

Modelling Energy-Aware Task Allocation in Mobile Workflows

Bo Gao and Ligang He

Department of Computer Science,
University of Warwick,
Coventry, CV4 7AL, UK,
`liganghe@dcs.warwick.ac.uk`

Abstract. Mobile devices are becoming the platform of choice for both business and personal computing needs. For a group of users to efficiently collaborate over the execution of a set workflow using their mobile devices, the question then arises as to which device should run which task of the workflow and when? In order to answer this question, we study two common energy requirements: in the *minimum group energy cost problem (MGECP)* we build the model as a quadratic 0-1 program and solve the optimisation problem with the objective to minimise the total energy cost of the devices as a group. In the *minimum max-utilisation problem (MMUP)* we aim to improve the fairness of the energy cost within the group of devices and present two adjustment algorithms to achieve this goal. We demonstrate the use of a Mixed Integer Quadratic Programming (MIQP) solver in both problem's solutions. Simulation result shows that both problems are solved to good standards. Data generated by different workload pattern also give us a good indication of the type of workflow that benefit the most from MMUP. The model used in this work can also be adapted for other energy critical scenarios.

Key words: mobile computing, energy-aware, collaboration, workflow

1 Introduction

Recent years have seen significant growth in the size of the mobile computing market, and yet the rarest commodity in the world of mobile computing remains to be its battery power. Development in battery technology is slow compared to other components of a mobile smart device. Hence, despite the moderately improved battery capacity on modern smart devices, with increasingly more complex functionality required from the user, developments of mobile applications remain largely energy-constrained [17].

In less than a decade, mobile devices have enriched their functionalities from being a simple dialling device to a hub of rich media applications. It is predicted that by 2015, mobile application development projects will outnumber desktop projects by a ratio of 4:1 [10]. The unique portability of a mobile device coupled with its ever growing hardware capability brings business and ad hoc workflows

that are traditionally supported by fixed location resources to be implemented over wireless mobile platforms.

Researches show that in a mobile environment, communication tasks are especially energy-demanding compared to local computation tasks[16]. Hence, this type of applications, namely *mobile workflows*, which has a particular emphasis on collaboration between users, is likely to be more energy-demanding than others and requires to be managed in an energy-efficient manner. Furthermore, unlike its desktop counterparts, mobile computing devices are often exposed to the open environment. Changing conditions in data connection, sudden drain of battery caused by user actions can bring disruption to a device’s availability.

Our research investigate ways to model and analyse the energy efficiency of such workflows running atop a group of mobile devices. Our goal is to provide an energy efficient execution platform for mobile workflows, while utilising fair share of each mobile device’s energy. Our objective is two-fold: First, in the MGECP, we aim to minimise the workflow’s total energy consumption. Second, in the MMUP, our objective is to minimize the maximum device utilisation in the group while keeping the overall energy cost close to the minimum.

In this paper, we first give use cases from possible application areas of a mobile workflow and discuss related work. We then construct the system energy model in section 3. The allocation problem is modelled as a quadratic 0-1 program, and its two objectives (MGECP and MMUP) are studied in sections 4 and 5. We conclude the paper in section 6 with a simulation study the result of which verifies our formulation and compares the results of our algorithms when applied to different types of workflows.

2 Applications and Related Work

Mobile workflow can be found when a group of mobile users are to share or communicate with each other in order to accomplish a certain task. Such scenario commonly exists in a business environment. With growing adaptation of mobile devices within their business models[10], modern *enterprise applications* often include or are entirely based on mobile devices. For instance, in a supply chain business, as illustrated in Figure 1a, the commencement of a workflow is triggered




Fig. 1: Examples of mobile workflow use cases. In the centre of both figures, the tasks between the pair of “)” are not restricted to be executed on any specific device. Allocation of these tasks can affect the energy-efficiency of the workflow.

by a member of staff registering receipt/sales of goods on their mobile devices. The system database is then updated via a query module¹. A forecast module is then evoked to produce a forecast based on the update, which is then projected onto the manager’s smartphone as a live trend graph or a production plan. In order to lower the overall energy cost of the workflow, our objective in section 4, the forecast module which requires complex computation for data mining purposes is more suitable to be allocated to a device that has a fast processor and low energy draw while running computation tasks. Additionally, it is preferable that modules that communicate frequently, e.g. invoice and query, are allocated to the same device to reduce communication cost.

Another use case illustrated in Figure 1b includes the use of three smartphones and a tablet, and expands on the idea of a popular *consumer application* [12] which lets its users to play darts with their mobile devices. During the game, a tablet is used to display the dartboard, participating players use their smartphones as darts. The workflow starts when a player throws a dart (by a throwing gesture from the phone towards the tablet). Sensor readings (accelerometer and gyroscope) are then taken from the phone and fed into a calculation module to work out where the dart should land on the board. Result from the calculation is then passed on to the display module on the tablet.

Like all multi-player competitions, the game can only function until its weakest player withdraws, which in this case, is the device that runs out of battery first. Although the calculation module does not require much energy at each run, repeated execution is required. As the game goes on, the battery of the device to which the calculation task is allocated drains faster than the others’. A fair task allocation, which we study in section 5, is needed in such scenarios to balance the contribution made by participating members of the mobile workflow.

A workflow engine is often required to oversee the execution of mobile workflows. In [15] a detailed mobile workflow engine is implemented and tested on Nokia devices. A decentralised workflow coordination architecture designed for mobile devices is presented in [1] for use in biological studies and the supply-chain industry. Authors of [14] propose a rapid application development framework based on a dynamic workflow engine for creating mobile web services.

Several researches has been carried out in workflow management issues in Mobile *Social Content Sharing applications* [6] [11] [18]. A mobile P2P social content sharing framework was proposed in [6]. In [18], a Java API based mobile workflow system was proposed. A content distribution protocol was proposed in [13] for vehicular ad hoc networks (VANET). Clusters of mobile devices has been proposed in [22] to support the execution of parallel applications.

The common approach towards an allocation problems often model the problem as a linear program [7] [19] [21]. A linear program is suitable for modelling situations where communication time is not considered or when there are only two devices involved in the process. However, in the cases of mobile workflows, communication tasks are an essential part of the workload and occurs significant

¹ E.g. we can assume that the support system is similar to that of an Excel application with embedded VBA macro modules. Data is stored in the local spreadsheet.

amount of energy cost[20]. Thus we construct our model as a quadratic program in order to accurately capture the communication costs.

Several recent researches has developed methods to measure the energy cost of mobile applications [8] [16] [20]. The difference in the current draw between sender and receiver in a wireless network can be read at [9]. [3] includes a detailed characteristics of a WiFi network's energy pattern. Our energy model draws ideas from these researches.

3 System Model

3.1 Mobile Platform Model

We consider a mobile platform MP consisting of m mobile devices, M_1, \dots, M_m , and denote a device profile as $M_i (s_i, e_i^{cmp}, e_i^{snd}, e_i^{rcv}), i \in \{1, 2, \dots, m\}$ with parameters defined as follows:

- s_i Peak processing speed of M_i , measured in the number of clock cycles available in a millisecond;
- e_i^{cmp} Current draw from the battery when the device is executing computation tasks at peak speed;
- $e_i^{snd/rcv}$ Current draw from the battery when the device is sending/receiving data to/from the data network.

These devices are interconnected via a network, and we use b_{ij} to denote the bandwidth between devices M_i and $M_j, i, j \in \{1, 2, \dots, m\}$. Thus, we have an m -matrix $B = (b_{ij})_{m \times m}$ which holds all of the bandwidth information of the underlying network of the MP . When two adjacent tasks are assigned to the same device, we assume that they share the same memory address space on the device. Therefore, we assign positive infinite values to the principal diagonal elements of B , that is $b_{ii} = +\infty, i \in \{1, 2, \dots, m\}$.

3.2 Workflow Model

The workflow hosted on MP is represented by a directed acyclic graph $W = (T, R)$ whose vertex set $T = \{t_1, \dots, t_n\}$ denotes the set of *tasks* of the workflow. We assume that all tasks are defined via a service-oriented architecture and that all services are available from each device. An n -matrix $D = (d_{a,b})_{n \times n}$ denotes the weighted adjacency matrix of W , where $d_{a,b}$ is the size of the data package that is to be sent from t_a to t_b for $(t_a, t_b) \in R$. The acyclic property of W implies that D has all principle diagonal elements zero.

Each task has profile $t_a (d_{(\cdot,a)}, d_{(a,\cdot)}, c_a), a \in \{1, \dots, n\}$ where $d_{(\cdot,a)}$ and $d_{(a,\cdot)}$ are the a -th column and the a -th row of D which represent the incoming and outgoing data respectively. c_a denotes the size/workload of the task.

3.3 Mobile Energy Model

Given an allocation scheme $\psi : T \rightarrow M$, we first derive the energy cost of computing $t_a, a \in \{1, \dots, n\}$ to be

$$\mathcal{E}_{a\psi(a)}^{cmp} = e_{\psi(a)}^{cmp} \times \frac{c_a}{s_{\psi(a)}} \quad (1)$$

where $\psi(a)$ is the device to which t_a is assigned. Secondly, we have the energy cost of transferring $d_{ab}, (t_a, t_b) \in R$ as

$$\mathcal{E}_{ab\psi(a)\psi(b)}^{tran} = \underbrace{e_{\psi(a)}^{snd} \times \frac{d_{ab}}{b_{\psi(a)\psi(b)}}}_{\text{sender's cost}} + \underbrace{e_{\psi(b)}^{rcv} \times \frac{d_{ab}}{b_{\psi(a)\psi(b)}}}_{\text{receiver's cost}} \quad (2)$$

4 Minimum Group Energy Cost Problem (MGECP)

In this section, we first show that the Minimum Group Energy Cost Problem can be modelled as a generalised Quadratic Assignment Problem (QAP) [5] and then we convexify the objective function in order to solve it using a MIQP solver.

To represent an allocation scheme ψ , we first construct an $n \times m$ binary matrix $X = (x_{ai})$, such that

$$x_{ai} = \begin{cases} 1 & \text{if } \psi(a) = i, \\ 0 & \text{otherwise.} \end{cases} \quad (3)$$

We call matrix X an *assignment matrix* and a valid assignment must satisfy the following constraints

$$\sum_{i=1}^m x_{ai} = 1, \quad a = 1, 2, \dots, n, \quad (4)$$

$$x_{ai} \in \{0, 1\}, \quad a = 1, 2, \dots, n, \quad i = 1, 2, \dots, m. \quad (5)$$

(4) ensures that every task must be assigned to one and only one device. (5) states that all tasks are indivisible.

4.1 Quadratic Program Formulation

With (1) (2) and (3), we can derive the total energy cost function as

$$\sum_{b=1}^n \sum_{j=1}^m \sum_{a=1}^n \sum_{i=1}^m (e_i^{snd} + e_j^{rcv}) \frac{d_{ab}}{b_{ij}} x_{ai} x_{bj} + \sum_{a=1}^n \sum_{i=1}^m e_i^{cmp} \frac{c_a}{s_i} x_{ai} \quad (6)$$

The quadratic terms in (6) gives the total energy cost for data transmission, whereas the linear term gives the total energy cost for executing computing tasks. We introduce $(nm)^2$ coefficients q_{aibj}

$$q_{aibj} := \begin{cases} e_i^{cmp} \frac{c_a}{s_i} + (e_i^{snd} + e_j^{rcv}) \frac{d_{ab}}{b_{ij}} & \text{if } (a, i) = (b, j), \\ e_i^{snd} \frac{d_{ab}}{b_{ij}} & a < b \\ e_i^{rcv} \frac{d_{ba}}{b_{ij}} & a > b \end{cases} \quad (7)$$

and with (7) we can transform (6) to

$$\sum_{b=1}^n \sum_{j=1}^m \sum_{a=1}^n \sum_{i=1}^m q_{aibj} x_{ai} x_{bj} \quad (8)$$

Theorem 1. *Let coefficients q_{aibj} be the entries of an $mn \times mn$ matrix Q , such that q_{aibj} is on row $(i-1)n + a$ and column $(j-1)n + b$, and $x = \text{vec}(X) = (x_{11}, x_{12}, \dots, x_{1n}, x_{21}, \dots, x_{mn})^T$ be the vector formed from the columns of X . Equivalent formulations for the minimum workflow energy cost problem's objective function are given by (8) and*

$$\text{vec}(X)^T Q \text{vec}(X) \quad (9)$$

Proof. From the construction of $\text{vec}(X)$, we observe that its u -th element $\text{vec}(X)_u = x_{ai} \Leftrightarrow u = (i-1)n + a$. Furthermore, given $u = (i-1)n + a$ and $v = (j-1)n + b$, $u, v \in \{1, 2, \dots, mn\}$, we also get $Q_{uv} = q_{aibj}$. Hence,

$$\begin{aligned} (8) &= \sum_{v=1}^{mn} \sum_{u=1}^{mn} \text{vec}(X)_u^T Q_{uv} \text{vec}(X)_v \\ &= \sum_{b=1}^n \sum_{j=1}^m \sum_{a=1}^n \sum_{i=1}^m x_{ai} q_{aibj} x_{bj} = (9) \end{aligned}$$

□

4.2 Convexification

In order to exploit the power of modern MIQP solvers, we first need to preprocess the problem and convexify the objective function [4]. There are a number of ways of convexification. Our process is similar to that use in [2].

Theorem 2. *Let $Q^* := 1/2(Q + Q^T) + \alpha I$, where I is the $mn \times mn$ identity matrix, then Q^* is positive definite if scalar $\alpha = 1 + \|Q\|_\infty$*

Proof. Due to the length of this paper, interested reader are referred to the appendix of [2] (on a negative definite matrix) for a similar proof. □

Addition of a constant on the main diagonal of Q only add a constant to (9) which does not change its optimal solution. Hence we can rewrite our objective function as

$$\min: \text{vec}(X)^T Q^* \text{vec}(X) \quad (10)$$

This together with (4) and (5) completes the formulation of the optimisation problem of MGECP. The positive definite property of Q^* ensures that (10) is strictly convex and a global minimum can be found by an MIQP solver.

5 Minimum Max-Utilisation Problem (MMUP)

While MGECP ensures that a workflow consumes minimum amount of energy from the mobile devices as a group, it does not consider the stress it has on individual devices. This causes unfair energy cost distribution within the MP , and creates *over-utilised* devices. Having such workflow executed repeatedly over time without adjustment to its task allocation scheme could lead to early retirement of the over-utilised devices. In a business environment, it is common to have authorisation constrained tasks taking critical roles within workflows. In such cases, the MP 's inclusion of these authorised devices is critical to the fulfilment of the workflow's functionality. This requires the workflow engine to shift its priority from reducing the total energy cost of the device group to ensuring the availability of key devices.

Hence in this section of the paper, we investigate ways to adjust the task allocation provided by the MGECP so that the availability period of a workflow can be lengthened. We refer to this class of problem as the Minimum Max-Utilisation Problem (MMUP). We first introduce the measure of utilisation:

Definition 1. Given an allocation scheme ψ , the utilisation of M_i , denoted U_i^ψ , equals $\mathcal{E}_i^\psi / \mathcal{E}_i^R$, for $1 \leq i \leq m$, where \mathcal{E}_i^ψ is the energy cost of M_i under ψ and \mathcal{E}_i^R is the size of the residual energy in M_i .

The reciprocal of a device's utilisation, $(1/U_i)$, is the number of times the workflow can run with M_i before it runs out of battery. The availability of a workflow is hence constrained by the member with the highest value of utilisation. As illustrated in Figure 2, we introduce a *guide utilisation* value U^G to classify devices into two groups: *Over-Utilised (OU)* and *Under-Utilised (UU)*. The objective then is to shift workload from devices in the OU group to those in UU.

We present two adjustment methods, both utilising the quadratic program formed in MGECP and use the result it produces to apply tight constraints to both methods' variables so that the group's overall energy cost remains minimised to a good degree. Upon need, or periodically, the workflow engine executes the adjustment algorithm in order to map the workflow to an updated task allocation scheme so that no device is over stressed unnecessarily and thus improve the availability of the workflow.

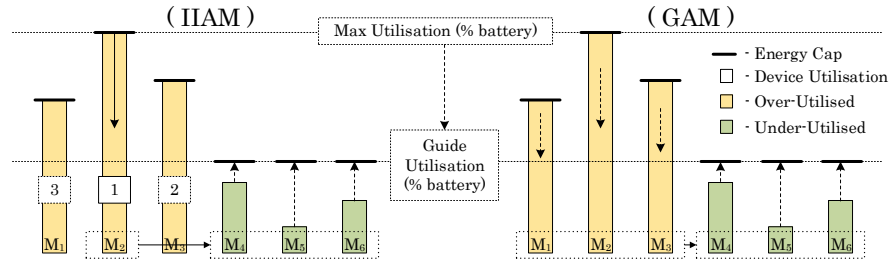


Fig. 2: MMUP Adjustment Algorithms

For both adjustment methods, in order to constraint each device's energy cost, we introduce a device specific cost matrix Q^i as an addition to the quadratic program formulated in MGECP.

Theorem 3. *Let*

$$Q_{uv}^i = \begin{cases} Q_{uv} & \text{if } n \times (i - 1) < u \leq n \times i, \\ 0 & \text{otherwise.} \end{cases} \quad (11)$$

Then given an allocation scheme ψ and its allocation matrix X^ψ , we have the energy cost of M_i to be $\mathcal{E}_i^\psi = \text{vec}(X^\psi)^T Q^i \text{vec}(X^\psi)$

Proof. Proof is similarly to that of Theorem 1 and can be worked out easily. \square

5.1 Iterative Individual Adjustment Method (IIAM)

In this method, we aim to reduce the energy cost of devices in OU individually (as illustrated in Figure 2). With (11), we formulate a quadratically constrained quadratic program (QCQP) with an objective function that minimise the energy cost of the device with the highest utilisation value. As constraints in the QCQP, we cap all other OU devices' energy cost to their current value and all UU devices to the guide utilisation value. (For brevity we use the average utilisation of the current allocation scheme as our guide value. This can be replaced with tailored values to suit the requirement of certain workflows).

If the solver returns a new allocation, we then update the OU and UU group and again select the highest utilised device to the objective function. If this device is same to the one we picked at the earlier iteration, this means that we have reached the optimum solution under the constraints and exit. Otherwise, we repeat the process with the updated group classification until no new device can be picked from the OU group and provide a new allocation.

The advantage of this method is that it pin-points the highest utilised device of the MP , and support its workload offload with the entirety of UU devices. The disadvantage of this method is that all other members of the OU group is capped at their current utilisation value, this restrains the workload offload on the objective device when communication tasks exist between these devices.

5.2 Group Adjustment Method (GAM)

Similar to IIAM, our second adjustment method also caps the contribution from the UU devices at the guide value which ensures that the relocation process does not produce a new OU device (as illustrated in Figure 2). It also apply cap to all OU devices to their current value. Unlike IIAM, the objective function in GAM include all devices in the OU group and also do not iterate.

The advantage of GAM over IIAM is that the workload offload is done between two groups of devices and thus encourages workload offload between members of each group. This in turn increase the possibility of producing a new allocation. The main disadvantage is that it does not prioritise on reducing the maximum utilisation which the workflow's availability is limited to.

6 Simulation

6.1 Environmental Settings

While it is intractable to cover all possible use cases of mobile workflows, we aim to base our simulation closely to the characteristics of an average mobile application and a modern smart device. We construct our simulation with the multiples of two essential building blocks: a *typical device* and a *unit workload*.

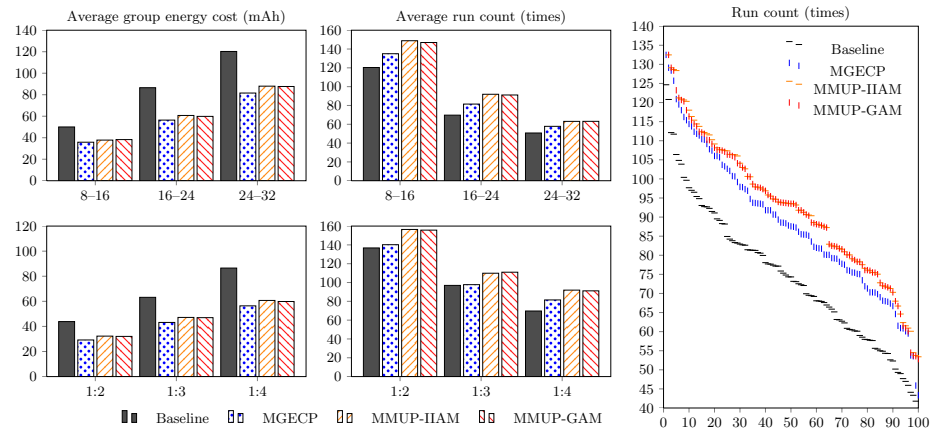
Definition 2. A typical mobile device has a battery capacity of 2000mAh, draws a current of 250-400mA during data transmission and 100-200mA when executing local computation tasks.

In order to accurately emulate the correlation between modern mobile applications and the behaviour and capability of a state-of-the-art smart device, we consult the data presented in recent researches ([8], [16] and [20])

Definition 3. A task has a unit workload if its execution takes 1 second to complete on a typical device.

In our simulation, we specify each task’s workload size using multiples of a unit workload. For instance, in the first plot of figure 3a, the tests are in 3 groups and the workflow generated in each test group has a task size that ranges in from 8 to 16, 16 to 24 and 24-32 units of a unit workload.

Apart from task size, many other factors (e.g. network bandwidth, etc.) also affects the energy cost of a workflow. Due to the length of this paper, we select to present the effect of different device to task ratios in our simulation settings to further verify our model and adjustment methods, as shown in the latter two plots of figure 3a. We also present the effect of different workload size distribution pattern in table 1. For each simulation setting 100 instances are randomly generated and worked on. Averages are taken for comparison. We use AMPL and CPLEX 12.5’s MIQP solver to solve the formulated problems.



(a) Lower group energy cost and higher run count. (b) Sorted run counts
 Fig. 3: Reduction in group energy cost and increase in workflow run count

Table 1: Comparison of adjustment methods

Workloads [‡]	Tests	MMUP-IIAM (MGECP)			MMUP-GAM (MGECP)				
		No.	Max. [†]	Avg. [†]	Dev. [†]	No.	Max. [†]	Avg. [†]	Dev. [†]
Exp. 6	100	35	0.311(0.332)	0.210(0.192)	0.069(0.103)	29	0.295(0.320)	0.204(0.191)	0.062(0.098)
Uni. 4-8	100	41	0.329(0.360)	0.229(0.214)	0.068(0.105)	38	0.334(0.367)	0.230(0.214)	0.068(0.106)
Exp. 8	100	33	0.399(0.436)	0.266(0.242)	0.092(0.141)	30	0.403(0.437)	0.266(0.244)	0.093(0.141)
Uni. 4-12	100	45	0.440(0.486)	0.308(0.287)	0.087(0.140)	42	0.440(0.488)	0.306(0.289)	0.087(0.140)
Exp. 12	100	29	0.651(0.728)	0.425(0.391)	0.151(0.238)	31	0.628(0.706)	0.420(0.391)	0.137(0.225)
Uni. 4-20	100	44	0.640(0.716)	0.464(0.427)	0.124(0.221)	42	0.646(0.727)	0.452(0.424)	0.131(0.226)
Exp. 20	100	26	0.987(1.050)	0.640(0.598)	0.243(0.329)	26	1.014(1.082)	0.651(0.601)	0.250(0.349)
Uni. 4-36	100	49	1.113(1.244)	0.773(0.714)	0.234(0.378)	50	1.097(1.202)	0.754(0.701)	0.226(0.365)

[†] - All utility values are percentages (%) of residual battery (mAh). [‡] - Distribution and task size.

6.2 Results and Analysis

Minimum Total Energy Cost The first group of our simulations aims to verify the formulation of MGECP. As a comparison, we use a baseline algorithm which attempts to reduce the total energy cost by distributing the number of tasks evenly across the *MP*. This algorithm provides a good baseline value because although it does not seek the benefit of using an energy efficient device, its chance of being able to take that advantage is consistent.

As shown in figure 3a, the total energy cost of a baseline allocation is reduced by 30-35% with MGECP applied. Both adjustment algorithms are applied to allocation produced by MGECP. As shown in figure 3a, the adjustments does not significantly increase the total energy cost of the workflow, but boost the workflow’s run count (the number of time the workflow can run before the first retirement from the *MP*). As discussed with the example illustrated in figure 1b, the fair distribution of workload amongst the *MP* is critical. One series of simulation (16-24 in figure 3a) is magnified and plotted in figure 3b to illustrate the effect our algorithms have in extending the run count of workflows.

Utilisation Adjustment This group of simulations focuses on the adjustment algorithms and their effect on workflows with different task workload range and distribution. Results (cf. table 1) show that first of all, not all MGECP allocations can be adjusted because of the tight constraints we apply in both IIAM and GAM. 26-50% of the 100 test instances generated in each setting can be adjusted in order to gain a lower maximum utility.

It is worth noting that workflows with uniformly distributed task sizes have better chance to be adjusted than those with exponential distribution pattern. Adjustments can be made when tasks can be offloaded or exchanged between devices without causing sizeable disturbance in each devices’ energy cost. Contrary to that of exponential distribution, when workload’s distribution is uniform within a set range, it is likely for a task to find another task that has similar workload size, thus a “minor” exchange of tasks is more likely to exist.

Device Energy Cost This group of simulation focuses on the energy cost of individual devices. Figure 4 shows that in order to reduce the maximum utilisation, tasks has to be offloaded or exchanged to a device where it will cost more energy to execute which increase the average utilisation of the group.

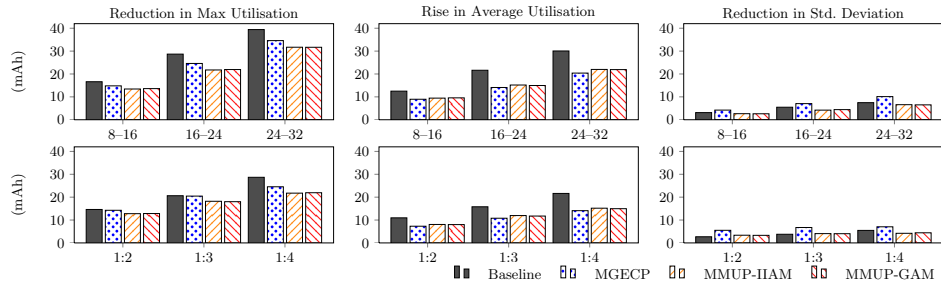


Fig. 4: Effect of adjustments within the *MP*.

Increase in the group’s standard deviation caused by MGECP shows that in order to minimise the group’s energy cost, devices with better energy-efficiency are required to take on more workload than the others. On the other hand, the reductions of this value from MGECP to MMUP show that the workflow’s energy cost is distributed more evenly within the *MP* after adjustments.

7 Conclusion

In this paper, we introduced a model that captures both computation and communication costs of a workflow with a 0-1 quadratic program. We demonstrated the use of MIQP solvers which produces exact solutions for the MGECP. We also investigated ways to adjust the allocation to lengthen the workflow’s availability with minimal impact on its overall energy cost. Our simulation produces good results for both problems and gives an insight into workflows of different characteristics. Our model is also applicable to other energy critical scenarios, its extension can be tailored for workflows of specific use cases.

Acknowledgement

This work is sponsored by the Research Project Grant of the Leverhulme Trust (Grant No. RPG-101).

References

1. Balasooriya, J., Joshi, J., Prasad, S.K., Navathe, S.: Distributed Coordination of Workflows over Web Services and Their Handheld-Based Execution. In: ICDCN’08 Proceedings of the 9th International Conference on Distributed Computing and Networking. pp. 39–53 (2008)
2. Bazaraa, M.S., Sherali, H.D.: On the Use of Exact and Heuristic Cutting Plane Methods for the Quadratic Assignment Problem. *Journal of the Operational Research Society* 33(11), 991–1003 (1982)
3. Bejerano, Y., Han, S.J., Li, L.E.: Fairness and Load Balancing in Wireless LANs Using Association Control. In: *MobiCom’04 Proceedings of the 10th Annual International Conference on Mobile Computing and Networking*. p. 315 (2004)

4. Billionnet, A., Elloumi, S.: Using a Mixed Integer Quadratic Programming Solver for the Unconstrained Quadratic 0-1 Problem. *Mathematical Programming* 109(1), 55–68 (2006)
5. Burkard, R.E., Pitsoulis, L.S., Linearization, J., Polytopes, Q.A.P.: The Quadratic Assignment Problem. In: *Handbook of Combinatorial Optimization* (1998)
6. Chang, C., Srirama, S.N., Ling, S.: An Adaptive Mediation Framework for Mobile P2P Social Content Sharing. In: *ICSOC'12 International Conference on Service-Oriented Computing. Lecture Notes in Computer Science* (2012)
7. Cuervo, E., Balasubramanian, A., Cho, D.k., Wolman, A., Saroiu, S., Chandra, R., Bahl, P.: MAUI: Making Smartphones Last Longer with Code Offload. In: *MobiSys'10 The 8th International Conference on Mobile Systems, Applications, and Services* (2010)
8. Dong, M., Zhong, L.: Self-Constructive High-Rate System Energy Modeling for Battery-Powered Mobile Systems. In: *MobiSys'11 The 9th International Conference on Mobile systems, Applications, and Services* (2011)
9. Feeney, L., Nilsson, M.: Investigating the Energy Consumption of a Wireless Network Interface in an Ad Hoc Networking Environment. In: *INFOCOM'01. Conference on Computer Communications. Twentieth Annual Joint Conference of the IEEE Computer and Communications Society* (2001)
10. Gartner Research: Gartner Reveals Top Predictions for IT Organizations and Users for 2012 and Beyond (2011), <http://www.gartner.com/it/page.jsp?id=1862714>
11. Huang, C.M., Hsu, T.H., Hsu, M.F.: Network-Aware P2P File Sharing over Wireless Mobile Networks. *IEEE Journal on Selected Areas in Communications* (2007)
12. Key Lime 314 LLC: KL Dartboard (2011), <https://itunes.apple.com/gb/app/kl-dartboard/id376234917?mt=8>
13. Lee, U., Park, J.S., Yeh, J., Pau, G., Gerla, M.: CodeTorrent: Content Distribution using Network Coding in VANET. In: *MobiShare'06 1st International Workshop on Decentralized Resource Sharing in Mobile Computing and Networking*(2006)
14. Mnaoue, A., Shekhar, A.: A Generic Framework for Rapid Application Development of Mobile Web Services with Dynamic Workflow Management. In: *SCC'04 IEEE International Conference on Services Computing* (2004)
15. Pajunen, L., Chande, S.: Developing Workflow Engine for Mobile Devices. In: *EDOC'07 11th IEEE International Enterprise Distributed Object Computing Conference* (2007)
16. Pathak, A., Hu, Y.C., Zhang, M.: Where is the energy spent inside my app? Fine Grained Energy Accounting on Smartphones with Eprof. In: *EuroSys'12 7th ACM european conference on Computer Systems*. ACM Press (2012)
17. Pentikousis, K.: In Search of Energy-Efficient Mobile Networking. *IEEE Communications Magazine* 48(1), 95–103 (2010)
18. Philips, E., Carreton, A.L., Joncheere, N., De Meuter, W., Jonckers, V.: Orchestrating Nomadic Mashups using Workflows. In: *Mashups '09/'10 The 3rd and 4th International Workshop on Web APIs and Services Mashups* (2010)
19. Shachnai, H., Tamir, T.: On Two Class-Constrained Versions of the Multiple Knapsack Problem. *Algorithmica* 29(3), 442–467 (2001)
20. Sharkey, J.: Coding for life - Battery Life, that is. *Google IO Conference* (2009)
21. Tang, C., Steinder, M., Spreitzer, M., Pacifici, G.: A Scalable Application Placement Controller for Enterprise Data Centers. In: *WWW'07 The 16th International Conference on World Wide Web* (2007)
22. Zong, Z., Nijim, M., Manzanares, A., Qin, X.: Energy Efficient Scheduling for Parallel Applications on Mobile Clusters. *Cluster Computing* 11(1), 91–113 (2007)