

1. Record Nr.	UNINA9910437571803321
Autore	Li Dawei
Titolo	Energy-aware scheduling on multiprocessor platforms // Dawei Li, Jie Wu
Pubbl/distr/stampa	New York, : Springer, 2013
ISBN	1-4614-5224-4
Edizione	[1st ed. 2013.]
Descrizione fisica	1 online resource (66 p.)
Collana	SpringerBriefs in computer science, , 2191-5768
Altri autori (Persone)	WuJie
Disciplina	004.21
Soggetti	Multiprocessors
Lingua di pubblicazione	Inglese
Formato	Materiale a stampa
Livello bibliografico	Monografia
Note generali	Description based upon print version of record.
Nota di bibliografia	Includes bibliographical references.
Nota di contenuto	Introduction -- System Model -- Scheduling on Homogeneous DVFS Multiprocessor Platforms -- Scheduling on Heterogeneous DVFS Multiprocessor Systems -- Related Work -- Conclusion and Future Directions.
Sommario/riassunto	Multiprocessor platforms play important roles in modern computing systems, and appear in various applications, ranging from energy-limited hand-held devices to large data centers. As the performance requirements increase, energy-consumption in these systems also increases significantly. Dynamic Voltage and Frequency Scaling (DVFS), which allows processors to dynamically adjust the supply voltage and the clock frequency to operate on different power/energy levels, is considered an effective way to achieve the goal of energy-saving. This book surveys existing works that have been on energy-aware task scheduling on DVFS multiprocessor platforms. Energy-aware scheduling problems are intrinsically optimization problems, the formulations of which greatly depend on the platform and task models under consideration. Thus, Energy-aware Scheduling on Multiprocessor Platforms covers current research on this topic and classifies existing works according to two key standards, namely, homogeneity/heterogeneity of multi-processor platforms and the task types considered. Under this classification, other sub-issues are also included, such as, slack reclamation, xed/dynamic priority scheduling, partition-based/global scheduling, and application-specific power consumption, etc.

