# Resource allocation for complex DAG tasks with probabilistic execution times

Paola Cappanera, Laura Carnevali, Leonardo Paroli, Riccardo Reali, Enrico Vicario
Department of Information Engineering, University of Florence, Italy
Email: {paola.cappanera, laura.carnevali, riccardo.reali, enrico.vicario}@unifi.it, leonardo.paroli1@stud.unifi.it

## I. Problem Formulation

In the context of real-time systems and applications, a relevant problem concerns the partition of sub-tasks of a task graph [1], [2] among resources, according to some optimality criteria that depends on a soft deadline. An example is the allocation of sub-tasks to different processor cores, with each allocation being characterized by a probabilistic Worst Case Execution Time (pWCET) [3] comprising a stochastic upper bound on the execution time of all possible execution conditions of the sub-task on the core. In particular, the allocation could be driven by the intent of minimising the energy effort of the involved cores, and attempting to meet a soft deadline on the task end-to-end response time. Let $G$ be a task graph structured as a non-recurrent directed acyclic graph (DAG), $T := \{T_0, T_1, ...T_N\}$ be the set of sub-tasks of $G$, $R := \{R_0, R_1, ...R_M\}$ be the set of resources that can serve the sub-tasks of $G$, $f_{ij}\left(t\right)$ be the generally-distributed (GEN) bounded-supported PDF of the execution time of sub-task $i$ on core $j$, $f_G\left(t\right)$ and $F_G\left(t\right)$ be the Probability Density Function (PDF) and the Cumulative Distribution Function (CDF), respectively, of the end-to-end response time of $G$ (to be computed), $d_G\left(t\right)$ be a penalty function whose support upper bound comprises a soft deadline for $G$, and $U(F_G(t), d_G(t))$ be the revenue function that depends on $F_G\left(t\right)$ and $d_G\left(t\right)$. We are interested in evaluating the partition of sub-tasks on resources, that maximizes the revenue function $U(F_G(t), d_G(t))$, i.e. we want to solve the following optimization problem:

$$\max_{\delta \in \Delta} U(F_G(t) \,|\, \delta, d_G(t)) \tag{1}$$

where $\Delta$ is the set of allocations between sub-tasks and resources, and $F_G(t) \,|\, \delta$ is the end-to-end response time CDF of $G$ conditioned on allocation $\delta$.

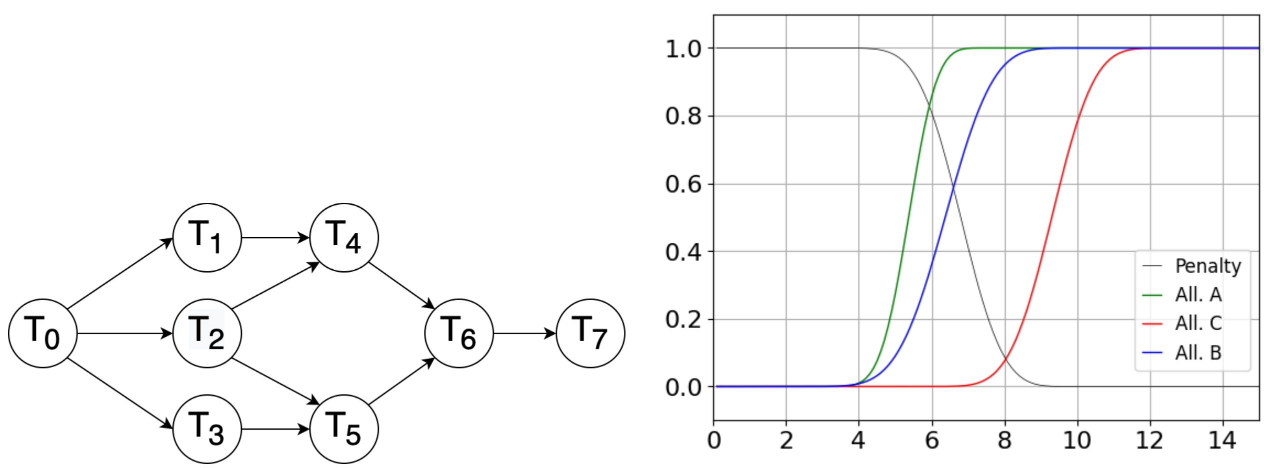


(a) (b)

Fig. 1: a) A task graph with 8 sub-tasks. b) A penalty function (grey line) and end-to-end response time CDFs (green, blue and red lines) of the task graphs obtained varying the probability distributions associated with sub-tasks.

In Fig. 1a, it is shown a task graph with 8 sub-tasks, organised in a DAG structure. Fig. 1b shows the CDFs of the end-to-end response time of the task graph, obtained considering three different resource allocations. The figure also shows an example of penalty function, which comprises a soft deadline on the execution time upper bound (which is 10 in the figure). For each allocation, all sub-tasks are associated with a uniform distribution. In particular, allocations A, B and C are associated with uniform distributions having support $[0.5, 1.5]$, $[0.3, 2]$ and $[1, 2.5]$, respectively. As it is shown, different allocations notably impact on the end-to-end response time of the task graph, and thus on the severity of the penalty function. For example, allocation A ends before the soft deadline with probability 1, while allocation C is more likely to end after the deadline. Hence, allocation C will be penalized more than allocation A, which turns out to be a better allocation.

## II. Scientific Ground

The problem described in Section I is strongly connected to the evaluation of the end-to-end response time of the task graph. Task graphs comprise workflows of activities orchestrated through precedence constraints and control flow constructs. PDFs of activity durations can be derived by fitting data, requiring the use of GEN PDFs for the sake of expressivity. Moreover, many real-time problems constrain sub-tasks to end within specific time bounds, leading to the requirement of PDFs with bounded supports. Workflows with GEN activities over bounded supports underlie non-Markovian processes, which can be evaluated through *regenerative transient analysis* based on the method of *stochastic state classes* [4]. The method is implemented in the SIRIO library of the Oris tool [5] for systems specified as Stochastic Time Petri Nets (STPNs).

When the model grows in the degree of concurrency, the proposed solutions do not scale efficiently. In these cases, a compositional approach can be considered, where models are decomposed into simpler sub-workflows, which are analysed efficiently in isolation and their response times are then recombined together, leading to a stochastic upper bound of the end-to-end response time of the starting workflow. Some preliminary results are shown in [6]. The paper proposes three compositional analysis heuristics demonstrating a good trade-off between analysis efficiency and accuracy. The compositional evaluation methods are implemented in a Java library that depends on the SIRIO library to define the STPN of a workflow and to perform regenerative transient analysis of the workflow STPN. The library is designed to represent generic workflows and evaluate their response time PDF. Nevertheless, by deploying Model Driven Engineering (MDE) practices, model-to-model transformations can be exploited to map task graphs into workflows, enabling evaluation of the task graph end-to-end response time PDF.

The described scientific ground allows the optimization problem in Eq. (1) to be solved. A first naive approach consists in enumerating the possible sub-task to resource allocations. Certainly, this strategy appears to be very inefficient, especially when the number of tasks and resources that can serve them increases. As shown in [7], more efficient strategies should leverage heuristics. An alternative heuristics could be defining different revenue functions (based on different features extracted from the PDFs of the sub-tasks, such as the mean value), solving the corresponding optimization problems, and selecting one of the obtained solutions (e.g., through a voting mechanism).

## III. Open Issues

The identification of the optimal allocation of resources on tasks of a system/application is a challenge that comes along with the following open issues:

- **OI1 - Evaluation of task graphs with recurrent sub-tasks.** We have considered taks graphs with no recurrent sub-tasks, whose evaluation relies on procedures that are tailored for this type of structure. How does the introduction of recurrent sub-tasks affect the evaluation technique, both in terms of efficiency and accuracy?
- **OI2 - Allocation of task graphs with recurrent sub-tasks.** Is the optimal solution of the allocation problem affected by the introduction of recurrent sub-tasks? Is it required to vary the revenue function to actually obtain the optimal (or sub-optimal) solution?
- **OI3 - Deadlines on single sub-tasks.** In optimizing the revenue function, no assumptions are made on the response times of individual sub-tasks. However, in many practical situations, soft deadlines and penalties are assigned to one or more individual sub-tasks. In this case, how is the allocation problem reformulated, in terms of constraints and revenue function?
- **OI4 - Deadline miss.** Allocation of sub-tasks is done a priori, based on time information such as the pWCET or data-based distribution. If a resource accumulates delay on one or more sub-tasks, the task graph can miss the deadline. In this case, how are critical tasks identified and what is their impact on the final response time?
- **OI5 - Allocation rescheduling.** If a task accumulates delay, what strategies can be adopted to reschedule the allocation of resources on the tasks following the delayed task?

## References


[1] Z. Houssam-Eddine, N. Capodieci, R. Cavicchioli, G. Lipari, and M. Bertogna, "The hpc-dag task model for heterogeneous real-time systems," *IEEE Transactions on Computers*, vol. 70, no. 10, pp. 1747–1761, 2020.

[2] A. Melani, M. Bertogna, V. Bonifaci, A. Marchetti-Spaccamela, and G. Buttazzo, "Schedulability analysis of conditional parallel task graphs in multicore systems," *IEEE Transactions on Computers*, vol. 66, no. 2, pp. 339–353, 2016.

[3] G. Bernat, A. Colin, and S. M. Petters, "Wcet analysis of probabilistic hard real-time systems," in *23rd IEEE Real-Time Systems Symposium, 2002. RTSS 2002.* IEEE, 2002, pp. 279–288.

[4] A. Horváth, M. Paolieri, L. Ridi, and E. Vicario, "Transient analysis of non-markovian models using stochastic state classes," *Performance Evaluation*, vol. 69, no. 7-8, pp. 315–335, 2012.

[5] M. Paolieri, M. Biagi, L. Carnevali, and E. Vicario, "The oris tool: quantitative evaluation of non-markovian systems," *IEEE Transactions on Software Engineering*, vol. 47, no. 6, pp. 1211–1225, 2019.

[6] L. Carnevali, M. Paolieri, R. Reali, and E. Vicario, "Compositional safe approximation of response time distribution of complex workflows," in *International Conference on Quantitative Evaluation of Systems*. Springer, 2021, pp. 83–104.

[7] H. Chen, G. Ding, S. Qin, and J. Zhang, "A hyper-heuristic based ensemble genetic programming approach for stochastic resource constrained project scheduling problem," *Expert Systems with Applications*, vol. 167, p. 114174, 2021.