS

The authors would like to thank Mihail Asavae, Florian Hauptenthal, Max John, Jan Reineke, and Reinhard Wilhelm for many comments and interesting discussions. This work was supported by the DFG as part of the Transregional Collaborative Research Centre SFB/TR 14 (AVACS) and by the Saarbrücken Graduate School of Computer Science which receives funding from the DFG as part of the Excellence Initiative of the German Federal and State Governments.

REFERENCES

- [1] R. Wilhelm *et al.*, “The worst-case execution-time problem — overview of methods and survey of tools,” *ACM Trans. Embed. Comput. Syst.*, vol. 7, no. 3, pp. 36:1–36:53, 2008.
- [2] S. Thesing, “Safe and precise WCET determination by abstract interpretation of pipeline models,” Ph.D. dissertation, 2004.
- [3] X. Li *et al.*, “Modeling out-of-order processors for WCET analysis,” *Real-Time Syst.*, vol. 34, no. 3, pp. 195–227, 2006.
- [4] A. Abel *et al.*, “Impact of resource sharing on performance and performance prediction: A survey,” in *Proceedings of the 24th Conference on Concurrency Theory*, 2013, pp. 25–43.
- [5] A. Schranzhofer *et al.*, “Timing analysis for TDMA arbitration in resource sharing systems,” in *Proceedings of the 16th IEEE Real-Time and Embedded Technology and Applications Symposium*, 2010, pp. 215–224.
- [6] R. Pellizzoni *et al.*, “Worst case delay analysis for memory interference in multicore systems,” in *Proceedings of the 13th Conference on Design, Automation and Test in Europe*, 2010, pp. 741–746.
- [7] A. Schranzhofer *et al.*, “Timing analysis for resource access interference on adaptive resource arbiters,” in *Proceedings of the 17th IEEE Real-Time and Embedded Technology and Applications Symposium*, 2011, pp. 213–222.
- [8] G. Giannopoulou *et al.*, “Timed model checking with abstractions: Towards worst-case response time analysis in resource-sharing manycore systems,” in *Proceedings of the 10th ACM International Conference on Embedded Software*, 2012, pp. 63–72.
- [9] Y. Liang *et al.*, “Timing analysis of concurrent programs running on shared cache multi-cores,” *Real-Time Systems*, vol. 48, pp. 638–680, 2012.
- [10] J. Nowotsch, “Interference-sensitive worst-case execution time analysis for multi-core processors,” Ph.D. dissertation, 2014.
- [11] D. Dasari *et al.*, “Response time analysis of COTS-based multicores considering the contention on the shared memory bus,” in *Proceedings of the 10th IEEE International Conference on Trust, Security and Privacy in Computing and Communications*, 2011, pp. 1068–1075.
- [12] D. Dasari and V. Nélis, “An analysis of the impact of bus contention on the WCET in multicores,” in *Proceedings of the 14th IEEE International Conference on High Performance Computing and Communication & the 9th IEEE International Conference on Embedded Software and Systems*, 2012, pp. 1450–1457.
- [13] S. Schliecker and R. Ernst, “Real-time performance analysis of multiprocessor systems with shared memory,” *ACM Trans. Embedded Comput. Syst.*, vol. 10, no. 2, p. 22, 2010.
- [14] S. Schliecker *et al.*, “Response time analysis in multicore ECUs with shared resources,” *IEEE Trans. Industrial Informatics*, vol. 5, no. 4, pp. 402–413, 2009.
- [15] S. Altmeyer *et al.*, “A generic and compositional framework for multicore response time analysis,” in *Proceedings of the 23rd International Conference on Real Time Networks and Systems*, 2015, pp. 129–138.
- [16] S. Chattopadhyay *et al.*, “A unified WCET analysis framework for multi-core platforms,” in *Proceedings of the 18th IEEE Real-Time and Embedded Technology and Applications Symposium*, 2012, pp. 99–108.
- [17] M. Jacobs *et al.*, “WCET analysis for multi-core processors with shared buses and event-driven bus arbitration,” in *Proceedings of the 23rd International Conference on Real Time Networks and Systems*, 2015, pp. 193–202.
- [18] S. Hahn *et al.*, “Towards compositionality in execution time analysis – definition and challenges,” in *Proceedings of the 6th International Workshop on Compositional Theory and Technology for Real-Time Embedded Systems*, 2013.
- [19] T. Lundqvist and P. Stenstrom, “Timing anomalies in dynamically scheduled microprocessors,” in *Proceedings of the 20th IEEE Real-Time Systems Symposium*, 1999, pp. 12–21.
- [20] T. Kelter and P. Marwedel, “Parallelism analysis: Precise WCET values for complex multi-core systems,” in *Revised Selected Papers of the 3rd International Workshop on Formal Techniques for Safety-Critical Systems*, 2014, pp. 142–158.
- [21] J. Reineke and R. Sen, “Sound and efficient WCET analysis in the presence of timing anomalies,” in *Proceedings of the 9th International Workshop on Worst-Case Execution Time Analysis*, 2009, pp. 98–108.
- [22] M. Jacobs, “Improving the precision of approximations in WCET analysis for multi-core processors,” in *Proceedings of the 7th Junior Researcher Workshop on Real-Time Computing*, 2013, pp. 1–4.
- [23] P. Cousot and R. Cousot, “Abstract interpretation: a unified lattice model for static analysis of programs by construction or approximation of fixpoints,” in *Conference Record of the 4th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, 1977, pp. 238–252.
- [24] R. Pellizzoni and M. Caccamo, “Impact of peripheral-processor interference on WCET analysis of real-time embedded systems,” *IEEE Transactions on Computers*, vol. 59, pp. 400–415, 2010.
- [25] Y.-T. S. Li and S. Malik, “Performance analysis of embedded software using implicit path enumeration,” in *Proceedings of the 32nd Annual ACM/IEEE Design Automation Conference*, 1995, pp. 456–461.
- [26] J. Gustafsson *et al.*, “The malmö WCET benchmarks - past, present and future,” in *Proceedings of the 10th International Workshop on Worst-Case Execution Time Analysis*, 2010.