

Maximum Data Gathering in Networked Sensor Systems*

Bo Hong and Viktor K. Prasanna

University of Southern California, Los Angeles CA 90089-2562

{bohong, prasanna}@usc.edu

Abstract

We focus on data gathering problems in energy-constrained networked sensor systems. We study store-and-gather problems where data are locally stored at the sensors before the data gathering starts, and continuous sensing and gathering problems that model time critical applications. We show that these problems reduce to maximization of network flow under vertex capacity constraint. This flow problem in turn reduces to a standard network flow problem. We develop a distributed and adaptive algorithm to optimize data gathering. This algorithm leads to a simple protocol that coordinates the sensor nodes in the system. Our approach provides a unified framework to study a variety of data gathering problems in networked sensor systems. The performance of the proposed method is illustrated through simulations.

*Supported by the National Science Foundation under award No. IIS-0330445 and in part by DARPA under contract F33615-02-2-4005.

1 Introduction

State-of-the-art sensors (e.g. Smart Dust [20]) are powered by batteries. Replenishing energy by replacing the batteries is infeasible since the sensors are typically deployed in harsh terrains. Also, the cost of replacing batteries can be prohibitively high. These sensors, which are usually unattended, need to operate over a long period of time after deployment. Energy efficiency is thus critical. Techniques ranging from low power hardware design [2, 15] and energy aware routing [8, 19] to application level optimizations [18, 21] have been proposed to improve energy efficiency of networked sensor systems.

An important application of networked sensor systems is to monitor the environment. Examples of such applications include vehicle tracking and classification in the battle field, patient health monitoring, pollution detection, etc. In these applications, a fundamental operation is to sense the environment and transmit the sensed data to the base station for further processing. In this paper, we study energy efficient data gathering in networked sensor systems from an algorithmic perspective.

Compared with sensing and computation, communication is the most expensive operation (in terms of energy consumption) in the context of data gathering [1]. Generally, data transfers are performed via multi-hop communications where each hop is a short-range communication. This is due to the well known fact that long-distance wireless communication is expensive in terms of both implementation complexity and energy dissipation, especially when using the low-lying antennae and near-ground channels typically found in networked sensor systems. Short-range communication also enables efficient spatial frequency re-use. A challenging problem with multi-hop communications is the efficient transfer of data through the system when the sensors have energy constraints.

Some variations of the problem have been studied recently. In [11], data gathering is assumed to be performed in *rounds* and each sensor can communicate (in a single hop) with the base station and all

other sensors. The total number of rounds is then maximized under a given energy constraint on the sensors. In [14], a non-linear programming formulation is proposed to explore the trade-offs between energy consumed and the transmission rate. It models the radio transmission energy according to Shannon's theorem. In [16], the data gathering problem is formulated as a linear programming problem and a $1 + \omega$ approximation algorithm is proposed. This algorithm further leads to a distributed heuristic.

Our study departs from the above with respect to the problem definition as well as the solution technique. For short-range communications, the difference in the energy consumption between sending and receiving a data packet is almost negligible. We adopt the reasonable approximation that sending a data packet consumes the same amount of energy as receiving a data packet [1]. The study in [14] and [16] differentiate the energy dissipated for sending and receiving data. Although the resulting problem formulations are indeed more accurate than ours, the improvement in accuracy is marginal for short-range communications.

In [11], each sensor generates exactly one data packet per round (a round corresponds to the occurrence of an event in the environment) to be transmitted to the base station. The system is assumed to be fully connected. The study in [11] also considers a very simple model of data aggregation where any sensor can aggregate all the received data packets into a single output data packet. In our system model, each sensor communicates with a limited number of neighbors due to the short range of the communications, resulting in a general graph topology for the system. We study *store-and-gather* problems where data are locally stored on the sensors before the data gathering starts, and *continuous sensing and gathering* problems that models time critical applications. A unified flow optimization formulation is developed for the two classes of problems.

Our focus in this paper is to maximize the throughput or volume of data received by the base station.

Such an optimization objective is abstracted from a wide range of applications in which the base station needs to gather as much information as possible. Some applications proposed for the networked sensor systems may have different optimization objectives. For example, the balanced data transfer problem [6] is formulated as a linear programming problem where a ‘minimum achieved sense rate’ is set for every individual node. In [5], data gathering is considered in the context of energy balance. A distributed protocol is designed to ensure that the average energy dissipation per node is the same throughout the execution of the protocol. However, these issues are not the focus of this paper.

By modeling the energy consumption associated with each send and receive operation, we formulate the data gathering problem as a constrained network flow optimization problem where each node u is associated with a capacity constraint w_u , so that the total amount of flow going through u (incoming plus outgoing flow) does not exceed w_u . We show that such a formulation models a variety of data gathering problems (with energy constraint on the sensor nodes).

The constrained flow problem reduces to the standard network flow problem, which is a classical flow optimization problem. Many efficient algorithms have been developed ([3]) for the standard network flow problem. However, in terms of decentralization and adaptation, these well known algorithms are not suitable for data gathering in networked sensor systems. In this paper, we develop a decentralized and adaptive algorithm for the maximum network flow problem. This algorithm is a modified version of the Push-Relabel algorithm [7]. In contrast to the Push-Relabel algorithm, it is adaptive to changes in the system. It finds the maximum flow in $O(n^2 \cdot |V|^2 \cdot |E|)$ time, where n is the number of adaptation operations, $|V|$ is the number of nodes, and $|E|$ is the number of links.

The above algorithm can be used to solve both store-and-gather problems and continuous sensing and gathering problems. For the continuous sensing and gathering problems, we develop a simple distributed

protocol based on the algorithm. The performance of this protocol is studied through simulations. Because the store-and-gather problems are by nature off-line problems, we do not develop a distributed protocol for this class of problems.

The rest of the paper is organized as follows. The data gathering problems are discussed in Section 2. We show that these problems reduce to network flow problem with constraint on the vertices. In Section 3, we develop a mathematical formulation of the constrained network flow problem and show that it reduces to a standard network flow problem. In Section 4, we derive a relaxed form for the network flow problem. A distributed and adaptive algorithm is then developed for this relaxed problem. A simple protocol based on this algorithm is presented in Section 4.3. Experimental results are presented in Section 5. Section 6 concludes this paper.

2 Data Gathering with Energy Constraint

2.1 System Model

Suppose a network of sensors is deployed over a region. The location of the sensors are fixed and known *a priori*. The system is represented by a graph $G(V, E)$, where V is the set of sensor nodes. $(u, v) \in E$ if $u \in V, v \in V$ and u is within the communication range of v . The set of successors of u is denoted as $\sigma_u = \{v \in V | (u, v) \in E\}$. Similarly, the set of predecessors of u is denoted as $\psi_u = \{v \in V | (v, u) \in E\}$. The event is sensed by a subset of sensors $V_c \subset V$. r is the base station to which the sensed data are transmitted. Sensors $V - V_c - \{r\}$ in the network does not sense the event but can relay the data sensed by V_c .

Among the three categories (sensing, communication, and data processing) of power consumption, a sensor node typically spends most of its energy in data communication. This includes both data

transmission and reception. Our energy model for the sensors is based on the first order radio model described in [9]. The energy consumed by sensor u to transmit a k -bit data packet to sensor v is $T_{uv} = \varepsilon_{elec} \times k + \varepsilon_{amp} \times d_{uv}^2 \times k$, where ε_{elec} is the energy required for transceiver circuitry to process one bit of data, ε_{amp} is the energy required per bit of data for transmitter amplifier, and d_{uv} is the distance between u and v . Transmitter amplifier is not needed by u to receive data and the energy consumed by u to receive a k -bit data packet is $R_u = \varepsilon_{elec} \times k$. Typically, $\varepsilon_{elec} = 50nJ/bit$ and $\varepsilon_{amp} = 0.1nJ/bit/m^2$. This effectively translates to $\varepsilon_{amp} \times d_{uv}^2 \ll \varepsilon_{elec}$, especially when short transmission ranges ($\simeq 1m$) are considered. For the discussion in the rest of this paper, we adopt the approximation that $T_{uv} = R_u$ for $(u, v) \in E$. We further assume that no data aggregation is performed during the transmission of the data.

Communication link (u, v) has transmission bandwidth c_{uv} . We do not require the communication links to be identical. Two communication links may have different transmission latencies and/or bandwidth. Symmetry is not required either. It may be the case that $c_{uv} \neq c_{vu}$. If $(u, v) \notin E$, then we define $c_{uv} = 0$.

An energy budget B_u is imposed on each sensor node u . We assume that there is no energy constraint on base station r . To simplify our discussions, we ignore the energy consumption of the sensors when sensing the environment. However, the rate at which sensor $u \in V_c$ can collect data from the environment is limited by the maximum sensing capability g_u . We consider both store-and-gather problems and continuous sensing and gathering problems. For the store-and-gather problems, B_u represents the total number of data packets that u can send and receive. For the continuous sensing and gathering problems, B_u represents the total number of data packets that u can send and receive in one unit of time.

2.2 Store-and-Gather Problems

In store-and-gather problems, the information from the environment is sensed (possibly over a long time period) and stored locally at the sensors. The data is then transferred to the base station during the data gathering stage. This represents those data-oriented applications (e.g. counting the occurrences of endangered birds in a particular region) where the environment changes slowly. There is typically no deadline (or the deadline is loose enough to be ignored) on the duration of data gathering for such problems, and we are not interested in the speed at which the data is gathered. But due to the energy constraint, not all the stored data can be gathered by the base station, and we want to maximize the amount of data gathered.

For each $u \in V_c$, we assume that u has stored d_u data packet before the data gathering starts. Let $f(u, v)$ represent the number of data packets sent from u to v .

For the simplified scenario where V_c contains a single node s , we have the following problem formulation:

Single Source Maximum Data Volume (SMaxDV) Problem:

Given: A graph $G(V, E)$. Source $s \in V$ and sink $r \in V$. Each node $u \in V - \{r\}$ has energy budget B_u .

Find: A real valued function $f : E \rightarrow R$

Maximize: $\sum_{v \in \sigma_s} f(s, v)$

Subject to:

$$f(u, v) \geq 0 \quad \text{for } \forall (u, v) \in E \quad (1)$$

$$\sum_{v \in \sigma_u} f(u, v) + \sum_{v \in \psi_u} f(v, u) \leq B_u \quad \text{for } u \in V - \{r\} \quad (2)$$

$$\sum_{v \in \sigma_u} f(u, v) = \sum_{v \in \psi_u} f(v, u) \quad \text{for } u \in V - \{s, r\} \quad (3)$$

B_u is the energy budget of u . Since we have normalized both T_{uv} and R_u to 1, the total number of data packets that can be sent and received by u is bounded from above by B_u . Condition 2 above represents the energy constraint of the sensors. Sensors $V - \{s, r\}$ do not generate sensed data, nor should they possess any data packets upon the completion of the data gathering. This is reflected in Condition 3 above. We do not model d_s , the number of data packets stored at s before the data gathering starts. This is because d_s is an obvious upper bound for the SMaxDV problem, and can be handled trivially.

$|V_c| > 1$ represents the general scenario where the event is sensed by multiple sensors. This multi-source data gathering problem is formulated as follows:

Multiple Source Maximum Data Volume (MMaxDV) Problem:

Given: A graph $G(V, E)$. The set of source nodes $V_c \subset V$ and sink $r \in V$. Each node $u \in V - \{r\}$ has energy budget B_u . Each node $v \in V_c$ has d_v data packets that are locally stored before the data gathering starts.

Find: A real valued function $f : E \rightarrow R$

Maximize: $\sum_{v \in \sigma_s} f(s, v)$

Subject to:

$$f(u, v) \geq 0 \quad \text{for } \forall (u, v) \in E \quad (1)$$

$$\sum_{v \in \sigma_u} f(u, v) + \sum_{v \in \psi_u} f(v, u) \leq B_u \quad \text{for } u \in V - \{r\} \quad (2)$$

$$\sum_{v \in \sigma_u} f(u, v) = \sum_{v \in \psi_u} f(v, u) \quad \text{for } u \in V - V_c - \{r\} \quad (3)$$

$$\sum_{v \in \sigma_u} f(u, v) \leq \sum_{v \in \psi_u} f(v, u) + d_u \quad \text{for } u \in S_c \quad (4)$$

Similar to the SMaxDV problem, the net flow out of the intermediate nodes ($V - V_c - \{r\}$) is 0 in the MMaxDV problem, as is specified in Condition 3. For each source node $u \in V_c$, the net flow out of u cannot exceed the number of data packets previously stored at u . This is specified in Condition 4.

2.3 Continuous Sensing and Gathering Problems

The continuous sensing and gathering problems model those time critical applications that need to gather as much information as possible from the environment while the nodes are sensing. Examples of such applications include battle field surveillance, target tracking, etc. We want to maximize the total number of data packets that can be gathered by the base station r in one unit of time. We assume that the communications are scheduled by time/frequency division multiplexing or channel assignment techniques. We consider the scenario in which B_u is the maximum power consumption rate allowed by u . Let $f(u, v)$ denote the number of data packets sent from u to v in one unit of time.

Similar to the store-and-gather problem, we have the following mathematical formulation when V_c contains a single node s .

Single Source Maximum Data Throughput (SMaxDT) Problem:

Given: A graph $G(V, E)$. Source $s \in V$ and sink $r \in V$. Each node $u \in V - \{r\}$ has energy budget B_u . Each edge $(u, v) \in E$ has capacity c_{uv} .

Find: A real valued function $f : E \rightarrow R$

Maximize: $\sum_{v \in \sigma_s} f(s, v)$

Subject to:

$$0 \leq f(u, v) \leq c_{uv} \quad \text{for } \forall (u, v) \in E \quad (1)$$

$$\sum_{v \in \sigma_u} f(u, v) + \sum_{v \in \psi_u} f(v, u) \leq B_u \quad \text{for } u \in V - \{r\} \quad (2)$$

$$\sum_{v \in \sigma_u} f(u, v) = \sum_{v \in \psi_u} f(v, u) \quad \text{for } u \in V - \{s, r\} \quad (3)$$

The major difference between the SMaxDV and the SMaxDT problem is the consideration of link capacities. In the SMaxDV problem, since there is no deadline for the data gathering, the primary factor that affects the maximum number of gathered data is the energy budgets of the sensors. But for the

SMaxDT problem, the number of data packets that can be transferred over a link in one unit of time is not only affected by the energy budget, but also bounded from above by the capacity of that link, as is specified in Condition 1 above. For the SMaxDT problem, we do not model the impact of g_u because g_u is an obvious upper bound of the throughput and can be handled trivially.

Similarly, we can formulate the multiple source maximum data throughput problem as follows.

Multiple Source Maximum Data Throughput (MMaxDT) Problem:

Given: A graph $G(V, E)$. The set of source nodes $V_c \subset V$ and sink $r \in V$. Each node $u \in V - \{r\}$ has energy budget B_u . Each edge $(u, v) \in E$ has capacity c_{uv} .

Find: A real valued function $f : E \rightarrow R$

Maximize: $\sum_{v \in \sigma_s} f(s, v)$

Subject to:

$$0 \leq f(u, v) \leq c_{uv} \quad \text{for } \forall (u, v) \in E \quad (1)$$

$$\sum_{v \in \sigma_u} f(u, v) + \sum_{v \in \psi_u} f(v, u) \leq B_u \quad \text{for } u \in V - \{r\} \quad (2)$$

$$\sum_{v \in \sigma_u} f(u, v) = \sum_{v \in \psi_u} f(v, u) \quad \text{for } u \in V - V_c - \{r\} \quad (3)$$

$$\sum_{v \in \sigma_u} f(u, v) \leq \sum_{v \in \psi_u} f(v, u) + g_u \quad \text{for } u \in V_c \quad (4)$$

Condition 4 in the above problem formulation takes into account the sensing capabilities of the sensors.

3 Flow Maximization with Constraint on Vertices

3.1 Problem Reductions

In this section, we present the formulation of the constrained flow maximization problem where the vertices have limited capacities (*CFM* problem). The CFM problem is an abstraction of the four problems discussed in Section 2.

In the CFM problem, we are given a directed graph $G(V, E)$ with vertex set V and edge set E . Vertex u has capacity constraint $w_u > 0$. Edge (u, v) starts from vertex u , ends at vertex v , and has capacity constraint $c_{uv} > 0$. If $(u, v) \notin E$, we define $c_{uv} = 0$. We distinguish two vertices in G , source s , and sink r . A flow in G is a real valued function $f : E \rightarrow R$ that satisfies the following constraints:

1. $0 \leq f(u, v) \leq c_{uv}$ for $\forall (u, v) \in E$. This is the capacity constraint on edge (u, v) .
2. $\sum_{v \in \sigma_u} f(u, v) = \sum_{v \in \psi_u} f(v, u)$ for $\forall u \in V - \{s, r\}$. This represents the flow conservation. The net amount of flow that goes through any of the vertices, except s and t , is zero.
3. $\sum_{v \in \sigma_u} f(u, v) + \sum_{v \in \psi_u} f(v, u) \leq w_u$ for $\forall u \in V$. This is the capacity constraint of vertex u . The total amount of flow going through u cannot exceed w_u . This condition differentiates the CFM problem from the standard network flow problem.

The *value* of a flow f , denoted as $|f|$, is defined as $|f| = \sum_{v \in \sigma_s} f(s, v)$, which is the net flow that leaves s . In the CFM problem, we are given a graph with vertex and edge constraint, a source s , and a sink r , and we wish to find a flow with the maximum value.

It is straight forward to show that the SMaxDV and the SMaxDT problems reduce to the CFM problem. By adding a hypothetical super source node, the MMaxDV and the MMaxDT problems can also be reduced to SMaxDV and SMaxDT, respectively.

It can be shown that the CFM problem reduces to a standard network flow problem. Due to the existence of condition 1, condition 3 is equivalent to $\sum_{v \in \sigma_u} f(u, v) \leq w_u/2$ for $u \in V - \{s, r\}$. This means that the total amount of flow out of vertex u cannot exceed $w_u/2$. Suppose we split u ($u \in V - \{s, r\}$) into two nodes u_1 and u_2 , re-direct all incoming links to u to arrive at u_1 and all the outgoing links from u to leave from u_2 , and add a link from u_1 to u_2 with capacity $w_u/2$, then the vertex constraint w_u is fully represented by the capacity of link (u_1, u_2) . Actually, such a split transforms all the vertex constraints to the corresponding link capacities, and effectively reduces the CFM problem to a standard network flow problem. The CFM problem has been studied in [12] where a similar reduction can be found.

The standard network flow problem is stated below:

Given: graph $G(V, E)$, source node $s \in V$, and sink node $r \in V$. Link (u, v) has capacity c_{uv} .

Maximize: $\sum_{v \in \sigma_s} f(s, v)$

Subject to:

$$0 \leq f(u, v) \leq c_{uv} \quad \text{for } \forall (u, v) \in E \quad (1)$$

$$\sum_{v \in \sigma_u} f(u, v) = \sum_{v \in \psi_u} f(v, u) \quad \text{for } u \in V - \{s, r\} \quad (2)$$

3.2 Relationship to Sensor Network Scenarios

The vertex capacity w_u in the CFM problem models the energy budget B_u of the sensor nodes. B_u does not have to be the total remaining energy of u . For example, when the remaining battery power of a sensor is lower than a particular level, the sensor may limit its contribution to the data gathering operation by setting a small value for B_u (so that this sensor still has enough energy for future operations). For another example, if a sensor is deployed in a critical location so that it is utilized as a gateway to relay

data packets to a group of sensors, then it may limit its energy budget for a particular data gathering operation, thereby conserving energy for future operations. These considerations can be captured by vertex capacity w_u in the CFM problem.

The edge capacity in the CFM problem models the communication rate (meaningful for continuous sensing and gathering problems) between adjacent sensor nodes. The edge capacity captures the available communication bandwidth between two nodes, which may be less than the the maximum available rate. For example, a node may reduce its radio transmission power to save energy, resulting in a less than maximum communication rate. This capacity can also vary over time based on environmental conditions. Our decentralized protocol results in an on-line algorithm for this scenario.

Because energy efficiency is a key consideration, various techniques have been proposed to explore the trade-offs between processing/communication speed and energy consumption. This results in the continuous variation of the performance of the nodes. For example, the processing capabilities may change as a result of dynamic voltage scaling [13]. The data communication rate may change as a result of modulation scaling [17]. As proposed by various studies on energy efficiency, it is necessary for sensors to maintain a power management scheme, which continuously monitors and adjusts the energy consumption and hence changes the computation and communication performance of the sensors. In data gathering problems, these energy related adjustments translate to changes of parameters (node/link capacities) in the problem formulations. Determining the exact reasons and mechanisms behind such changes is beyond the scope of this paper. Instead, we focus on the development of data gathering algorithms that can adapt to such changes.

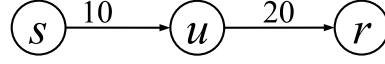


Figure 1. An example of the relaxed network flow problem where $c_{su} = 10$ and $c_{ur} = 20$.

4 Distributed and Adaptive Algorithm To Maximize Flow

In this section, we first show that the maximum flow remains the same even if we relax the flow conservation constraint. Then we develop a distributed and adaptive algorithm for the relaxed problem.

4.1 Relaxed Flow Maximization Problem

Consider the simple example in Figure 1 where s is the source, r is the sink, and u are the intermediate nodes. Obviously, the flow is maximized when $f(s, u) = f(u, r) = 10$. Suppose s , u , and r form an actual system and s has sent 10 data packets to u . Then u can send no more than 10 data packets to r even if u is allowed to transfer more to r . This means the actual system still works as if $f(u, r) = 10$ even if we set $f(u, r) \geq 10$.

This leads to the following relaxed network flow problem:

Given: graph $G(V, E)$, source node $s \in V$, and sink node $r \in V$. Link (u, v) has capacity c_{uv} .

Maximize: $|f| \stackrel{\text{def}}{=} \sum_{v \in \sigma_s} f(s, v)$

Subject to:

$$0 \leq f(u, v) \leq c_{uv} \quad \text{for } \forall (u, v) \in E \quad (1)$$

$$\sum_{v \in \sigma_u} f(u, v) \geq \sum_{v \in \psi_u} f(v, u) \quad \text{for } u \in V - \{s, r\} \quad (2)$$

Condition 2 differentiates the relaxed and the standard network flow problem. In the relaxed problem, the total flow out of a node can be equal to or larger than the total flow into the node. A feasible function f (which satisfies the two constraints above) to the relaxed flow problem is called a *relaxed flow* in graph

G . $|f|$ denotes the net amount of flow out of source s and is called the *value* of the relaxed flow. The following theorem shows the relation between the relaxed and the standard network flow problem.

Theorem 1. *Given graph $G(V, E)$, source s and sink r . If f^* is an optimal solution to the relaxed network flow problem, then there exists an optimal solution f' to the standard network flow problem such that $f'(u, v) \leq f^*(u, v)$ for $\forall (u, v) \in E$. Additionally, $|f^*| = |f'|$.*

Proof of the theorem is not difficult and omitted here due to space limitations. If we interpret $f^*(u, v)$ as the number of data units that we ask u to transfer and $f'(u, v)$ as the number of data units that u actually transfers, then this theorem essentially indicates that the solution to a relaxed flow problem can have an actual implementation that satisfies flow conservation.

4.2 The Algorithm

In this section, we develop a decentralized and adaptive algorithm for the relaxed network flow maximization problem. This algorithm is a modified version of the Push-Relabel algorithm [3] and is denoted as the *Relaxed Incremental Push-Relabel* (RIPR) algorithm.

The Push-Relabel algorithm is a well known algorithm for network flow maximization. It has a decentralized implementation where every node only needs to exchange messages with its immediate neighbors and makes decisions locally. But in order to be adaptive to the changes in the system, this algorithm has to be re-initialized and re-run from scratch each time when some parameters (weight of the nodes and edges in the graph) of the flow maximization problem change. Each time before starting to search for the new optimal solution, the algorithm needs to make sure that every node has finished its local initialization, which requires a global synchronization and compromises the property of decentralization.

In contrast to the Push-Relabel algorithm, our algorithm introduces the adaptation operation, which is performed upon the current values of $f(u, v)$ and $h(u)$ for $\forall u, v \in V$. In other words, our algorithm performs *incremental* optimization as the parameters of the system change. Our algorithm does not need global synchronizations. Another difference is that our algorithm applies to the relaxed network flow problem, rather than the standard one.

For the discussion below, let us first briefly re-state some notations for the network flow maximization problem. For notational convenience, if edge $(u, v) \notin E$, we define $c_{uv} = 0$; if the actual data transfer is from u to v , we define $f(v, u) = -f(u, v)$. With these two definitions, if neither (u, v) nor (v, u) belongs to E , then $c_{uv} = c_{vu} = 0$, which implies that $f(u, v) = f(v, u) = 0$. Of course, $f(u, u) = -f(u, u)$ implies that $f(u, u) = 0$, which essentially says that a node cannot send flow to itself. In this way, we can define $f(u, v)$ over $V \times V$, rather than being restricted to E . $f(u, v) = -f(v, u)$ also allows us to compute the total amount of flow into u as $\sum_{v \in \psi_u \cup \sigma_u} f(v, u)$, which equals $\sum_{v \in V} f(v, u)$ since $f(v, u) = f(u, v) = 0$ if $(u, v) \notin E$ and $(v, u) \notin E$. With the definition of $f(u, v)$ thus extended, it is easy to show that the relaxed network flow problem is equivalent to the following formulation:

Given: graph $G(V, E)$, source node $s \in V$, and sink node $r \in V$. Link (u, v) has capacity c_{uv} .
 $c_{uv} = 0$ if $(u, v) \notin E$.

Maximize: $|f| \stackrel{\text{def}}{=} \sum_{v \in V} f(s, v)$

Subject to:

$$f(u, v) = -f(v, u) \quad \text{for } u, v \in V \quad (1)$$

$$f(u, v) \leq c_{uv} \quad \text{for } u, v \in V \quad (2)$$

$$\sum_{v \in V} f(v, u) \leq 0 \quad \text{for } u \in V - \{s, r\} \quad (3)$$

Given a direct graph $G(V, E)$, function f is called a flow if it satisfies the three conditions in the

above problem; function f is called a pre-flow if it satisfies conditions 1 and 2. Given $G(V, E)$ and f , the *residual capacity* $c_f(u, v)$ is given by $c_{uv} - f(u, v)$, and the *residual network* of G induced by f is $G_f(V, E_f)$, where $E_f = \{(u, v) | u \in V, v \in V, c_f(u, v) > 0\}$. For each node $u \in V$, $e(u)$ is defined as $e(u) = \sum_{v \in V} f(v, u)$, which is the total amount of flow into u .

The algorithm is as follows:

1. *Initialization*: $h(u)$, and $f(u, v)$ are initialized as follows:

$$\begin{aligned}
 h(u) &= 0 && \text{for } \forall u \in V \\
 f(u, v) &= 0 && \text{for } \forall u, v \in E \\
 h(s) &= |V| \\
 f(s, v) &= l_{sv} && \text{for } \forall v \in V \\
 f(v, s) &= -l_{vs} && \text{for } \forall v \in V \\
 e(u) &= \sum_{v \in V} f(v, u) && \text{for } \forall u \in V
 \end{aligned}$$

2. *Search for maximum flow*:

Each node $u \in V - \{s, r\}$ conducts one of the following three operations as long as $e(u) \neq 0$:

(a) *Push*(u, v): applies when $e(u) > 0$ and $\exists (u, v) \in E_f$ s.t. $h(u) > h(v)$,

$$\begin{aligned}
 d &= \min(e(u), c_f(u, v)) \\
 f(u, v) &= f(u, v) + d \\
 f(v, u) &= -f(u, v) \\
 e(u) &= e(u) - d \\
 e(v) &= e(v) + d
 \end{aligned}$$

(b) *Relabel*(u): applies when $e(u) > 0$ and $h(u) \leq h(v)$ for $\forall (u, v) \in E_f$,

$$h(u) = \min_{(u,v) \in E_f} h(v) + 1$$

3. *Adaptation to changes in the system*: For the flow maximization problem, the only possible change that can occur in the system is the increase or decrease of the capacity of some edges. Suppose the value of l_{uv} changes to l'_{uv} , the following four scenarios are considered when performing the *Adaptation* (u, v) operation:

(a) if $l'_{uv} > l_{uv}$ and $f(u, v) < l_{uv}$, do nothing.

(b) if $l'_{uv} > l_{uv}$ and $f(u, v) = l_{uv}$, then

$$\begin{aligned} h(s) &= h(s) + 2|V| \\ f(s, v) &= l_{sv} && \text{for } \forall v \in \sigma_s \\ f(v, s) &= -l_{sv} && \text{for } \forall v \in \sigma_s \\ e(v) &= \sum_{v \in V} f(v, v) && \text{for } \forall v \in V \end{aligned}$$

(c) if $l'_{uv} < l_{uv}$ and $f(u, v) \leq l'_{uv}$, do nothing.

(d) if $l'_{uv} < l_{uv}$ and $f(u, v) > l'_{uv}$, then

$$\begin{aligned} h(s) &= h(s) + 2|V| \\ f(s, v) &= l_{sv} && \text{for } \forall v \in \sigma_s \\ f(v, s) &= -l_{sv} && \text{for } \forall v \in \sigma_s \\ e(v) &= \sum_{v \in V} f(v, v) && \text{for } \forall v \in V \\ f(u, v) &= l'_{uv} \\ f(v, u) &= -l'_{uv} \\ e(u) &= e(u) + (l_{uv} - l'_{uv}) \end{aligned}$$

$$e(v) = e(v) - (l_{uv} - l'_{uv})$$

The above algorithm defines an integer valued auxiliary function $h(u)$ for $u \in V$, which will be discussed below. The ‘adaptation’ is activated when some link capacity changes in the relaxed flow problem. Because link capacities in the relaxed flow problem map to either vertex or link capacities in the corresponding CFM problem, the adaptation operation actually reacts to capacity changes in both vertex and link capacities. The ‘Push’ and ‘Relabel’ operations are called the *basic operations*. Every node in the graph determines its own behavior based on the knowledge about itself and its neighbors (as can be seen, the Push and Relabel operations are triggered by variables whose value are locally known by the nodes). No central coordinator or global information about the system is needed. More importantly, unlike the Push-Relabel algorithm, no global synchronization is needed when the RIPR algorithm adapts to the changes in the system.

An intuitive explanation of the RIPR is as follows. $h(u)$ represents, intuitively, the shortest distance from u to t when $h(u) \leq h(s)$. When $h(u) > h(s)$, $h(u) - h(s)$ represents the shortest distance from u to s . Hence the RIPR algorithm attempts to push more flow from s to t along the shortest path; excessive flow of intermediate nodes are pushed back to s along the shortest path. Similar to the Edmonds-Karp algorithm[3], such a choice of paths can lead to an optimal solution.

Lemma 1. *During the execution of the RIPR Algorithm, for $\forall u \in V$, $h(u)$ never decreases.*

Proof: $h(s)$ is only changed by the Adaptation operation, during which $h(s)$ is increased by $2|V|$. $h(r)$ never changes. When $u \neq s$ and $u \neq r$, $h(u)$ is only changed by *Relabel*(u). *Relabel*(u) is applied when $e(u) > 0$ and $h(u) \leq h(v)$ for $\forall v \in \{v | c_f(u, v) > 0\}$. And $h(u_i) = \min_{(u,v) \in E_f} h(v) + 1$ after *Relabel*(u). Hence $h(u)$ is increased at least by 1. □

Lemma 2. *During the execution of the RIPR Algorithm, for $\forall u \in V$ s.t. $e(u) > 0$, there exists a simple path in E_f from u to a node v s.t. $e(v) < 0$.*

Proof: Suppose $e(u) > 0$. Let \hat{V} denotes the set of nodes that can be reached by u through a simple path in E_f . Note that $u \in \hat{V}$. For sake of contradiction, suppose $e(v) \geq 0$ for $\forall v \in \hat{V}$. Let $\bar{V} = V - \hat{V}$.

We claim that if $x \in \bar{V}$ and $y \in \hat{V}$, then $f(x, y) \leq 0$. Otherwise if $f(x, y) > 0$, then $c_f(y, x) = c_{yx} - f(y, x) = c_{yx} + f(x, y) > 0$, which means x can be reached from y in E_f , hence there exists a path from u to x in E_f . But this contradicts the choice that $x \in \bar{V}$.

It is fairly easy to show that $\sum_{v \in \hat{V}} e(v) = \sum_{x \in \bar{V}, y \in \hat{V}} f(x, y)$. Hence $\sum_{v \in \hat{V}} e(v) \leq 0$. But this contradicts the assumption that $e(v) \geq 0$ for $\forall v \in \hat{V}$ and $e(u) > 0$. \square

Lemma 3. *During the execution of the RIPR, for $\forall u \in V$, if $e(u) > 0$, then either $Relabel(u)$, or $Push(u, v)$ (where $v \in V$) can be applied.*

Proof: When $e(u) > 0$, if $\exists v$ s.t. $c_f(u, v) > 0$ and $h(u) > h(v)$, then $Push(u, v)$ can be applied; otherwise, $h(u) \leq h(v)$ for $\forall v \in \{v | c_f(u, v) > 0\}$, which means $Relabel(u)$ can be applied. \square

Lemma 4. *During the execution of the RIPR algorithm,*

$$h(u) \leq \max(h(v) + 1, h(s) - |V| - 1) \quad \text{for } \forall (u, v) \in E_f,$$

$$h(u) \leq h(s) + |V| - 1 \quad \text{for } \forall u \in V$$

Proof: We prove by induction on the number of adaptation operations.

- Base case:

Before any changes occur in the system, the adaptation operation will not be applied. At this stage, the Incremental Push-Relabel algorithm performs the exact operations as the Push-Relabel algorithm, hence we have

$$\begin{aligned}
h(u) &\leq \max(h(v) + 1, h(s) - |V| - 1) && \text{for } \forall (u, v) \in E_f, \\
h(u) &\leq h(s) + |V| - 1 && \text{for } \forall u \in V
\end{aligned}$$

before any adaptation operation is applied.

- Induction step:

Suppose after the adaptation has been applied $n - 1$ times and we still have

$$\begin{aligned}
h(u) &\leq \max(h(v) + 1, h(s) - |V| - 1) && \text{for } \forall (u, v) \in E_f, \\
h(u) &\leq h(s) + |V| - 1 && \text{for } \forall u \in V
\end{aligned}$$

and then the n^{th} adaptation, $Adaptation(u^*, v^*)$, is applied.

1. We first show that $h(u) \leq \max(h(v)+1, h(s)-|V|-1)$ for $\forall (u, v) \in E_f$ after $Adaptation(u^*, v^*)$.

Considering $Adaptation(u^*, v^*)$, if either scenario (a) or (c) occurs, no residual edge is added or removed, no node $w \in V$ has its $h(w)$ changed, either. If scenario (d) occurs, the change in the system removes (u^*, v^*) from E_f and hence the corresponding constraint on $h(u^*)$ and $h(v^*)$. If scenario (b) occurs, (u^*, v^*) is added to the E_f . By induction assumption, $h(u^*) \leq h(s) + |V| - 1$ before the adaptation operation. Because $h(u^*)$ does not change and $h(s)$ increases by $2|V|$ after the operation, $h(u^*) \leq h(s) - |V| - 1$ after the adaption operation. In summary, after the adaptation operation, $h(u) \leq \max(h(v) + 1, h(s) - |V| - 1)$ for $\forall (u, v) \in E_f$.

$Adaptation(u^*, v^*)$ changes the values of some $e(w)$, allowing new Push and Relabel operations to be applied. Yet these operations preserve the property that

$h(u) \leq \max(h(v) + 1, h(s) - |V| - 1)$ for $\forall (u, v) \in E_f$. This is shown by induction on the number of Push and Relabel operations.

(a) Suppose $Push(u, v)$ is applied.

This may add edge (v, u) into E_f or remove the edge (u, v) from E_f . In the former case, for edge $(v, u) \in E_f$, we have $h(v) < h(u)$ because otherwise the push will not be applied. In the latter case, the removal of (u, v) from E_f removes the corresponding constraint on $h(u)$ and $h(v)$. In both cases, we still have $h(u) \leq \max(h(v) + 1, h(s) - |V| - 1)$ for any $(u, v) \in E_f$.

(b) Suppose $Relabel(u)$ is applied.

For a residual edge (u, v) that leaves u , we have $h(u) = \min_{w \in \{w | (u, w) \in E_f\}} h(w) + 1$ after the Relabel operation, which means $h(u) \leq h(v) + 1$. For a residual edge (w, u) that enters u , $h(w) \leq \max(h(u) + 1, h(s) - |V| - 1)$ before the relabel operation. According to Lemma 1, $h(w) \leq \max(h(u) + 1, h(s) - |V| - 1)$ after the relabel operation. Therefore, after a relabel operation, we have $h(u) \leq \max(h(v) + 1, h(s) - |V| - 1)$ for any $(u, v) \in E_f$.

2. Now we need to show that $h(u) \leq h(s) + |V| - 1$ for $\forall u \in V$.

Let \hat{V} denote the set of $u \in V$ s.t. there exists a simple path from u to s in E_f . $\bar{V} = V - \hat{V}$.

For any node $u \in \hat{V}$, suppose the simple path to s in E_f is $\{u, u_1, \dots, u_k\}$, where $u_k = s$ and $k \leq |V| - 1$. We have

$$h(u) \leq \max(h(u_1) + 1, h(s) - |V| - 1)$$

$$h(u_1) \leq \max(h(u_2) + 1, h(s) - |V| - 1)$$

...

$$h(u_{k-1}) \leq \max(h(u_k) + 1, h(s) - |V| - 1)$$

Combining these inequalities, we have

$$\begin{aligned} h(u) &\leq \max(h(u_k) + k, h(s) - |V| - 1 + k - 1) \\ &= \max(h(s) + k, h(s) - |V| + k - 2) \\ &\leq \max(h(s) + |V| - 1, h(s) - 3) \\ &= h(s) + |V| - 1 \end{aligned}$$

For any node $u \in \bar{V}$, according to Lemma 2, there exist a simple path in E_f from u to a node z s.t. $e(z) < 0$ and $z \neq s$. Suppose the simple path is $\{u, u_1, \dots, u_k\}$, where $u_k = z$ and $k \leq |V| - 1$. We have

$$\begin{aligned} h(u) &\leq \max(h(u_1) + 1, h(s) - |V| - 1) \\ h(u_1) &\leq \max(h(u_2) + 1, h(s) - |V| - 1) \\ &\dots \\ h(u_{k-1}) &\leq \max(h(u_k) + 1, h(s) - |V| - 1) \end{aligned}$$

Note that $e(u) \geq 0$ immediately after the initialization for $\forall u \in V$. The only operation that can bring $e(z)$ below 0 is the adaptation operation (when scenario (d) occurs). Suppose $e(z)$ becomes negative as the result of m^{th} adaptation operation ($m \leq n$). Since the Relabel operation (which is the only operation that can increase the value of $h(z)$) is applied only if $e(z) > 0$, then $e(z) < 0$ means that $h(z)$ has not been increased after the m^{th} , and hence the n^{th} adaptation operation. Therefore, $h(z) \leq h(s) + |V| - 1$ before the n^{th} adaptation means that $h(z) \leq h(s) - |V| - 1$ thereafter, since $h(s)$ is increased by $2|V|$.

Combining these inequalities, we have

$$\begin{aligned}
h(u) &\leq \max(h(V_{i_k}) + k, h(s) - |V| - 1 + k - 1) \\
&= \max(h(V_s) + k, h(s) - |V| - 2 + k) \\
&\leq \max(h(V_0) - |V| - 1 + |V| - 1, h(s) - |V| - 2 + |V| - 1) \\
&\leq h(V_0) + |V| - 1
\end{aligned}$$

□

Corollary 1. *During the execution of the RIPR algorithm, for any node $u \in V$, if $e(u) < 0$, then $h(u) \leq h(s) - |V| - 1$.*

The proof of Corollary 1 is included in the proof of Lemma 4.

Lemma 5. *During the execution of the RIPR algorithm, for $\forall u \in V - \{s, r\}$, if there exists a simple path from s to u in E_f , then $e(u) \geq 0$.*

Proof: Suppose for the sake of contradiction that there exists a node $u \in V - \{s, r\}$ such that $e(u) < 0$ and there exists a simple path $\{s, u_1, \dots, u_k, u\}$ in E_f . Without loss of generality, this is a simple path and $k \leq |V| - 2$.

According to Lemma 4,

$$\begin{aligned}
h(u_1) &\leq \max(h(u_2) + 1, h(s) - |V| - 1) \\
&\dots \\
h(u_{k-1}) &\leq \max(h(u_k) + 1, h(s) - |V| - 1) \\
h(u_k) &\leq \max(h(u) + 1, h(s) - |V| - 1)
\end{aligned}$$

According to Corollary 1, $h(u) \leq h(s) - |V| - 1$ since $e(u) < 0$. Combining these inequalities, we can see that

$$\begin{aligned}
h(u_1) &\leq \max(h(u) + k, h(s) - |V| - 2 + k) \\
&\leq \max(h(s) - |V| - 1 + k, h(s) - |V| - 2 + k) \\
&\leq h(s) - |V| - 1 + |V| - 2 \\
&< h(s)
\end{aligned}$$

On the other hand, consider the first hop (s, u_1) along this path. $(s, u_1) \in E_f$ implies that $f(s, u_1) < c_{su_1}$. Recall that the value of $f(s, u_1)$ is set to c_{su_1} after the initialization and each adaptation operation. The only operation that can reduce the value of $f(s, u_1)$ is a push from u_1 to s . However, $Push(u_1, s)$ is applied only when $h(u_1) > h(s)$. This contradicts the claim $h(u_1) < h(s)$ that we just derived.

□

Similar to the standard flow problem, for the relaxed flow problem, a *cut* is defined as a binary partition (S, R) of V such that $S \cup R = V$, $S \cap R = \emptyset$, $s \in S$ and $r \in R$. The capacity c_{SR} of a cut (S, R) is defined as $c_{SR} = \sum_{u \in S, v \in R} c_{uv}$. The next lemma shows that the value of a relaxed flow cannot exceed the capacity of any cuts.

Lemma 6. *Given graph $G(V, E)$ with source s and sink r , a relaxed flow f , and an arbitrary cut (S, R) of G , $\sum_{v \in V} f(s, v) \leq c_{SR}$.*

Proof: we have $e(u) \leq 0$ for $u \in V - \{s, r\}$. Therefore

$$\begin{aligned}
\sum_{u \in S - \{s\}} e(u) &= \sum_{v \in V, u \in S - \{s\}} f(v, u) \leq 0 \\
&\Rightarrow \sum_{v \in S, u \in S - \{s\}} f(v, u) + \sum_{v \in R, u \in S - \{s\}} f(v, u) \leq 0 \\
&\Rightarrow \sum_{u \in S - \{s\}} f(s, u) + \sum_{v \in S - \{s\}, u \in S - \{s\}} f(v, u) + \sum_{v \in R, u \in S - \{s\}} f(v, u) \leq 0 \\
&\Rightarrow \sum_{u \in S - \{s\}} f(s, u) + \sum_{v \in S - \{s\}, u \in S - \{s\}} f(v, u) \leq \sum_{v \in R, u \in S - \{s\}} f(u, v) \\
&\Rightarrow \sum_{u \in S - \{s\}} f(s, u) \leq \sum_{v \in R, u \in S - \{s\}} f(u, v)
\end{aligned}$$

$$\begin{aligned} &\Rightarrow \sum_{u \in S - \{s\}} f(s, u) + \sum_{u \in R} f(s, u) \leq \sum_{v \in R, u \in S - \{s\}} f(u, v) + \sum_{u \in R} f(s, u) \\ &\Rightarrow \sum_{u \in V - \{s\}} f(s, u) \leq \sum_{u \in S, v \in R} f(u, v) \end{aligned}$$

Since $f(u, v) \leq c_{uv}$ for $\forall u, v \in V$, additionally, because $f(s, s) = 0$, we have

$$\sum_{v \in V} f(s, v) = \sum_{u \in V - \{s\}} f(s, u) \leq \sum_{u \in S, v \in R} c_{uv} = c_{SR}$$

□

The next lemma shows that the RIPR algorithm finds the maximum relaxed flow if it terminates. After proving this lemma, we will show that the RIPR algorithm indeed terminates.

Lemma 7. *If the RIPR algorithm terminates, it finds the maximum relaxed flow.*

Proof:

According to Lemma 3, if the algorithm terminates, then $e(u) \leq 0$ for $\forall u \in V - \{s, r\}$. Hence f is a flow if the algorithm terminates.

Given such an f , we construct a cut of G as follows:

$$S = \{u \in V \mid \text{there exists a simple path from } s \text{ to } u \text{ in } E_f\}$$

$$R = V - S$$

According to Lemma 5, $e(u) \geq 0$ for $u \in S - \{s\}$. Note that $e(u) \leq 0$ for $\forall u \in V - \{s, r\}$ upon termination of the algorithm. Hence $e(u) = 0$ (i.e. $\sum_{v \in V} f(v, u) = 0$) for $\forall u \in S - \{s\}$. Then it is easy to show that $\sum_{u \in S} f(s, u) = \sum_{u \in S - \{s\}} f(s, u) = \sum_{u \in S - \{s\}, v \in R} f(u, v)$.

Therefore,

$$\sum_{v \in V} f(s, v) = \sum_{v \in S} f(s, v) + \sum_{v \in R} f(s, v)$$

$$\begin{aligned}
&= \sum_{u \in S - \{s\}, v \in R} f(u, v) + \sum_{v \in R} f(s, v) \\
&= \sum_{u \in S, v \in R} f(u, v)
\end{aligned}$$

We claim that $f(u, v) = c_{uv}$ for $\forall u \in S, v \in R$, because otherwise $f(u, v) \leq c_{uv}$ implies that edge $(u, v) \in E_f$, hence v can be reached by s in E_f . But this contradicts the definition of R .

Therefore,

$$\sum_{v \in V} f(s, v) = \sum_{u \in S, v \in R} c(u, v) = c_{SR}$$

According to Lemma 6, such a relaxed flow f is a maximum relaxed flow. \square

Now we show that the algorithm indeed terminates.

Lemma 8. *If the adaptation is applied n ($n \geq 0$) times, then the number of Relabel operations that can be performed is less than $(2n + 2)|V|^2$.*

Proof: If the adaptation is applied n times, then $h(s) = (2n + 1)|V|$. According to Lemma 4, $h(u) \leq h(s) + |V| - 1 = (2n + 2)|V| - 1$ for $\forall u \in V$. Each time $Relabel(u)$ is applied, $h(s)$ is increased at least by 1. Since $h(s) = 0$ initially, $Relabel(u)$ is applied at most $(2n + 2)|V| - 1$ times. There are $|V|$ nodes in the system, hence then the total number of Relabel operations that can be performed is at most $((2n + 2)|V| - 1)|V|$, which is less than $(2n + 2)|V|^2$. \square

If a $Push(u, v)$ operation removes (u, v) from E_f (i.e. $c_f(u, v) = 0$ after the operation), it is a *saturated push*, otherwise it is a *non-saturated push*.

Lemma 9. *If the adaptation is applied n ($n \geq 0$) times, then the number of saturated push operations that can be performed is less than $(n + 1)|V| \cdot |E|$.*

Proof: Consider edge $(u, v) \in E$. Suppose a saturated push $Push(u, v)$ is first applied. For a second saturated push to be applied over (u, v) , $Push(v, u)$ must be applied before the second saturated push.

Because $h(u) > h(v)$ for the first push (otherwise the first push will not be applied), then $h(v)$ must be increased at least by 2, otherwise $h(v) \leq h(u)$ then $Push(v, u)$ will not be applied. Similarly, $h(u)$ must be increased at least by 2 for the second saturated push to occur over edge (u, v) . So on so force. Because $h(u) \leq h(s) + |V| - 1 = (2n + 2)|V| - 1$, $h(u)$ and $h(v)$ cannot increase to infinity. It is easy to show that saturated push can occur at most $((2n + 2)|V| - 1)/2$ times for edge (u, v) . There are $|E|$ edges in the graph. The total number of saturated push operations is less than $((2n + 2)|V| - 1)/2 \cdot |E|$, which is less than $(n + 1)|V| \cdot |E|$. \square

Lemma 10. *If the adaptation is applied n ($n \geq 0$) times, then the number of non-saturated push operations that can be performed is less than $(n^2 + 2n + 1) \cdot (4|V|^3 + 2|V|^2|E|)$.*

Proof: Define a potential function $\Phi = \sum_{e(u)>0} h(u)$. $\Phi = 0$ initially. Obviously, $\Phi \geq 0$.

According to Lemma 4, $h(u) \leq h(s) + |V| - 1 = (2n + 2)|V| - 1$, hence a relabel operation increases Φ by at most $(2n + 2)|V| - 1$. According to Lemma 8, there can be at most $(2n + 2)|V|^2$ Relabel operations. The increase in Φ induced by all relabel operations is at most $((2n + 2)|V| - 1) \cdot (2n + 2)|V|^2$. A saturated push $Push(u, v)$ increases Φ by at most $(2n + 2)|V| - 1$ since $e(v)$ may become positive after the push, and $(2n + 2)|V| - 1$ is the highest value that $h(v)$ can be. According to Lemma 9, the increase in Φ induced by all saturated push is at most $((2n + 2)|V| - 1) \cdot ((n + 1)|V| \cdot |E|)$.

For a non-saturated push $Push(u, v)$, $e(u) > 0$ before the push and $e(u) = 0$ after the push, hence $h(u)$ is excluded from Φ after the push. If $e(v) > 0$ after the push, Φ is decreased at least by 1 because $h(u) - h(v) > 0$. If $e(v) \leq 0$ after the push, then Φ is decreased by $h(u) \geq 1$.

Therefore, the total increase in Φ is at most $((2n + 2)|V| - 1) \cdot (2n + 2)|V|^2 + ((2n + 2)|V| - 1) \cdot ((n + 1)|V| \cdot |E|) < (n^2 + 2n + 1) \cdot (4|V|^3 + 2|V|^2|E|)$, while each non-saturated push decreases Φ at least by 1. Therefore, the total number of non-saturated push operations that can be performed is at

most $(n^2 + 2n + 1) \cdot (4|V|^3 + 2|V|^2|E|)$. □

Theorem 2. *The RIPR algorithm finds the maximum flow for the relaxed flow problem with $O(n^2 \cdot |V|^2 \cdot |E|)$ basic operations, where n is the number of adaptation operations performed, $|V|$ is the number of nodes in the graph, and $|E|$ is the number of edges in the graph.*

Proof: Immediate from Lemma 7, 8, 9, and 10. □

4.3 A Simple Protocol for Data Gathering

In this section, we present a simple on-line protocol for SMaxDT problem based on the RIPR algorithm in Section 4.

In this protocol, each node maintains a data buffer. Initially, all the data buffers are empty. The source node s senses the environment and fills its buffer continuously. At any time instance, let β_u denote the amount of buffer used by node u . Each node $u \in V$ operates as follows:

1. Contact the adjacent node(s) and execute the RIPR algorithm.
2. While $\beta_u > 0$, send message ‘request to send’ to all successors v of u s.t. $f(u, v) > 0$. If ‘clear to send’ is received from v , then set $\beta_u \leftarrow \beta_u - 1$ and send a data packet to v at rate $f(u, v)$. (recall that $f(u, v)$ is the flow rate at which data should be sent from u to v according to the RIPR algorithm.)
3. Upon receiving ‘request to send’, u acknowledges ‘clear to send’ if $\beta_u \leq U$. Here U is a pre-set threshold that limits the maximum number of data packets a buffer can hold.

For node s , it stops sensing if $\beta_s > U$. The nodes execute the RIPR algorithm and find the rate $f(u, v)$ for sending the data. Meanwhile, the nodes transfer the data according to the values of $f(u, v)$, without

waiting for the RIPR algorithm to terminate. Two types of data are transferred in the system: the control messages that are used by the RIPR algorithm, and the sensed data themselves. The control messages are exchanged among the nodes to query and update the values of $f(u, v)$ and h_u when executing the RIPR algorithm. The control messages and the sensed data are transmitted over the same links and higher priority is given to the control messages in case of a conflict.

For the MMaxDT problem, the situation is a bit more complicated. Since the MMaxDT problem is reduced to the SMaxDT problem by adding a hypothetical super source node s' , the RIPR algorithm needs to maintain the flow out of s' as well as the value of function $h(s')$. Additionally, the values of $f(s', v)$ ($v \in V_c$) and $h(s')$ are needed by all nodes $u \in V_c$ during the execution of the algorithm. Because s' is not an actual sensor, sensors in V_c therefore need to maintain a consistent image of s' . This requires some regional coordination among sensors in V_c and may require some extra cost to actually implement such a consistent image.

The SMaxDV and MMaxDV problems are by nature off-line problems and we do not develop online protocols for the two problems.

5 Experimental Results

Simulations were conducted to illustrate the effectiveness of the RIPR algorithm and the data gathering protocol. For the sake of illustration, we present simulation results for the SMaxDT problem.

The systems were generated by randomly scattering the sensor nodes in a unit square. The base station was located at the lower-left corner of the square. The source node was randomly chosen from the sensor nodes. B_u 's are uniformly distributed between 0 and B_{max} . B_{max} was set to 100. We assume a signal decaying factor of r^{-2} . The flow capacity between sensor nodes u and v is determined by Shannon's

theorem as $l_{uv} = W \log(1 + \frac{P_{uv}r_{uv}^{-2}}{\eta})$ where W is the bandwidth of the link, r_{uv} is the distance between u and v , P_{uv} is the transmission power on link (u, v) , and η is the noise in the communication channel. In all the simulations, W was set to 1KHz , P_{uv} was set to $10^{-3}mW$, and η was set to $10^{-6}mW$. U was set to 2. Each data packet was assumed to contain 32 bytes. Each control message was assumed to be transferred in $1ms$.

The RIPR algorithm described in Section 4 adapts to every single change that occurs in the system. The adaptation is initiated by source node s , which increases $h(s)$ by $2|V|$ and pushes flow to every node in σ_s . However, the adaptation can be performed in batch mode, i.e. source node s initiates the adaptation after multiple changes have occurred in the system. Since the proof of Theorem 2 does not utilize any information about the number of changes occurred, the correctness and complexity of the RIPR algorithm still holds even if the adaptation is performed in batch mode. We have observed that the RIPR algorithm always finds the optimal solution, regardless of the number of changes occurred before the adaptation is performed.

The result in Figure 2 illustrates the cost of adaptation (in terms of the total number of basic operations) vs the number of changes occurred before the adaptation. In each experiment, a randomly generated system with 40 nodes was deployed in a unit square. After the system stabilized and found the optimal solution, the bandwidth of a certain number of links was changed. Then adaptation was performed and the system stabilized again (and found a new optimal solution) after executing a certain number of basic operations. For each experiment, we record the number of basic operations executed by the system to find the new optimal solution. Each data point in Figure 2 is averaged over 50 experiments. We can see that the required number of basic operations increases as the number of changes (per adaptation) increases.

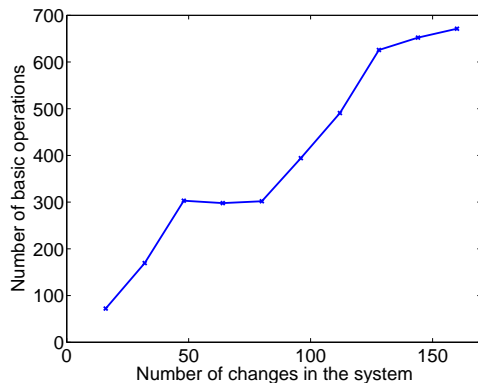


Figure 2. Adaptation performed in batch mode

So far the performance of the RIPR algorithm is evaluated in terms of the total number of basic operations. We do not expect the individual nodes to execute the same number of operations since the RIPR algorithm is not designed for load balancing. But interestingly, the following simulation results show that the RIPR algorithm is pretty well-balanced in terms of the number of basic operations executed by different nodes. For each experiment, a randomly generated system is initialized and the number of basic operations executed by the system to stabilize was recorded. The basic operations were reclassified into two categories: local updates and control message exchanges. Each $push(u, v)$ operation consists of one local update at u , one message transfer (send) at u , and one message transfer (receive) at v . Each $relabel(u)$ operation consists of one local update at u , one message transfer (broadcast $h(u)$ to σ_u) at u , and one message transfer (receive $h(u)$) at each $v \in \sigma_u$. Figure 3 shows the number of local updates and control message executed/transferred by the nodes. We report the maximum and the mean number of local updates, and the maximum and mean number of control message exchanges. Each data point is averaged over 100 experiments. Figure 3 shows that the maximum number of local updates is only about 2 times the mean number of local updates. The maximum number of control message exchanges is also about 2 times the mean number of control message exchanges. This result shows that

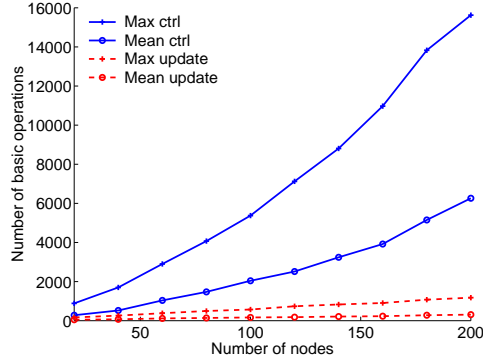


Figure 3. The maximum and mean cost per node for executing the RIPR algorithm

the RIPR algorithm is quite well-balanced in terms of per node cost.

The second set of simulation results illustrate the convergence and adaptivity of the proposed protocol. In each experiment, a certain number (between 40 and 100) of nodes were randomly deployed in the unit square. Communication radius ranging from 0.2 to 0.5 units were tested. For each experiment, the data gathering process lasted 30 seconds.

The *steady state throughput* is calculated as the average throughput during the last 10 seconds of data gathering. Table 1 shows the steady state throughput of the protocol. The results have been normalized to the optimal throughput. The optimal throughput was calculated offline. Each data point in Table 1 is averaged over 50 systems. The results show that the steady state throughput of the proposed protocol approaches the optimal throughput, regardless of the number of nodes and the communication radius.

In the protocol, data is transferred when the RIPR algorithm is being executed. Hence the start-up time of the system needs to be evaluated from two aspects: the *execution time* of the RIPR algorithm (i.e. how fast the RIPR algorithm terminates), and the time for the data transfer to reach steady state throughput.

For each experiment in the second set of simulations, we monitored the activities of each individual

Table 1. Normalized Steady-state Throughput. r is the communication radius. n is the number of nodes.

	$r = 0.2$	$r = 0.3$	$r = 0.4$	$r = 0.5$
$n = 40$	0.9641	0.9557	0.9446	0.9423
$n = 60$	0.9317	0.9322	0.9208	0.9075
$n = 80$	0.9239	0.9262	0.9186	0.9315
$n = 100$	0.9264	0.9184	0.9247	0.9080

node. The termination of the RIPR algorithm was detected when none of the nodes needed to execute any of the basic operations. Note that such a global monitoring is made available in the simulations for performance analysis only. It may be very costly to implement this monitoring function in actual deployment. Let $N(t)$ denote the number of data packets received by the base station from time 0 to time t . The *instantaneous throughput* at time instance t is defined as $(N(t + 0.1) - N(t - 0.1))/0.2$. The *start-up time of the protocol* is defined as the time period for the instantaneous throughput to reach 85% of the steady-state throughput.

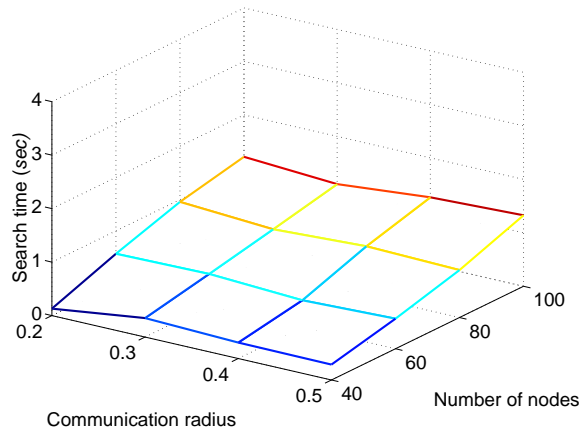


Figure 4. Execution time of the RIPR algorithm

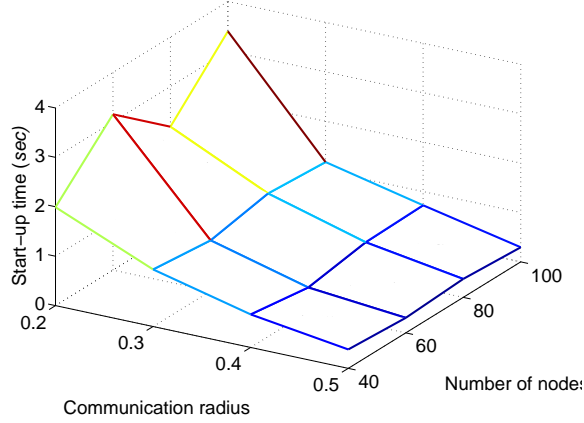


Figure 5. Start-up time of the proposed protocol

The impact of the number of nodes and the communication radius on the execution time of the RIPR algorithm is shown in Figure 4. The execution time increases as the number of nodes increases. The execution time also increases as the communication radius increases, which leads to an increase in the number of links in the system. Such a trend is expected from Theorem 2.

The start-up time of the protocol is shown in Figure 5. The result shows that for a given communication radius, the start-up time of the protocol increases as the number of nodes increases; and interestingly, for a given number of nodes, the start-up time decreases as the communication radius increases. Such a behavior is due to the fact that a larger communication radius leads to a smaller diameter of the graph. The diameter of a graph is defined as the largest distance (in terms of the number hops) between any two nodes in the graph. In systems with small diameter, the base station is closer to the source node. Hence the data can be transferred sooner to the base station during the start-up time.

We have also observed that in some experiments, the system throughput reached steady state even before the RIPR algorithm terminated. This is not a contradiction. Actually, when such scenarios occurred, the RIPR algorithm was pushing excessive flow (node u is said to have excessive flow when $e(u) > 0$, i.e. when u has more incoming flow than outgoing flow) back to the source node. During this

time period, the RIPR algorithm was still executing, but the net flow from the source to the sink did not increase. In other words, the RIPR algorithm had already found the maximum flow if the excessive flow had been eliminated. Meanwhile, data was transferred when the RIPR algorithm was still executing. Because each node maintained a data buffer which prevented the node from accumulating excessive data, the excessive flow did not cause the nodes to accumulate data. Consequently, the protocol was able to drive the system throughput to steady state before the RIPR algorithm terminated.

The above results illustrate the behavior of the the protocol and the RIPR algorithm. Awareness of such behaviors is useful for system synthesis. For example, in order to reduce the start-up time of the protocol, we can deploy the nodes so that they can reach the sink in a small number of hops. To reduce the cost (both time and energy) of executing the RIPR algorithm, we can restrict the communication of each node to a subset of its neighbors (thereby reducing $|E|$).

Note that the observed execution time of the RIPR algorithm (less than 1.3 seconds) and the start-up time of the protocol (less than 4.3 seconds) depends on the bandwidth settings of the links. In our simulations, the bandwidth of the links is around 10kbps, which is around 40 data packets per second because each data packet is 32 bytes. The shortest path (in terms of the transfer time of one data packet) from the source node to the base station ranges from 0.05 seconds to 0.13 seconds. The execution and start-up time will be much shorter if the links have higher bandwidth. For example, if the system is built with Telos [4] wireless sensors that can communicate at 250 kbps, we can expect about 20 times speed up in both the execution time of RIPR algorithm and the start-up time of the protocol.

Adaptivity of the proposed protocol is shown in Figure 6. The system consisted of 40 nodes randomly deployed in the unit square. Communication radius was set to 0.4. The system activities during the first 40 seconds are shown. At time $t = 20$ sec, we changed the bandwidth of a randomly selected set of

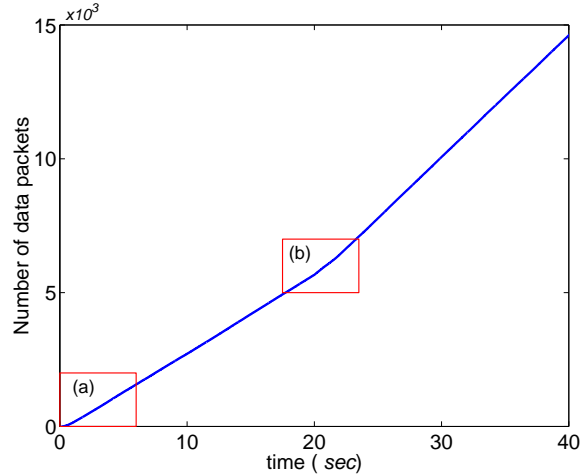


Figure 6. Illustration of the start-up and the adaptation of the proposed protocol. Framed block (a) is zoomed in figure 7(a), framed block (b) is zoomed in figure 7(b).

links, each of which was increased by 100%. Consequently, the optimal throughput (calculated off-line) changed from 314 to 492 (data packets/sec). As such changes occurred, the adaptation procedure was activated and the system operated at a new steady state throughput after the adaptation was completed. Figure 6 shows the number of data packets received by the base station as time advances. The throughput actually achieved by the protocol is reflected by the slope of the curve, which is 293 (93% of the optimal) before $t = 20$ and 452 thereafter (92% of the optimal). For this experiment, we define the start-up time as the time period for the instantaneous throughput to reach 85% of the first steady state throughput 293, starting from $t = 0$; and the adaptation time as the time period for the instantaneous throughput to reach 85% of the second steady state throughput 452, starting from $t = 20$. In this experiment, the shortest path (in terms of overall transfer time) to send a data packet (from the source node) to the base station consists of 3 hops and requires 0.06 sec. By using our protocol, the first data packet was received by the base station 0.12 seconds after the system started; the start-up time is 1.13 seconds; and the adaptation time

is 1.4 seconds. The system activities during the start-up and adaptation period are shown in more detail in Figure 7. An important observation from Figure 7 is that the system started (at $t = 0$) and continued (at $t = 20$) to gather data while the RIPR algorithm was still executing. The system did not wait until the optimal solution was found. Actually, because the protocol was executed in a distributed fashion, none of the nodes would know the completion of the RIPR algorithm unless a global synchronization was performed.

6 Conclusion

In this paper, we studied a set of data gathering problems in energy-constrained networked sensor systems. We reduced such problems to a network flow maximization problem with vertex capacity constraint, which can be further reduced to the standard network flow problem. After deriving a relaxed formulation for the standard network problem, we developed a distributed and adaptive algorithm to maximize the flow. This algorithm can be approximated as a simple data gathering protocol.

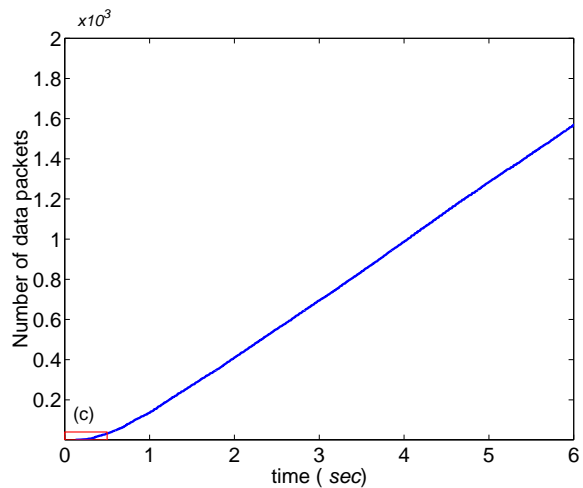
One of the future directions is to design distributed algorithms that do not generate excessive flow at the nodes (i.e. $e(u)$ does not become positive) during the execution. Our formulation of constrained flow optimizations can be applied to problems beyond the four problems discussed in this paper. For example, the system model considered in [11] gathers data in rounds. In each round, every sensor generates one data packet and the data packets from all the sensors need to be collected by the sink. The goal is to maximize the total number of rounds the system can operate under energy constraints on the nodes. This problem can be described by our constrained flow formulation and an optimal solution can be developed [10].

References

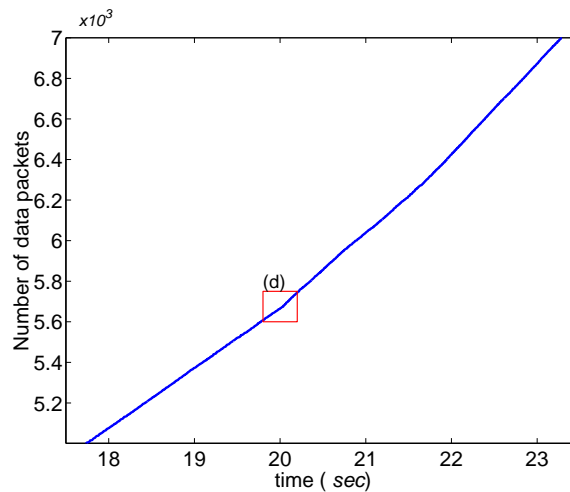
- [1] I. F. Akyildiz, W. Su, Y. Sankarasubramaniam, and E. Cyirci. Wireless Sensor Networks: A Survey. *Computer Networks*, 38(4):393–422, 2002.
- [2] G. Asada, T. Dong, F. Lin, G. Pottie, W. Kaiser, and H. Marcy. Wireless Integrated Network Sensors: Low Power Systems on a Chip. In *Proceedings of European Solid State Circuits Conference*, 1998.
- [3] T. H. Cormen, C. E. Leiserson, and R. L. Rivest. *Introduction to Algorithms*. MIT Press, 1992.
- [4] Moteiv Corporation. Telos Wireless Sensor Module. <http://www.moteiv.com>.
- [5] C. Efthymiou, S. Nikolettseas, and J. Rolim. Energy Balanced Data Propagation in Wireless Sensor Networks. *4th International Workshop on Algorithms for Wireless, Mobile, Ad Hoc and Sensor Networks (WMAN '04), hold in conjunction with IPDPS 2004*, April 2004.
- [6] E. Falck, P. Floreen, P. Kaski, J. Kohonen, and P. Orponen. Balanced Data Gathering in Energy-Constrained Sensor Networks. In *Sotiris Nikolettseas and Jose D. P. Rolim, editors, Algorithmic Aspects of Wireless Sensor Networks (First International Workshop, ALGOSENSORS 2004)*, July 2004.
- [7] A. V. Goldberg and R. E. Tarjan. A New Approach to the Maximum Flow Problem. *Journal of Association for Computing Machinery*, 35:921–940, 1988.
- [8] W. B. Heinzelman. An Application-Specific Protocol Architecture for Wireless Microsensor Networks. *IEEE Transactions on Wireless Communications*, 1(3), 2002.

- [9] W. R. Heinzelman, A. Chandrakasan, and H. Balakrishnan. Energy Efficient Communication Protocol for Wireless Micro-sensor Networks. In *Proceedings of IEEE Hawaii International Conference on System Sciences*, 2000.
- [10] B. Hong and V. K. Prasanna. Optimizing System Life time for Data Gathering in Networked Sensor Systems. *AlgorithmS for Wireless and Ad-hoc Networks (A-SWAN 2004) (Held in conjunction with MobiQuitous 2004)*, August 2004.
- [11] K. Kalpakis, K. Dasgupta, and P. Namjoshi. Maximum Lifetime Data Gathering and Aggregation in Wireless Sensor Networks. *IEEE Networks '02 Conference*, 2002.
- [12] E. L. Lawler. *Combinatorial Optimization : Networks and Matroids*. Holt, Rinehart and Winston, 1976.
- [13] R. Min, T. Furrer, and A. Chandrakasan. Dynamic Voltage Scaling Techniques for Distributed Microsensor Networks. *Workshop on VLSI (WVLSI '00)*, April 2000.
- [14] F. Ordonez and B. Krishnamachari. Optimal Information Extraction in Energy-Limited Wireless Sensor Networks. *to appear in IEEE Journal on Selected Areas in Communications, special issue on Fundamental Performance Limits of Wireless Sensor Networks, 2004*.
- [15] J. Rabaey, J. Ammer, T. Karalar, S. Li, B. Otis, M. Sheets, and T. Tuan. PicoRadios for Wireless Sensor Networks: The Next Challenge in Ultra-Low-Power Design. In *Proceedings of the International Solid-State Circuits Conference*, 2002.
- [16] N. Sadagopan and B. Krishnamachari. Maximizing Data Extraction in Energy-Limited Sensor Networks. *IEEE Infocom 2004*, 2004.

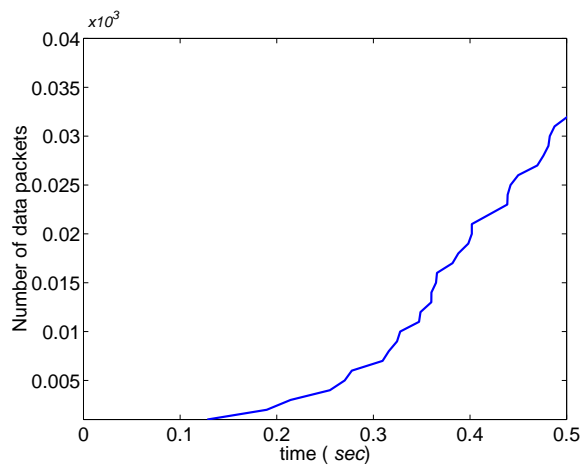
- [17] C. Schurgers, O. Aberthorne, and M. Srivastava. Modulation Scaling for Energy Aware Communication Systems. In *International Symposium on Low Power Electronics and Design*, August 2001.
- [18] M. Singh and V. K. Prasanna. Optimal Energy Balanced Algorithm for Selection in Single Hop Sensor Network. *IEEE International Workshop on Sensor Network Protocols and Applications (SNPA) ICC*, 2003.
- [19] S. Singh and C. Raghavendra. PAMAS: Power Aware Multi-Access protocol with Signalling for Ad Hoc Networks. *ACM Computer Communications Review*, 1998.
- [20] B. Warneke, M. Last, B. Liebowitz, and K. S. J. Pister. Smart Dust: Communicating with a Cubic-Millimeter Computer. *Computer*, 34(1):44–51, 2001.
- [21] Y. Yu and V. K. Prasanna. Energy-Balanced Task Allocation for Collaborative Processing in Networked Embedded System. *ACM Conference on Language, Compilers, and Tools for Embedded Systems (LCTES)*, 2003.



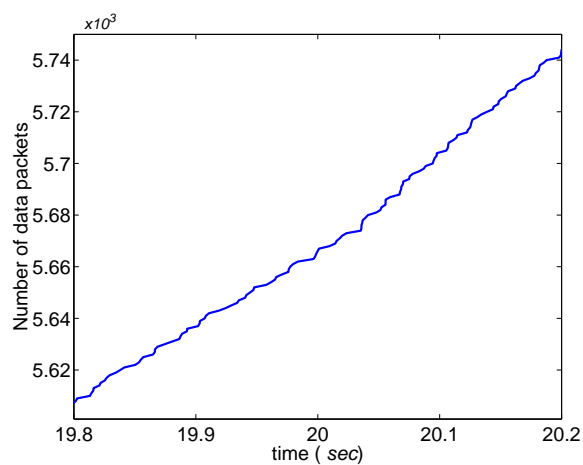
(a) $0 \leq t \leq 6$, the framed block is zoomed in (c)



(b) $17.5 \leq t \leq 23.5$, the framed block is zoomed in (d)



(c) $0 \leq t \leq 0.5$



(d) $19.8 \leq t \leq 20.2$

Figure 7. Detailed illustration of the start-up and the adaptation of the proposed protocol