Stabilization of Flood Sequencing Protocols in Sensor Networks

Young-ri Choi and Mohamed G. Gouda

Department of Computer Sciences, The University of Texas at Austin,
1 University Station C0500, Austin, Texas 78712-0233, U.S.A.
{yrchoi, gouda}@cs.utexas.edu

Abstract. Flood is a communication primitive that can be used by the base station of a sensor network to send a copy of a message to every sensor in the network. When a sensor receives a flood message, the sensor needs to check whether it has received the message for the first time and so the message is fresh, or it has received the same message earlier and so the message is redundant. In this paper, we discuss a family of four flood sequencing protocols that use sequence numbers to distinguish between fresh and redundant flood messages. They are a sequencing free protocol, a linear sequencing protocol, a circular sequencing protocol, and a differentiated sequencing protocol. We analyze the self-stabilization properties of these four flood sequencing protocols. We also compare the performance of these flood sequencing protocols, using simulation, over various settings of sensor networks. We conclude that the differentiated sequencing protocol has better stabilization property and provides better performance than those of the other three protocols.

Keywords: Self-stabilization, Flood sequencing protocol, Sequence numbers, Sensor networks.

1 Introduction

Flood is a communication primitive that can be used by the base station of a sensor network to send a copy of a message to every sensor in the network. The execution of a flood starts by the base station sending a message to all its neighbors. When a sensor receives a message, the sensor needs to check whether it has received this message for the first time or not. Only if the sensor has received the message for the first time, the sensor keeps a copy of the message and may forward it to all its neighbors. Otherwise, the sensor discards the message.

To distinguish between "fresh" flood messages that a sensor should keep and "redundant" flood messages that a sensor should discard, the base station selects a sequence number and attaches it to a flood message before the base station broadcasts the message. When a sensor receives a flood message, the sensor determines based on the sequence number in the received message whether the message is fresh or redundant. The sensor accepts the message if it is fresh and discards the message if it is redundant. We call a protocol that uses sequence

T. Masuzawa and S. Tixeuil (Eds.): SSS 2007, LNCS 4838, pp. 139–153, 2007.

numbers to distinguish between fresh and redundant flood messages a *flood sequencing protocol*.

In a flood sequencing protocol, when a fault corrupts the sequence numbers stored in some sensors in a sensor network, the network can enter an illegitimate state where the sensors discard fresh flood messages and accept redundant flood messages. Therefore, a flood sequencing protocol should be designed such that if the protocol ever reaches an illegitimate state due to some fault, the protocol is guaranteed to converge back to its legitimate states where every sensor accepts every fresh flood message and discards every redundant flood message.

In this paper, we discuss a family of four flood sequencing protocols. They are a *sequencing free* protocol, a *linear sequencing* protocol, a *circular sequencing* protocol, and a *differentiated sequencing* protocol. We analyze the stabilization properties of these four protocols. For each of the protocols, we first compute an upper bound on the convergence time of the protocol from an illegitimate state to legitimate states. Second, we compute an upper bound on the number of fresh flood messages that can be discarded by each sensor during the convergence. Third, we compute an upper bound on the number of redundant flood messages that can be accepted by each sensor during the convergence.

The rest of the paper is organized as follows. In Section 2, we discuss related work and motivation of the flood sequencing protocols. In Section 3, we present a model of sensor networks. In Section 4, we give an overview of a flood protocol. We present the four flood sequencing protocols in Sections 5, 6, 7, and 8. We analyze their stabilization properties and compare them with each other in Section 9. In Section 10, we show the simulation results of these protocols. We finally make concluding remarks in Section 11.

2 Related Work and Motivation

The practice of using sequence numbers to distinguish between fresh and redundant flood messages has been adopted by most flood protocols in the literature. In other words, most flood protocols “employ” some flood sequencing protocols to distinguish between fresh and redundant flood messages. A flood sequencing protocol can be designed in various ways, depending on several design decisions such as how the next sequence number is selected by the base station, how each sensor determines based on the sequence number in a received message whether the received message is fresh or redundant, and what information the base station and each sensor stores in its local memory. Unfortunately, flood sequencing protocols have been used without full investigation of their design decisions in the literature.

The flood protocols discussed in [1,2,3,4] assume that when a sensor receives a flood message, the sensor can figure out whether the sensor has received this message for the first time or not, without specifying any mechanism to achieve this. In [5,6], it was suggested to associate a sequence number with each flood message, but any details on how sequence numbers are used by sensors (i.e. the design decisions of their flood sequencing protocols) were not specified. The flood protocols discussed in [7,8] propose to attach a unique identifier to each flood message and make each sensor maintain a list of identifiers that the sensor has received recently. Similarly, it was suggested in [9] that each sensor maintains a list of flood messages received by the sensor recently. However, any details such as how many identifiers or messages each sensor maintains and when a sensor deletes an identifier or a message from the list were not discussed.

A flood sequencing protocol is important, since the fault tolerance property of a sensor network is affected by a flood sequencing protocol used in the network. When a fault corrupts the sequence number stored in some sensor in the network, the sensor may discard fresh flood messages and accept redundant flood messages. The number of fresh flood messages discarded by the sensor and the number of redundant flood messages accepted by the sensor, before the network reaches a legitimate state, are different depending on which flood sequencing protocol is used in the network. Therefore, we need to study various flood sequencing protocols and analyze the stabilization properties of these protocols. The stabilization properties of the flood sequencing protocols are useful for sensor network designers or developers to select a proper flood sequencing protocol that satisfies the needs of a target sensor network.

In practice, a flood sequencing protocol is used with a flood protocol that may use other techniques to improve the performance of flood such as reliability or efficiency. In this paper, each of the flood sequencing protocols is described focusing on how sequence numbers are used by sensors, and it is not described as a specific flood protocol. Note that the stabilization property of a flood protocol is affected by that of a flood sequencing protocol used in the flood protocol. If the flood protocol does not maintain any extra state such that it is based on probability [2,5], the stabilization property of the flood protocol is the same as that of the used flood sequencing protocol. If the flood protocol maintains extra state such that it is based on neighbor information [1,5], the stabilization property of the flood protocol also depends on how the extra state in each sensor is stabilized.

3 Model of Sensor Networks

In this section, we describe a formal model of the execution of a sensor network, which was introduced first in [10]. This model accommodates several characteristics of sensor networks such as unavoidable local broadcast, probabilistic message transmission, asymmetric communication, message collision, and timeout actions and randomization steps. We use the model to specify our flood sequencing protocols, verify the stabilization properties of these protocols, and develop our simulation of these protocols.

The *topology* of a sensor network is a directed graph where each node represents a distinct sensor in the network and where each directed edge is labeled with some probability. A directed edge (u,v), from a sensor u to a sensor v, that is labeled with probability p (where $p > 0$) indicates that if sensor u sends a message, then this message arrives at sensor v with probability p (provided that neither sensor v nor any "neighboring sensor" of v sends another message at the same time). If the topology of a sensor network has a directed edge from a sensor u to a sensor v, then u is called an *in-neighbor* of v and v is called an *out-neighbor* of u.

We assume that during the execution of a sensor network, the real-time passes through discrete instants: instant 1, instant 2, instant 3, and so on. The time periods between consecutive instants are equal. The different activities that constitute the execution of a sensor network occur only at the time instants, and not in the time periods between the instants. We refer to the time period between two consecutive instants t and $t + 1$ as a *time unit* $(t,\ t + 1)$. (The value of a time unit is not critical to our current presentation of a sensor network model, but we estimate that the value of the time unit is around 100 milliseconds.)

A sensor is specified as a program that has global constants, local variables, one timeout action, and one receiving action.

At a time instant t, if the timeout of a sensor u expires, then u executes its timeout action at t. Executing the timeout action of sensor u at t causes u to update its local variables, and to send at most one message at t. It also causes u to execute the statement "timeout-after <expression>" which causes the timeout of u to expire (again) after k time units, where k is the value of <expression> at the time unit $(t,\ t + 1)$. The timeout action of sensor u is of the following form:

```
timeout-expires ->
    <update local variables of u>;
    <send at most one message>;
    <execute timeout-after <expression>>
```

To keep track of its timeout, each sensor u has an implicit variable named "timer.u". In each time unit between two consecutive instants, timer.u has a fixed positive integer value. If the value of timer.u is k, where $k > 1$, in a time unit $(t - 1,\ t)$, then the value of timer.u is $k - 1$ in the time unit $(t,\ t + 1)$. On the other hand, if the value of timer.u is 1 in a time unit $(t - 1,\ t)$, then sensor u executes its timeout action at instant t. Moreover, since sensor u executes the statement "timeout-after <expression>" as part of executing its timeout action, the value of timer.u in the time unit $(t,\ t + 1)$ is the value of <expression> in the same time unit.

If a sensor u executes its timeout action and sends a message at an instant t, then an out-neighbor v of u receives a copy of the message at t, provided that the following three conditions hold.

i. A random integer number is uniformly selected in the range 0 .. 99, and this selected number is less than $100 * p$, where p is the probability label of edge (u,v) in the network topology.
ii. Sensor v does not send any message at instant t.
iii. For each in-neighbor w of v, other than u, if w sends a message at t, then a random integer number is uniformly selected in the range 0 .. 99, and this selected number is at least $100 * p'$, where p' is the probability label of edge (w,v) in the network topology.

If v sends a message at t, or if w sends a message at t and for v, selects a random number that is less than 100*p', then this message *collides* with the message sent by u with the net result that v receives no message at t.

If a sensor u receives a message at instant t, then u executes its receiving action at t. Executing the receiving action of sensor u causes u to update its own local variables. It may also cause u to execute the statement "timeout-after <expression>". The receiving action of sensor u is of the following form:

```
rcv <msg> ->
    <update local variables of u>;
    <may execute timeout-after <expression>>
```

A *state* of a sensor network protocol is defined by a value for each variable and timer.u for each sensor u in the protocol. We use the notation <var>.u to denote the value of variable <var> at some sensor u.

During the execution of a sensor network protocol, several faults can occur, resulting in corrupting the state of the protocol arbitrarily. Examples of these faults are wrong initialization, memory corruption, message corruption, and sensor failure and recovery. We assume that these faults do not continuously occur in the network.

4 Overview of a Flood Protocol

In this section, we give an overview of a flood protocol that is used with our flood sequencing protocols. Consider a network that has n sensors. In this network, sensor 0 is the base station and can initiate floods over the network. To initiate the flood of a message, sensor 0 sends a message of the form data($hmax$), where $hmax$ is the maximum number of hops to be made by this data message in the network.

If sensor 0 initiates one flood and shortly after initiates another flood, some forwarded messages from these two floods can collide with one another causing many sensors in the network not to receive the message of either flood, or (even worse) not to receive the messages of both floods.

To prevent message collision across consecutive flood messages, once sensor 0 broadcasts a message, it needs to wait enough time until this message is no longer forwarded in the network, before broadcasting the next message. The time period that sensor 0 needs to wait after broadcasting a message and before broadcasting the next message is called the *flood period.* The flood period consists of f time units. A lower bound on the value of f is computed as $(hmax - 1) * tmax + 1$. (Due to space limit, this bound is presented without proof. We refer the reader to [11] for proof.) Thus, after sensor 0 broadcasts a message, it sets its timeout to expire after f time units in order to broadcast the next message.

When a sensor receives a data(h) message, the sensor decides whether the sensor accepts the message and forwards it as a data($h - 1$) message, provided $h > 1$. To reduce the probability of message collision, any sensor u, that decides to forward a message, chooses a random period whose length is chosen uniformly from the range $1..tmax$, and sets its timeout to expire after the chosen random period, so that u can forward the received message at the end of the random period. This random time period is called the *forwarding period*.

To analyze each of the four flood sequencing protocols, we use the following value for the flood period f:

$$f = hmax * tmax + 1.$$

(We choose this value for f, instead of the minimum value $(hmax-1)*tmax+1$, to keep our proofs of the stabilization properties simple.)

Note that the above flood period is computed to guarantee that no two consecutive flood messages ever collide with each other. In a typical execution of the protocol, each sensor chooses its forwarding period at random in the range $1..tmax$, and so most sensors likely receive the flood messages within $(hmax-1)* tmax/2$ time units, instead of $(hmax-1)*tmax$ time units. Therefore, the half (or even less) of the flood period may be used without significantly degrading the stabilization property and performance of a flood sequencing protocol.

5 First Protocol: Sequencing Free

In this section, we discuss a first flood sequencing protocol where no sequence number is attached to each flood message, and so a sensor cannot distinguish between fresh and redundant flood messages, resulting that the sensor accepts every received message. This protocol is called the *sequencing free* protocol.

To initiate the flood of a new message, sensor 0 sends a data($hmax$) message, and then sets its timeout to expire after f time units to broadcast the next message.

Each sensor u that is not sensor 0 maintains a variable called new. The value of new is true only when u is in the forwarding period (i.e. u has a flood message that has been received earlier but has not been forwarded yet). When sensor u receives a data(h) message, u always accepts the message. Sensor u forwards the message as data($h-1$), if $h > 1$ in the received message and $new = false$ in u. (A formal specification of the sequencing free protocol can be found in [11].)

Note that in all the flood sequencing protocols presented in this paper, the value of timer.0 is at most f time units, and the value of timer.u is at most $tmax$. This is maintained by the executions of all the protocols.

A state S of the sequencing free protocol is *legitimate* iff either S is a state where the predicate

(timer.0= 1) $\wedge$ (for all u, $u \neq 0$, new.u=$false$)

holds or S is a state that is reachable from a state, where this predicate holds, by some execution of the protocol.

It follows from this definition that if the protocol is executed starting from a legitimate state, then every time sensor 0 initiates a new flood, previous flood messages (whether initiated by sensor 0 legitimately or other sensors illegitimately due to some fault) are no longer forwarded in the network.

6 Second Protocol: Linear Sequencing

In this section, we discuss a second flood sequencing protocol where each flood message carries a unique sequence number that is linearly increased, and so a sensor accepts a flood message that has a sequence number larger than the last sequence number accepted by the sensor. This protocol is called the *linear sequencing* protocol.

```
sensor 0                                  {base station}
const  hmax    : integer,                 {max hop count}
       f       : integer                  {flood period}
var    slast   : integer                  {last seq number}
begin
       timeout-expires →   {generate new msg}
                           slast := slast + 1;
                           send data(hmax,slast);
                           timeout-after f
end
```

Fig. 1. A specification of sensor 0 in the linear sequencing protocol

Each flood message in this protocol is of the form data(h,s), where field h is the remaining number of hops to be made by this message, and field s is the unique sequence number of this message.

Whenever sensor 0 broadcasts a new message, sensor 0 increases the sequence number of the last message by one, and attaches the increased sequence number to the message. A formal specification of sensor 0 is given in Fig. 1.

Each sensor u that is not sensor 0 keeps track of the last sequence number accepted by u in a variable called *slast*. When sensor u receives a data(h, s) message, sensor u accepts the message if $s > slast$, and forwards it if $h > 1$. A formal specification of sensor u is given in Fig. 2. (Each sensor u also maintains a received data message that u will forward later, even though this is not explicitly specified in the specification.)

A state S of the linear sequencing protocol is *legitimate* iff either S is a state where the predicate

$$(\text{timer.0}{=}\ 1) \wedge (\text{for all } u,\ u \neq 0,\ \text{new.}u{=}false \wedge \text{slast.}u \leq \text{slast.0})$$

holds or S is a state that is reachable from a state, where this predicate holds, by some execution of the protocol.

It follows from this definition that if the protocol is executed starting from a legitimate state, then every time sensor 0 initiates a new flood, previous flood messages are no longer forwarded in the network, and the new flood message has a sequence number that is larger than every slast.u in the network, so that every u accepts the message.

```
sensor u:1 .. n − 1
const hmax      : integer,       {max hop count}
      tmax      : integer        {max forwarding period}
var   h,hlast   : 1 .. hmax,     {rcvd,last hop count}
      s, slast  : integer,       {rcvd,last seq number}
      new       : boolean        {true if u has msg to forward}
begin
   timeout-expires →   if new →        new := false;
                                       send data(hlast, slast)
                       [] ¬ new →      skip
                       fi; timeout-after random(1,tmax)

   [] rcv data(h, s) → if s > slast →{accept msg} slast := s;
                          if h>1 →     new := true;
                                       hlast := h − 1
                          [] h≤ 1 →    skip
                          fi
                       [] s ≤ slast →    {discard msg} skip
                       fi
end
```

Fig. 2. A specification of sensor u in the linear sequencing protocol

7 Third Protocol: Circular Sequencing

In this section, we discuss a third flood sequencing protocol where each flood message carries a sequence number that is circularly increased within a limited range, and so a sensor accepts a flood message that has a sequence number "logically" larger than the last sequence number accepted by the sensor. This protocol is called the *circular sequencing* protocol.

Each flood message is augmented with a sequence number that has a value in the range 0 .. $smax$, where $smax > 1$. We assume that $smax$ is an even number (to keep our presentation simple).

Whenever sensor 0 broadcasts a new message, sensor 0 increases the sequence number of the last message by one circularly within the range 0 .. $smax$, i.e. $slast := (slast + 1) \bmod (smax{+}1)$, and attaches the increased sequence number to the message.

From the viewpoint of each sequence number s in the range 0 .. $smax$, the range can be divided into two subranges, where one subrange consists of the sequence numbers that are logically "smaller" than s, and the other subrange consists of the sequence numbers that are logically "larger" than s. Thus, sequence number s has $\frac{smax}{2}$ numbers logically smaller than it and $\frac{smax}{2}$ numbers logically larger than it. For example, if $smax = 8$, number 0 is logically smaller than 1, 2, 3, and 4, and is logically larger than 5, 6, 7, and 8.

When a sensor u receives a data(h, s) message, sensor u checks if s is logically larger than $slast$. Sensor u calls the function "Larger$(s, slast)$" that returns true if s is logically larger than $slast$, and otherwise returns false. Sensor u accepts the message if Larger$(s, slast)$ returns true, and forwards it if $h > 1$. Otherwise, sensor u discards the message.

To prove the stabilization property of the circular sequencing protocol, we make an assumption of bounded message loss as follows:

> *Bounded message loss*: Starting from any state, if sensor 0 broadcasts $\frac{smax}{2}$ consecutive flood messages, then every sensor in the network receives at least one of those flood messages.

Two explanations concerning the above assumption are in order. First, the protocol may not be self-stabilizing without any bound on message loss. For example, consider a scenario where $smax$=8. Assume that sensor 0 sends a flood message with sequence number 0 and a sensor u accepts the message. If sensor u does not receive the next 4 (i.e. $\frac{smax}{2}$) consecutive messages with sequence numbers 1, 2, 3 and 4, and later receives a fresh message with sequence number 5, it discards the message since sequence number 5 is not logically larger than sequence number 0. Sensor u also discards the next flood messages with sequence numbers 6, 7, 8, and 0, if it receives them. In this scenario, if sensor u does not receive flood messages with sequence numbers 1, 2, 3 and 4, it keeps discarding fresh flood messages. Thus, some assumption of bounded message loss is necessary for the stabilization property of the protocol.

Second, the above assumption becomes acceptable if the value of $smax$ is reasonably large enough for a given network setting. Selecting an appropriate value for $smax$ depends on the size of the network, the topology of the network, and a flood sequencing protocol used in the network. (In Section 10, we show how different values are selected for $smax$ depending on these factors.)

A state S of the circular sequencing protocol is *legitimate* iff either S is a state where the predicate

$$
\begin{aligned}
&(\text{timer}.0{=}1) \wedge \\
&(\text{for all } u,\ u \neq 0, \\
&\qquad (\text{new}.u{=}\text{false}) \wedge \\
&\qquad (\text{slast}.u = \text{slast}.0 \vee \\
&\qquad\ \ \text{slast}.u = (\text{slast}.0{-}1) \bmod (smax{+}1) \vee \\
&\qquad\ \ \ldots \\
&\qquad\ \ \text{slast}.u = (\text{slast}.0{-}\tfrac{smax}{2}{+}1) \bmod (smax{+}1) \\
&\qquad) \\
&) \wedge \\
&(\text{sensor 0 has already initiated at least } \tfrac{smax}{2} + 2 \text{ floods})
\end{aligned}
$$

holds or S is a state that is reachable from a state, where this predicate holds, by some execution of the protocol.

It follows from this definition that if the protocol is executed starting from a legitimate state, then every time sensor 0 initiates a new flood, previous flood messages are no longer forwarded in the network, and the new flood message has a sequence number that is logically larger than every slast.u in the network, so that every u accepts the message.

8 Fourth Protocol: Differentiated Sequencing

In this section, we discuss the last flood sequencing protocol where the sequence numbers of flood messages are in a limited range, similar to the circular sequencing protocol. However, in this protocol, a sensor accepts a flood message if the sequence number of the message is different from the last sequence number accepted by the sensor. This protocol is called the *differentiated sequencing* protocol.

Each flood message is augmented with a sequence number that has a value in the range 0 .. $smax$, where $smax > 0$. We assume that $smax$ is an even number (to keep our presentation simple).

Sensor 0 in this protocol is identical to the one in the circular sequencing protocol. However, when a sensor u receives a data(h, s) message, sensor u accepts the message if s is different from $slast$, i.e. $s \neq slast$, and forwards the message if $h > 1$. Otherwise, sensor u discards the message.

Similar to the circular sequencing protocol, if a sensor does not receive a large number of consecutive flood messages, the differentiated sequencing protocol may not be self-stabilizing. Thus, the proofs of the stabilization property of this protocol are based on the assumption of bounded message loss described in Section 7.

A state S of the differentiated sequencing protocol is *legitimate* iff either S is a state where the predicate

(timer.0=1) $\wedge$
(for all u, $u \neq 0$,
  (new.u=false) $\wedge$
  (slast.u = slast.0 $\vee$
   slast.u = (slast.0−1) mod ($smax$+1) $\vee$
   ...
   slast.u = (slast.0$-\frac{smax}{2}$+1) mod ($smax$+1)
)
)

holds or S is a state that is reachable from a state, where this predicate holds, by some execution of the protocol.

It follows from this definition that if the protocol is executed starting from a legitimate state, then every time sensor 0 initiates a new flood, previous flood messages are no longer forwarded in the network, and the new flood message has a sequence number that is different from every slast.u in the network, so that every u accepts the message.

9 Stabilization of the Protocols

In this section, we analyze the stabilization properties of the four flood sequencing protocols. For each of the protocols, we first compute an upper bound on the

Table 1. Stabilization properties of the flood sequencing protocols

	Convergence time (time units)	Max # of fresh msgs discarded by u until convergence	Max # of redundant msg accepted by u until convergence	Stabilization property
free	$2 * f$	0	$2 * f$	good
lin	unbounded	unbounded	$n - 1$	bad
cir	$(smax + 2) * f$	$(smax + 2) * f$	$f + 1$	good
dif	$(\frac{smax}{2} + 2) * f$	$(\frac{smax}{2} + 2) * f$	$f + 1$	good

convergence time of the protocol from an illegitimate state to legitimate states. Second, we compute an upper bound on the number of fresh flood messages that can be discarded by each sensor during the convergence. Third, we compute an upper bound on the number of redundant flood messages that can be accepted by each sensor during the convergence.

The stabilization properties of the four protocols are shown in Table 1. (Due to space limit, we present the stabilization properties without proof. We refer the reader to [11] for proof.) We also analyze the properties of the protocols after convergence (or starting from a legitimate state) in Table 2. We call these properties the *stable properties* of the protocols. In these tables, "free", "lin", "cir", and "dif" represent the sequencing free, linear sequencing, circular sequencing, and differentiated sequencing protocols, respectively. Note that the properties of the circular sequencing and differentiated sequencing protocols are analyzed under the assumption of bounded message loss.

Starting from an illegitimate state, the sequencing free protocol converges to legitimate states faster than the other three protocols do. However, even starting from any legitimate state, a sensor cannot distinguish between fresh and redundant messages, and so the sensor accepts every received message. The number of redundant copies of the same message accepted by a sensor depends on the value of *hmax* and the network topology. In worst case, the sensor can accept a redundant copy of the same message at each time instant during the flood period of the message. Thus, starting from any legitimate state, every sensor accepts at most f redundant copies of the same message.

In the linear sequencing protocol, sensors are required to use unbounded sequence numbers. Thus, this protocol is very expensive to implement for sensor networks that have limited resources. However, once the protocol starts its execution from any legitimate state, every sensor accepts every fresh message and discards every redundant message under any degree of message loss. On the other hand, in the circular sequencing and differentiated sequencing protocols, sensors use bounded sequence numbers. Thus, starting from any legitimate state, every sensor accepts every fresh message and discards every redundant message under the assumption of bounded message loss.

From the above results, we conclude that overall the differentiated sequencing protocol has better stabilization and stable properties than those of the other three protocols.

Table 2. Stable properties of the flood sequencing protocols

	Max # of fresh msgs discarded by u after convergence	Max # of redundant copies of the same msg accepted by u after convergence	Stable property
free	0	f	bad
lin	0	0	good
cir	0	0	good
dif	0	0	good

10 Simulation Results

We have developed a simulator that can simulate the execution of the four flood sequencing protocols, based on our model described in Section 3. In this simulator, a network is an $N * N$ grid where N is the number of sensors in each side of the grid, and the distance between a sensor (i, j) and each of $(i + 1, j)$, $(i, j + 1)$, $(i - 1, j)$, and $(i, j - 1)$, if it exists, where $0 \leq i, j < N$, is 1.

For the purpose of simulation, sensor 0 is (0,0) which is located at the left-bottom conner in a grid, and the following two types of topologies that have different network density were used.

- A topology for a sparse network: The edge probability between two sensors is labeled with a high probability 0.95 if their distance is at most 1, and with a low probability 0.5 if their distance is larger than 1 and less than 2. Otherwise, there is no edge between the two sensors. In this topology, each sensor (i,j) that is not on or near the boundary of the grid generally has 8 neighbors.
- A topology for a dense network: The edge probability between two sensors is labeled with probability 0.95 if their distance is at most 1.5, and with probability 0.5 if their distance is larger than 1.5 and less than 3. Otherwise, there is no edge between the two sensors. In this topology, each sensor (i,j) that is not on or near the boundary of the grid generally has 24 neighbors.

(The used probabilities, 0.95 and 0.5, were chosen based on some experiments on sensors. We refer the reader to [10] for details.)

The performance of a flood sequencing protocol can be measured by the following two metrics:

i. *Reach:* The percentage of sensors that receive a message sent by sensor 0.
ii. *Communication:* The total number of messages forwarded by all sensors in the network.

We ran simulations of the four flood sequencing protocols, and measured the above two metrics in 10*10 and 20*20 grids for both sparse and dense network topologies. In our simulations, we do not consider other techniques that can improve the performance of a flood protocol based on extra information such as probability, location, and neighbor information.

Table 3. Performance of the sequencing free and linear sequencing protocols

	sparse 10*10			sparse 20*20			dense 10*10			dense 20*20		
	hmax	Reach	Com.	hmax	Reach	Com.	hmax	Reach	Com.	hmax	Reach	Com.
free	13	99%	351.3	27	99.2%	2885.7	7	99.8%	200.5	13	99%	1262
lin	15	98.5%	97.8	28	98.5%	390.3	7	98.5%	87.5	14	98.8%	376.4

First, we studied the performance of the sequencing free protocol and the linear sequencing protocol starting from a legitimate state. The result of each simulation in this study represents the average value over the simulations of 100,000 floods.

Table 3 shows the reach and communication of the sequencing free and linear sequencing protocols in sparse and dense networks. In these simulations, $tmax = 6$ was used for a sparse network, and $tmax = 7$ was used for a dense network. From the above results, one can observe that the sequencing free protocol requires the sensors to send much more messages than those that the linear sequencing protocol does. Note that when the value of $smax$ is reasonably large for a given network setting, the performance of the circular sequencing and differentiated sequencing protocols becomes same as that of the linear sequencing protocol.

Next, we studied the stabilization properties of the circular sequencing and differentiated sequencing protocols, and their performance while stabilizing. We simulated the sequences of floods starting from 1000 different illegitimate states, and computed the average reach for each i-th flood. We attempted to select an appropriate value for $smax$ for each network setting such that the assumption of bounded message loss becomes acceptable, while the convergence time of each protocol is minimized.

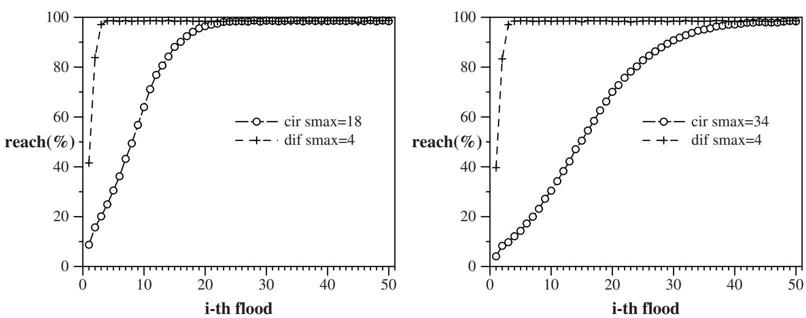

(a) A 10*10 network (b) A 20*20 network

Fig. 3. Reach of the circular and differentiated sequencing protocols starting from an illegitimate state in a sparse network

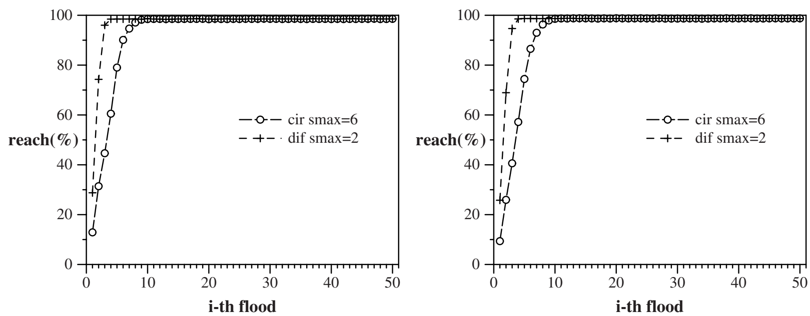

(a) A 10*10 network (b) A 20*20 network

Fig. 4. Reach of the circular and differentiated sequencing protocols starting from an illegitimate state in a dense network

Figures 3 and 4 show the reach of the circular sequencing and differentiated sequencing protocols starting from an illegitimate state in a sparse network and in a dense network, respectively. During the convergence time, each sensor has a higher probability to accept a received fresh message in the differentiated sequencing protocol than that in the circular sequencing protocol. Thus, in all simulated network settings, the differentiated sequencing protocol reaches a legitimate state faster than the circular sequencing protocol does.

In summary, starting from a legitimate state, the performance of any flood sequencing protocol that attaches a sequence number to a flood message is better than that of the sequencing free protocol in terms of communication. Starting from an illegitimate state, the differentiated sequencing protocol converges to a legitimate state quickly in all simulated network settings.

11 Concluding Remarks

In this paper, we discussed a family of the four flood sequencing protocols, namely the sequencing free, linear sequencing, circular sequencing, and differentiated sequencing protocols. We analyzed the stabilization and stable properties of these four protocols, and also studied their performance, using simulation, over various settings of sensor networks. We concluded that the differentiated sequencing protocol has better overall performance in terms of communication and stabilization and stable properties compared to those of the other three protocols.

Acknowledgment

This work was supported in part by the US National Science Foundation under Grant No. 0520250.

References

1. Stojmenovic, I., Seddigh, M., Zunic, J.: Dominating Sets and Neighbor Elimination-Based Broadcasting Algorithms in Wireless Networks. IEEE Transactions on Parallel and Distributed Systems 13(1), 14–25 (2002)
2. Sasson, Y., Cavin, D., Schiper, A.: Probabilistic Broadcast for Flooding in Wireless Mobile Ad hoc Networks. In: WCNC 2003. Proceedings of IEEE Wireless Communications and Networking Conference, pp. 1124–1130. IEEE Computer Society Press, Los Alamitos (2003)
3. Li, J., Mohapatra, P.: A Novel Mechanism for Flooding Based Route Discovery in Ad Hoc Networks. In: GLOBECOM. Proceedings of the IEEE Global Telecommunications Conference, pp. 692–696 (2003)
4. Ganesan, D., Krishnamurthy, B., Woo, A., Culler, D., Estrin, D., Wicker, S.: An Empirical Study of Epidemic Algorithms in Large Scale Multihop Wireless Networks. IRP-TR-02-003 (2002)
5. Ni, S., Tseng, Y., Chen, Y., Sheu, J.: The Broadcast Storm Problem in a Mobile Ad Hoc Network. In: MOBICOM. Proceedings of the ACM/IEEE International Conference on Mobile Computing and Networking, pp. 151–162. IEEE Computer Society Press, Los Alamitos (1999)
6. Heissenbttel, M., Braun, T., Waelchli, M., Bernoulli, T.: Optimized Stateless Broadcasting in Wireless Multi-hop Networks. In: IEEE INFOCOM, pp. 1–12 (2006)
7. Williams, B., Camp, T.: Comparison of Broadcasting Techniques for Mobile Ad Hoc Networks. In: MOBIHOC. Proceedings of the ACM International Symposium on Mobile Ad Hoc Networking and Computing, pp. 194–205. ACM Press, New York (2002)
8. Johnson, D.B., Maltz, D.A.: Dynamic Source Routing in Ad Hoc Wireless Networks. In: Imielinski, T., Korth, H. (eds.) Mobile Computing, vol. 353, pp. 153–181. Kluwer Academic Publishers, Dordrecht (1996)
9. Sun, M., Feng, W., Lai, T.: Location Aided Broadcast in Wireless Ad Hoc Networks. In: Proceedings of the IEEE GLOBECOM 2001, pp. 2842–2846. IEEE Computer Society Press, Los Alamitos (2001)
10. Gouda, M., Choi, Y.: A State-based Model of Sensor Protocols. In: Anderson, J.H., Prencipe, G., Wattenhofer, R. (eds.) OPODIS 2005. LNCS, vol. 3974, pp. 246–260. Springer, Heidelberg (2006)
11. Choi, Y., Gouda, M.: Stabilization of Flood Sequencing Protocols in Sensor Networks. Technical Report TR-06-58, Department of Computer Sciences, The University of Texas at Austin (2006)