

AN ECONOMICS-BASED POWER-AWARE PROTOCOL FOR COMPUTATION DISTRIBUTION IN MOBILE AD-HOC NETWORKS[‡]

Li Shang, Robert P. Dick, and Niraj K. Jha
Department of Electrical Engineering
Princeton University, Princeton, NJ 08544, USA
{lshang,dickrp,jha}@ee.Princeton.edu

ABSTRACT

In this paper, we present a new economics-based power-aware protocol, called the *distributed economic subcontracting protocol*, that dynamically distributes task computation among mobile devices in an ad-hoc wireless network. Mobile computation devices may be energy buyers, contractors, or subcontractors. Tasks are transferred between devices via distributed bargaining and transactions. When additional energy is required, buyers and contractors negotiate energy prices within their local markets. Contractors and subcontractors spend communication and computation energy to relay or execute buyers' tasks. Buyers pay the negotiated price for this energy. Our experimental results indicate that markets based on our protocol and decision-making algorithms fairly and effectively allocate energy resources among different tasks in both cooperative and competitive scenarios.

KEY WORDS

Economics-based, Power-aware, Ad-hoc networks

1. INTRODUCTION

In ad-hoc wireless networks, mobile computation devices are usually battery-powered. A limited energy budget constrains the computation and communication capacity of each device. Energy resources and computation workloads have different distributions within the network. Some mobile devices have spare energy. Devices that expend all their energy can only be recharged when they leave the network. Therefore, it is beneficial to redistribute spare energy resources to satisfy unevenly distributed workloads in the network. In this paper, we propose a protocol for computation distribution that solves this dynamic energy resource allocation problem.

In wireless networks, the ratio of computation energy consumption to communication energy consumption

varies in a wide range, depending on application type. In some applications, e.g., micro-sensor networks, communication dominates energy consumption [1]. In other application domains and applications, e.g., simulation, classification, artificial intelligence, target detection, handwriting recognition, and voice recognition, computation energy consumption generally dominates communication energy consumption [2]-[6]. If devices with excess computation-intensive tasks can, for a fee, transfer these tasks to devices with spare energy and time, both buyer and seller devices benefit; sellers may use their earnings to buy energy in the future.

In distributed computing systems, economics-based techniques have been used to balance resource allocation and utilization [7]-[11]. Our *distributed economic subcontracting protocol* (DESP) dynamically distributes task computation among mobile or fixed-position devices in an ad-hoc network. This is the first work to propose a power-aware market-based computation distribution protocol. Tasks that may be transferred between devices for less energy than is required for local computation are distributed via on-line bargaining within each local market. Sellers may be contractors or subcontractors. They automatically adjust their energy prices based upon market conditions. Local market sizes are dynamically adjusted in order to balance communication energy and the lowest prices available to buyers. DESP supports a new class of economic agents, called subcontractors. Subcontracting allows transitive transfers of task execution among devices; subcontractors tie local markets together into a global market. Subcontracting can be seen as a computational version of multihop communication. We propose policies to handle both cooperative and competitive scenarios. We believe this first study demonstrates the effectiveness of economics-based approach as an energy resource allocation mechanism for ad-hoc mobile networks.

The rest of this paper is organized as follows. In Section 2, we present related concepts and formalize our problem definition. In Section 3, we introduce the economics-based protocol in detail. We experimentally demonstrate the feasibility of our approach in Section 4. Finally, we conclude in Section 5.

[‡]Acknowledgments: This work was supported by DARPA under contract no. DAAB07-02-C-P302.

2. PROTOCOL OVERVIEW

DESP performs dynamic allocation of energy resources in ad-hoc wireless mobile networks, through on-line transactions within markets. As illustrated in Figure 1, DESP consists of the following elements:

Buyers: A device that intends to purchase energy from other devices is a buyer. A buyer uses an advertising broadcast to construct a local market, in which it may purchase energy.

Sellers: A device that is willing to sell spare energy to other devices joins one or more local markets as a seller.

Contractors: In a local market, sellers compete with each other. The winner signs a contract with the buyer: it is a contractor. A contractor may decide to execute a buyer's task. However, it may decide to create another local market to find subcontractors. A contractor is a seller. However, if it uses a subcontractor, it is also a buyer.

Subcontractors: A subcontractor is a contractor that sells to another contractor or subcontractor, instead of selling directly to a buyer.

Local market: Every energy transaction occurs within a local market. Each local market is dynamically constructed by a market owner that may be a buyer, contractor, or subcontractor. The market owner's advertising broadcast energy controls the market's area. Multiple sellers within the local market send out their encrypted offers to the market owner, which chooses the winner and signs a contract.

We use the wireless communication path loss model to calculate transmission energy consumption [12,13]. In this model, the received signal power is proportional to $1/d^n$, where d is the transmission distance, and n is a path loss exponent dependent on the environment. In our model, different mobile devices are assumed to be equipped with low-power global positioning system (GPS) receivers to provide position information.

3. THE SUBCONTRACTOR MARKET

In this section, we explain the transaction protocols and corresponding optimization algorithms for the economic agents in our protocol.

3.1 Transaction Protocols

In DESP, there are energy transactions between buyers and sellers. Each device bases its judgment about market conditions on a history of its recent transactions.

3.1.1 Transaction Protocol for the Buyer Market

Figure 2 shows the buyer transaction protocol. First, the buyer analyzes its pending tasks, remaining energy, remaining money, and transaction history. Based on this information, it decides whether to execute a pending task or become a buyer and pay other devices to execute the task for it. A buyer makes an advertising broadcast to

construct a local market. Advertising broadcast energy controls advertising range and, thereby, market area. The sellers, within the buyer's local market, may make bids. The buyer accepts offers until its bid deadline. After the bid deadline, and before the decision deadline, the buyer may choose one of the bids it has received and send out an acceptance message. It then signs a contract with the corresponding seller, thereby changing the seller to a contractor. Finally, the buyer sends its tasks to the contractor, receives the computation results, and pays the contractor. At the end of the transaction, the local market automatically closes.

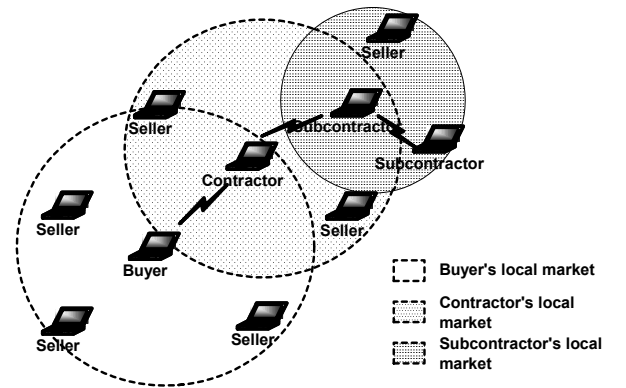


Figure 1: DESP example

Figure 3 shows the seller transaction protocol. First, a device that is willing to sell energy becomes a seller and begins to monitor the advertising channel. If a seller receives an advertisement, it analyzes the incoming task, its energy budget, and transaction history. Based on this information, the seller sends back its bid, including price and position information. It then waits for the buyer's decision until the buyer's decision deadline. If the seller's offer is not accepted by this time, it assumes the offer is rejected, and the transaction is closed. If, instead, its offer is accepted, it signs a contract and receives the task from the buyer, thereby becoming a contractor. This contractor may decide to construct another, overlapping, local market to find a subcontractor. After the resulting data have been computed, either by the contractor or by a subcontractor, the contractor sends them to the buyer. Finally, the seller receives its payment and pays a subcontractor, if necessary.

3.1.2 Transaction Protocol for the Contractor and Subcontractor Markets

When a seller becomes a contractor, it may construct its own local market to find subcontractors. The contractor transaction protocol is similar to the buyer protocol. In essence, the contractor becomes a relay node between the buyer and the subcontractor, transferring tasks from the buyer to the subcontractor and returning the results. For this work, the contractor earns the difference between the buyer's payment and the subcontractor's bid. This

protocol allows contractors and subcontractors to cooperate in providing resources to a buyer and share the buyer's payment.

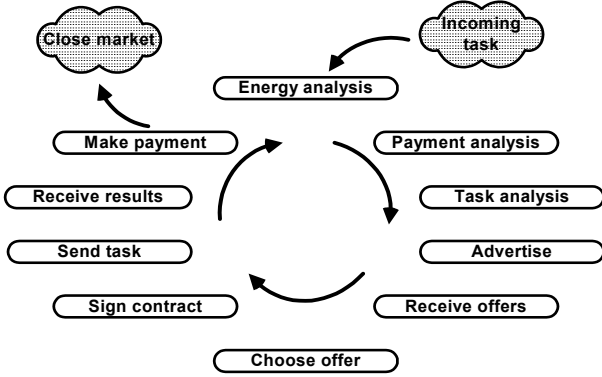


Figure 2: Transaction protocol for buyers

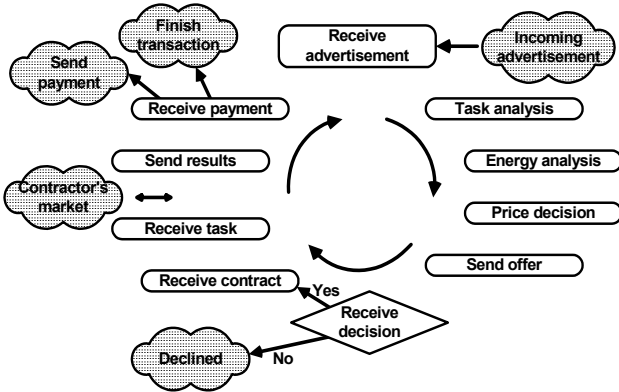


Figure 3: Transaction protocol for sellers

3.2 Transaction Policies for Buyers

In the absence of a central controller, mobile devices must make their own energy purchasing decisions. In DESP, buyers do local advertising broadcasts. A buyer may only carry out direct transactions with sellers in its advertising area. It is desirable to reduce communication energy and price. However, these costs conflict with each other, i.e., it is often possible to decrease one only by increasing the other. Communication energy is the energy expended by a buyer during advertising broadcast and task transmission for remote computation. It is correlated with advertising broadcast area. Price is correlated with the energy scarcity of the available seller devices. In other words, increasing the number of sellers in a market will, on average, reduce the minimum price available. It is necessary to decide upon a broadcast range that results in a good trade-off between price and communication energy. In our protocols, buyers dynamically adjust their advertising distance in the following way:

1. For task k , the buyer calculates an upper bound on communication distance D_k , subject to the constraint that communication energy is lower than

computation energy. The buyer also predicts the communication distance lower bound d_n , based on previous successful transactions. If $D_k < d_n$, then the buyer executes the task locally. Otherwise, it proceeds to step 2.

2. If the last transaction succeeded, the buyer multiplies the advertising range by a user-defined constant, e.g., 0.9. Otherwise, the advertising range is similarly increased, under the constraint that advertising range is less than D_k .
3. Periodically, the buyer doubles its broadcast distance to probe for a better offer, under the constraint that the advertising range is less than D_k .

Each bid has two costs, price and communication energy. Therefore, buyers need to choose a bidder that offers the best trade-off between price and communication energy. From all the received offers, buyers calculate the average unit energy price, p_e , in the transaction history. For each offer, they calculate the equivalent total price P_j^* and choose the offer with the lowest equivalent price based on the following equation:

$$P_j^* = p_e \times E_{comm_j} + P_j$$

where E_{comm_j} is the communication energy for offer j , and P_j is offer j 's price.

3.3 Transaction Policies for Sellers

Multiple sellers may exist within a local market, each competing to maximize its own optimization criterion. In this subsection, two optimization criteria are proposed: one for competitive scenarios and one for cooperative scenarios.

3.3.1 Competitive Sellers

In competitive ad-hoc mobile networks, sellers have the goal of maximizing their total profits subject to their energy budgets and lifetime constraints. The mobile network is a dynamic system; guaranteeing optimal profit is a hard problem. In reality, each device has only imperfect information and must base its predictions on its recent transaction history. Therefore, in this work, we use an incremental greedy derivative-following strategy to maximize profit.




Figure 4: (a) Competitive optimization curve, and (b) adaptive step-size transformation function

We assume that the utility function is always concave, as shown in Figure 4(a). Intuitively, initial increases in price do not substantially reduce sales, allowing an increase in profit. Eventually, price increases result in a significant reduction in sales, reducing total profit. The point between these regions is the unit price resulting in maximal profit. Marginal utility is equivalent to the profit gradient, which is positive at the beginning and non-increasing. Maximum profit is achieved when the profit gradient is zero, i.e., given that $E_i(t)$ is the remaining energy at time t , $x_i(t)$ is seller i 's unit energy price for the transaction at time t , $f_i(x_i(t), t)$ is the energy consumption rate at time t , $cost_i$ is the unit cost of device i 's energy, and $T_i(t)$ is device i 's remaining time in this network:

$$profit_{\max} = \max_{x_i(t)} \left\{ (x_i(t) - cost_i) \times f_i(x_i(t), t) \times \min \left\{ \frac{E_i(t)}{f_i(x_i(t), t)}, T_i(t) \right\} \right\}$$

We define the equivalent lifetime $T_i(t)^*$ of device i as follows:

$$T_i(t)^* = \min \left\{ \frac{E_i(t)}{f_i(x_i(t), t)}, T_i(t) \right\}$$

If $T_i(t)^* < T_i(t)$, it implies that, given the current energy consumption rate, device i will use all its spare energy before it leaves the network.

The incremental greedy derivative-following algorithm has the following properties. It does boundary checks to guarantee that the bid price is higher than the energy cost. It increases or decreases its unit energy price if, based on its transaction history, this is expected to increase profit. After arriving at a stable unit energy price, it dynamically probes and adapts to changing market conditions.

We use an adaptive step-size strategy to change the seller's unit energy price:

$$price_{j+1} = price_j + price_j \times \text{sign}(price_j - price_{j-1}) \times g \left(\frac{profit_j}{profit_{j-1}} \right)$$

where $price_{j+1}$ is the predicted unit energy price to be used in next transaction $j+1$, $price_j$ and $price_{j-1}$ are unit energy price estimates, and $profit_j$ and $profit_{j-1}$ are profit estimates. These estimates are based on the transaction history. Function $\text{sign}(x) = -1$ if x is negative, otherwise $\text{sign}(x) = +1$.

We use a transformation function, $g(u)$, as shown in Figure 4(b), to dynamically adapt the step-size. Our strategy ensures that, when the change in profit is small, the change in unit energy price is also small. To ensure stability, we bound changes to unit energy price during rapid profit changes.

3.3.2 Cooperative Sellers

In a *fair* market, a rational decision maker receives a quantity of service proportional to the amount of money it spends. DESP can be tailored to optimize fairness. Within a wireless market, energy price is determined by energy supply and demand. An increase in demand, relative to

supply, increases price. Therefore, market price can be used to regulate buyer policy. A low price indicates that more energy is available; buyers react by migrating more tasks to sellers. A high price indicates that less energy is available. Therefore, buyers can only afford to buy energy for most important tasks and delay or drop others.

In the cooperative market scenario, a seller adjusts its price to finish expending its energy at the moment it exits the network, instead of attempting to maximize its total profit. The seller dynamically adjusts its price to maintain an energy consumption rate E_r , defined as E_i/T_i , where E_i is its remaining energy and T_i is its lifetime, as shown in Figure 5. A cooperative seller attempts to provide energy to buyers at a constant rate. This stability helps the seller achieve fairness. In addition to changing its bid price, a seller reacts to a change in its energy consumption rate by appropriately adjusting the bid price it will tolerate from subcontractors.

We use an incremental greedy goal-directed strategy for energy resource allocation. Each seller decides its pricing policy based on the following algorithm.

1. During each transaction, this algorithm performs a boundary check to guarantee that the offer price is higher than the monetary cost of carrying out the necessary transactions.
2. Compute the recent energy consumption rate based on the transaction history. Use this rate as a predictor for future market conditions. If this energy consumption rate is higher (lower) than E_r , increase (decrease) the unit energy price.




Figure 5: Energy consumption rate

We use an adaptive step-size strategy to change the seller's unit energy price. Given that $price_{j+1}$ is the predicted unit energy price to be used in next transaction $j+1$, $price_j$ is a unit energy price estimate based on the transaction history, E_{rate_j} is the energy consumption rate from the transaction history, E_{rem_j} is the remaining energy, and T_{rem_j} is the remaining lifetime:

$$price_{j+1} = price_j + price_j \times g \left(\frac{E_{rate_j} \times T_{rem_j}}{E_{rem_j}} \right)$$

We dynamically adjust the step size with the same transformation function, $g(u)$, described in Section 3.3.1.

3.4 Transaction Policies for Contractors

In the competitive scenario, the contractor tries to maximize its total profit subject to its lifetime and energy budget constraints. Although collaboration requires the contractor to share the buyer's payment with a subcontractor, subcontracting may allow it to reach a higher equivalent unit price, than by executing every task by itself. The contractor's only cost is the communication energy required to relay the task and computation results. As a result, collaborating with a subcontractor can increase a contractor's equivalent lifetime, allowing a higher profit.

In the cooperative scenario, the contractor's decision is based on the following criteria. If the contractor's current energy consumption rate is higher than E_r , finding a subcontractor may extend its equivalent lifetime. Collaborating with subcontractors provides the additional advantage of making prices and network's workload distribution in the network more homogeneous; local regions, in which the price decided by the balance between supply and demand is extreme, are dispersed.

4. EXPERIMENTAL RESULTS

In this section, we present experimental results to evaluate the performance of DESP.

4.1 Dynamic Pricing of Competitive Sellers

In this subsection, we evaluate the dynamic pricing strategies of competitive sellers in three different market scenarios. The relationship between price and energy demand is a step function. When a seller's price is less than a buyer-defined upper-bound, the energy demand is a positive constant, otherwise, the energy demand drops to zero. Figure 6 contains the simulation results for dynamic pricing of competitive sellers. The simulation period is 3,500 seconds. In this figure, three different market configurations are studied. In the first configuration, the buyer-defined upper-bound on price is a continuous function that decreases from 400 to 50 during the simulation. In the second configuration, the upper-bound on price is a concave function. Its initial value is 50, it increases to 400, and then decreases to 50. In the third configuration, the upper-bound price is a step function that starts at 100, changing to 200 at time 501, 300 at time 1001, 400 at time 1501, 300 at time 2001, 100 at time 2501, and 50 at time 3001.

From the simulation results, it is clear that, in each configuration, sellers using DESP dynamically adjust their prices to reach the buyer-defined upper-bound on price, thereby maximizing their total profits. Similarly, they dynamically adapt their prices to changes in the buyer-defined upper-bound on price. The slight oscillations around the optimal prices result from

continuously probing the market conditions.




Figure 6. Dynamic pricing policy for competitive sellers

4.2 Dynamic Pricing of Cooperative Sellers

In this subsection, we evaluate the dynamic pricing strategies of cooperative sellers. we examine the fairness of energy allocation in this scenario. As described in Section 3.3.2, in a fair market, the quantity of energy that a rational decision maker receives is proportional to the amount of money it spends. We examined the amount of energy allocated to buyers with different monetary budgets. Table 1 shows the network setup. In this table, the *funding ratio* column contains the ratio between the starting money held by three different classes of buyers. The *finished task energy ratios* column shows, for the three classes of buyers, the ratios between the amounts of energy used for task execution. As we can see from the table, cooperative sellers allocate their energy in a manner that approximates their funding ratios, i.e., they achieve fair energy allocation. The deviations of the energy allocation ratios from the funding ratios are caused by numerous factors, some of which are the uneven spatial and temporal distributions of energy as well as the discrete nature of transactions.

Table 1: Fair energy allocation

Funding ratios	Finished task energy ratios		
	3 buyers	30 buyers	100 buyers
1:1:1	1:0.99:1	0.97:1:1	0.97:0.97:1
3:2:1	2.94:1.98:1	2.90:1.97:1	2.90:1.96:1
10:5:1	9.57:4.87:1	9.36:4.75:1	8.98:4.60:1

4.3 Effectiveness in Cooperative Markets

Network *effectiveness* is the proportion of task volume that a network is able to execute. To determine the impact of subcontractors on effectiveness, we consider two scenarios. In the first, subcontracting is allowed; in the second, it is forbidden. In addition, we examine the effect of varying the ratio between computation and communication energy. We simulate an ad-hoc network composed of 100 buyers and 1,000 sellers. The average speed of each device is 5 meters per second. The average distance between neighboring devices is 50 meters. We vary the ratio of computation to communication energy, for devices separated by this average distance, in a range

from 1 to 100.

Figure 7 shows the simulation results for DESP under four different market conditions: advertising distances (Adv.) of 30 m and 50 m, with and without subcontractors (sub.). These results indicate that DESP made a good trade-off between energy demands and communication energy. When the computation energy to communication energy ratio is high, DESP allocates energy resources from sellers outside a buyer's local market. As the ratio decreases, the energy overhead associated with subcontractor collaboration also increases. As a result, subcontractors are used less frequently. When the ratio reaches one, communication energy has the same cost as computation energy. In this case, buying energy from sellers is not beneficial.




Figure 7. Effectiveness of energy allocation (cooperative)

4.4 Effectiveness in Competitive Markets

Figure 8 shows the effectiveness of DESP in competitive markets. In such markets, each seller tries to maximize its total profit. DESP allows better allocation of spare energy resources to buyers, and higher seller profits, than a market without subcontractors because sellers outside a local market are sometimes willing to provide their energy at lower prices than contractors. Therefore, when a contractor's energy level is low, it can increase its profit by collaborating with subcontractors.




Figure 8. Effectiveness of energy allocation (competitive)

5. CONCLUSIONS

This paper has presented a novel economics-based protocol, called DESP, that dynamically allocates energy

resources in ad-hoc wireless mobile networks. DESP is a scalable, distributed approach: it requires no central coordinator. We have provided and analyzed buyer and seller decision strategies for cooperative and competitive scenarios. Experimental results indicate that the DESP fairly and effectively allocates energy resources to devices in wireless ad-hoc mobile networks.

REFERENCES

- [1] J. L. da Silva, Jr. et al., Design methodology for picoradio networks, *Proc. Design, Automation & Test in Europe Conf.*, Mar. 2001, pp. 314-323.
- [2] M. Othman et al., Power conservation strategy for mobile computers using load sharing, *Mobile Computing & Communications Review*, vol. 2, no. 1, pp. 19-26, Jan. 1998.
- [3] U. Kremer, J. Hicks, and J. M. Rehg, A compilation framework for power and energy management on mobile computers, *Proc. Workshop Compilers & Operating Systems for Low Power*, Oct. 2000.
- [4] A. Rudenko, P. Reiher, G. Popek, and G. Kuenning, The remote processing framework for portable computer power saving, *Proc. ACM Symp. Applied Computing*, Feb. 1999, pp. 365-372.
- [5] J. Flinn, D. Narayanan, and M. Satyanarayanan, Self-tuned remote execution for pervasive computing, *Proc. 8th Workshop Hot Topics in Operating Systems*, May 2001, pp. 61-66.
- [6] A. Vahdat, A. Lebeck, and C. S. Ellis, Every Joule is precious: The case for revisiting operating system design for energy efficiency, *Proc. ACM SIGOPS European Workshop*, Sept. 2000, pp. 31-36.
- [7] R. G. Smith, The contract net protocol: High-level communication and control in a distributed problem solver, *IEEE Trans. Computers*, vol. 29, no. 12, pp. 1104-1113, Dec. 1980.
- [8] O. Regev and N. Nisan, The popcorn market – An online market for computational resources, *Proc. Int. Conf. Info. & Comp. Economics*, Oct. 1998, pp. 148-157.
- [9] J. S. Chase et al., Managing energy and server resources in hosting centers, *Proc. ACM Symp. Operating Systems Principles*, Oct. 2001, pp. 103-116.
- [10] V. Shah et al., Power control for wireless data based on utility and pricing, *Proc. Int. Symp. Personal, Indoor, Mobile Radio Communications*, Sept. 1998, pp. 1427-1432.
- [11] A. B. Machkenzie and S. B. Wicker, Game-theory in communications: Motivation, explanation, and application to power control, *Proc. GlobeCom*, Nov. 2001, pp. 821-826.
- [12] T. S. Rappaport, *Wireless Communication: Principles and Practice*, Prentice Hall, Englewood Cliffs, NJ, pp. 69-122, pp. 139-196, 1996.
- [13] V. Rodoplu and T. H. Meng, Minimum energy mobile wireless networks, *IEEE J. Selected Areas Comm.*, vol. 17, no. 8, pp. 1333-1344, Aug. 1999.