

Specification and Analysis of Resource Utilization Policies for Human-Intensive Systems (Extended Abstract)

Seung Yeob Shin

University of Massachusetts, Amherst MA 01003, USA

`shin@svv.lu`

https://wwwen.uni.lu/snt/people/seung_yeob.shin

Abstract. Societal processes, such as those used in healthcare, typically depend on the effective utilization of resources, both human and non-human. Sound policies for the management of these resources are crucial in assuring that these processes achieve their goals. But complex utilization policies may govern the use of such resources, increasing the difficulty of accurately incorporating resource considerations into complex processes. This dissertation presents an approach to the specification, allocation, and analysis of the management of such resources.

Keywords: resource utilization policy, discrete-event simulation, model checking, human-intensive systems, process modeling, process analysis

1 Introduction

Complex real-world processes with human, software, and hardware resources integrated to perform key functions, play an important role in our society. Because access to these resources is usually limited both by their quantity and by restrictions on their availability, contention for them is often a serious problem. Understanding how such contention arises requires that these process specifications be sufficiently powerful and precise to specify just how process activities make use of these resources. This is further complicated by the need to also precisely define potentially complex and numerous resource allocation policies that are driven by process goals, regulations, or the need to satisfy the interests of different stakeholders. For instance, hospital emergency department (ED) processes use diverse resources, such as medical staff, beds, and devices that are often severely limited in quantity and are restricted by policies that often conflict. For example, in many EDs, a patient should be cared for by the same doctor and nurse for the entire stay, but under unusual circumstances this policy may be violated to improve efficiency in patient care. Suboptimal policies for resolving hospital resource management conflicts can result in such problems as overcrowding, inefficient staff utilization, and long length of patient stay.

Process resource management has been widely studied through modeling and analysis, but most prior work has used relatively simple resource models

that have difficulty specifying the necessary complex resource concerns flexibly and rigorously. This dissertation provides a framework for the specification and analysis of resources and their utilization policies in such complex processes. The research makes the following main contributions: (1) a precise specification language for resources and resource utilization policies, (2) a process- and resource-aware discrete-event simulation (DES) system that supports the performance of process simulations that adhere to the specifications, (3) process- and resource-aware model checking tools that can verify process properties and the adherence of process simulations to these policies, (4) elaborated patient care process models based on real-world data and domain expert knowledge, and (5) a case study applying these approaches to validating and verifying the soundness of resource allocations to patient care processes in an ED.

2 Resource and Resource Utilization Policy Modeling

Our resource modeling approach relies on a precise, well-defined process model. For the ED domain, this means a detailed model of the process by which patients are treated. We use the Little-JIL language [6] to specify this model. Little-JIL process definitions use hierarchical decomposition, exception handling, concurrency and human choice to structure process activities, where each activity incorporates the specification of the resources needed to perform the activity.

Our resource specification is orthogonal to, and separate from, activity and dataflow specifications. A resource is modeled as the composition of a set of attributes and a set of capabilities. A resource’s attributes, e.g., age, experience, model number, memory footprint, and skill level, are used in deciding which resource instance is assigned to a requesting activity. One particularly important attribute in an ED domain is the work shift, specifying the times when a resource can be allocated to an activity. A resource’s capabilities are the activities that the resource can perform. For example, a set of capabilities of a doctor (MD) includes prescribing medications and ordering tests.

Given the existence of such resource specifications, this approach enables specification of three different kinds of resource utilization policies: permission constraint policies, scheduling policies, and conflict resolution policies. A permission constraint policy specifies the permissibility of a resource to handle a request as restricted by a specified guard (defined as a Boolean expression). For example, the **SameMD** policy can constrain the choice of MD who can handle a request a specific patient, and the **ShiftMD** policy can constrain the times when an MD can treat a patient. Schedule policies support the specification of contention and selection policies. Contention policies specify precedence among requests.

In a hospital ED, when multiple patient care activities for different patients require the service of more MDs than are currently available for allocation, an appropriate scheduling policy is necessary to resolve the contention problem among the requests. The selection policy complements the contention policy, supporting specification of precedence among the resources that are able to handle a resource request. For instance, when a new patient arrives in an ED, there are

usually more than two MDs who can assess the patient. An appropriate workload policy can balance the workloads of the MDs. Conflict resolution policies specify how to deal with policies that come into conflict with each other and with the situation where two or more policies cannot be enforced simultaneously. For instance, patient handoff in a hospital ED is specified by using a conflict resolution policy because MDs are not able to satisfy both the `SameMD` and `ShiftMD` policies at the ends of their shifts [2, 4].

3 Dynamic Analysis: Discrete-Event Simulation

The process- and resource-aware DES framework supports the execution of simulations of processes whose activities are modeled in Little-JIL, and whose resources are modeled in the resource modeling approach presented in this dissertation by extending JSim [3]. The extension enables the simulator to support scenarios such as (1) allowing MDs to have varying shift constraints, (2) enforcing the `SameMD` policy unless it contradicts the `ShiftMD` policy, at which point the patient is handed off to a new MD, and (3) enabling a variety of scheduling policies, such as handing the sickest patient first, as opposed to using the least utilized resource first [4, 5].

The dissertation presents a novel constraint-aware resource scheduling approach that consists of three steps. First, an algorithmic control method embedded within the resource-aware DES framework computes resource requirements, such as how many of each resource must be present at each time epoch to meet user-specified resource utilization requirements. Second, deterministic Integer Linear Programming (ILP) produces a resource schedule that satisfies those resource requirements and user-specified constraints on resource utilization. Finally, the resource-aware DES computes how the resource schedule affects statistical estimates of the system’s runtime properties.

4 Static Analysis: Model Checking

This process- and resource-aware DES framework shows promise as a vehicle for evaluating diverse resource utilization policies in a flexible manner. However, the dynamic approach verifies only that a single simulation adheres to the resource utilization constraints, but it is inherently unable to assure that all possible simulations adhere to the constraints. The dissertation develops a novel approach to analyzing the application of resource utilization policies statically.

Model checking is a static analysis technology that has long been used to either demonstrate that all possible executions of a program adhere to specified properties, or to identify counterexample executions that violate a property, e.g., FLAVERS [1]. Typically, the properties are event sequence specifications characterizing program functionality. This dissertation uses this technology to statically verify that all simulation executions adhere to resource utilization properties, increasing the credibility of the simulation studies and their applicability in practice.

As with other model checking tools, FLAVERS requires three kinds of inputs: a detailed and rigorously-defined flow graph, a property finite state machine (FSM), and a set of constraint FSMs. Given the precise specifications of resources and their utilization policies, the process- and resource-aware model checking approach presented in the dissertation augments the three activity-focused inputs to FLAVERS with resource utilization specifications, rigorously defining feasible and infeasible resource allocations and utilization policies. This enables FLAVERS to verify that resource utilization policies must always be adhered to by all possible simulation runs. As an example, the dissertation presents a verification that a specified ED patient care process satisfies the `SameMD` and `ShiftMD` policies. This static analysis approach has also been used to support reachability analyses that have detected the presence in some example processes of resource-driven deadlocks and conflicts among policies.

5 Results

The dissertation presents an evaluation of the resource modeling and analysis approach by applying it to the challenging study of a detailed hospital ED process. The ED process model is based on real-world data and the knowledge of a domain expert with extensive experience as an emergency department physician and ED manager at the Baystate Medical Center (BMC), in Springfield, MA, USA. MD and nurse (RN) resources

are scheduled simultaneously while adhering to the complex characteristics of the ED such as time-varying patient arrivals, constraints and policies of medical providers, other hospital resource (e.g., beds) utilization, and complex patient care processes. The evaluation demonstrates that this scheduling approach creates better staffing than existing real-world staffing in terms of balancing resource utilization over a 24-hour period (see Fig. 1).

The DES framework provides detailed trace information about resource utilization to support the dynamic analysis of different resource utilization policies. However, this dynamic analysis is valid only when all the simulations of resource utilization adhere to the defined resource utilization policies. Therefore, the dissertation also includes a demonstration that the formality of the resource specification supports verification of resource properties such as the absence of violations of resource utilization policies and the absence of deadlock. Given the entire ED specifications, the static analysis approach verifies the adherence to the `SameMD` and `ShiftMD` policies; and finds conflict among hospital policies or detects a resource deadlock under various specification of the numbers of MDs, RNs, beds, and patients.




Fig. 1. Utilization comparison over 24 hours a day: RN Utilization (BMC) RN utilization of Baystate Medical Center, RN Utilization (New) RN utilization derived by the constraint-aware scheduling approach

6 Contributions

Complex societal processes integrate multiple types of resources, such as humans, software, and hardware, to support the performance of process activities. Participation by these resources is often restricted by diverse resource utilization policies. Further, policies may at times conflict with each other, requiring conflict resolution strategies that add extra complexity. Resource utilization policies usually significantly impact system behaviors and results. Therefore, resource utilization policies should be thoroughly evaluated and rigorously analyzed.

This dissertation presents a novel approach to create specifications of policies that are sufficiently precise and detailed to support static and dynamic analyses of how these policies affect the properties of processes that are governed by these policies. The presented process- and resource-aware DES framework supports simulations of such processes that adhere to resource specifications. The framework’s evaluation of hospital ED processes demonstrated that it supports considerable flexibility in resource and utilization policy specification and provides powerful dynamic analyses. However, the approach is inherently limited because of the impossibility of exhaustive simulations of all scenarios. To complement the dynamic simulation approach, this work develops a process- and resource-aware static analysis approach that globally verifies system properties and adherence to resource utilization policies. In doing this, the dissertation demonstrates how the specification and analysis framework can be effective in guiding domain experts towards sound decisions about policies for the management of hospital resources, while also providing rigorously-based assurances that the guidance is reliable and well-founded.

References

1. M. B. Dwyer, L. A. Clarke, J. M. Cobleigh, and G. Naumovich. Flow analysis for verifying properties of concurrent software systems. *ACM Transactions on Software Engineering and Methodology*, 13(4):359–430, Oct. 2004.
2. P. L. Henneman, S. Shin, Y. Brun, H. Balasubramanian, F. Blank, and L. J. Osterweil. Using computer simulation to study nurse-to-patient ratios in an emergency department. *The Journal of Nursing Administration*, 45(11):551–556, 2015.
3. M. S. Raunak, L. J. Osterweil, A. Wise, L. A. Clarke, and P. L. Henneman. Simulating patient flow through an emergency department using process-driven discrete event simulation. In *Software Engineering in Health Care*, 2009.
4. S. Shin, Y. Brun, H. Balasubramanian, P. L. Henneman, and L. J. Osterweil. Discrete-event simulation and integer linear programming for constraint-aware resource scheduling. *IEEE Transactions on Systems, Man, and Cybernetics: Systems*, 47, 2017.
5. S. Shin, Y. Brun, L. J. Osterweil, H. Balasubramanian, and P. L. Henneman. Resource specification for prototyping human-intensive systems. In *International Conference on Fundamental Approaches to Software Engineering*, pages 332–346, 2015.
6. A. Wise. Little-JIL 1.5 language report. Technical Report 2006–051, Department of Computer Science, University of Massachusetts, Amherst, 2006.