

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
1 September 2011 (01.09.2011)

(10) International Publication Number
WO 2011/104620 A1

- (51) **International Patent Classification:**
H04L 12/56 (2006.01) *H04L 12/28* (2006.01)
- (21) **International Application Number:**
PCT/IB2011/000393
- (22) **International Filing Date:**
25 February 2011 (25.02.2011)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
2010/01400 26 February 2010 (26.02.2010) ZA
- (71) **Applicant (for all designated States except US):** **UNIVERSITY OF CAPE TOWN** [ZA/ZA]; Lovers Walk, Rondebosch, Cape Town 7700 (ZA).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** **LUSILAO-ZODI, Guy-Alain** [CD/ZA]; 32 Fairglen Mews, Heathfield, Bergvliet, 7945 (ZA). **MORRISON, Norman** [ZA/ZA]; A7 Ridgmont, Horak Avenue, Camps Bay, 8005 (ZA).
- (74) **Agent:** **VON SEIDELS**; P O Box 440, 7446 Century City, (Cape Town) (ZA).
- (81) **Designated States (unless otherwise indicated, for every kind of national protection available):** AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO,

DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) **Designated States (unless otherwise indicated, for every kind of regional protection available):** ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) **Title:** A SYSTEM AND METHOD FOR ESTIMATING ROUND-TRIP TIME IN TELECOMMUNICATION NETWORKS

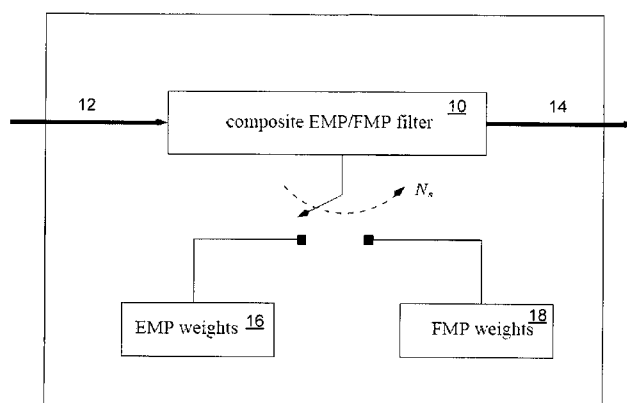


Figure 1

(57) **Abstract:** A method for estimating round-trip time (RTT) values for data packets travelling in a telecommunication network is disclosed. During a first time interval (206) which begins with the start of a data transmission, estimated RTT values (14) are generated by filtering a periodic sequence of RTT samples (12) with a self-initialising expanding memory polynomial (EMP) filter (16) and, after a first switch-over time at the end of the first time interval, the estimated RTT values are generated by filtering the RTT samples with a fading memory polynomial (FMP) filter (18). The first switch-over time is determined by comparing an FMP estimation error and an EMP estimation error, where the FMP estimation error is the difference between an output of the FMP filter and the RTT samples, and the EMP estimation error is the difference between an output of the EMP filter and the RTT samples, and triggering the first switch-over time when the FMP estimation error becomes equal to or less than the EMP estimation error.



WO 2011/104620 A1

A SYSTEM AND METHOD FOR ESTIMATING ROUND-TRIP TIME IN TELECOMMUNICATION NETWORKS

FIELD OF THE INVENTION

This invention relates to round-trip time estimation in telecommunication
5 networks, and more particularly to a system and method for estimating
round-trip time in packet switched telecommunication networks.

BACKGROUND TO THE INVENTION

- 10 The Internet is increasingly being used for the transmission of high volumes
of multimedia data generated from applications like video on demand, video
telephony and videoconferencing. Reliable transmission of video over packet
switched telecommunication networks such as the Internet requires stringent
bandwidth demand as well as low delay and packet loss requirements.
- 15 Currently, the Internet does not include any particular function in the network
layer to support these requirements. Instead, the service quality functions are
provided for at the application layer of the protocol stack or in the
Transmission Control Protocol (TCP) found at the stack's transport layer.
- 20 TCP incorporates algorithms that are implemented on top of the network
layer to control network congestion by adjusting the transmission rate based
on the packet loss ratio and the round-trip time (RTT) experienced by data
during transmission. To regulate network congestion it is therefore important
to have an accurate estimation of the round-trip time. Other protocols that
25 regulate real-time network traffic also make use of RTT estimates.

The round-trip time is the time interval that elapses between the sending of a

packet and the receipt of its acknowledgement. Individual RTT samples can be obtained by a protocol that sends packets from the source to destination and tracks their delivery. Although it is fairly straightforward to design such a protocol, the accuracy of individual samples is not a very good representation of the actual RTT due to various noise sources such as the difference in synchronization between the sender's and the receiver's clocks. Rather than using the sampled value directly, most traffic controllers use a prediction of the future values, as the predicted values give better results than the sampled values.

10

Prediction of round-trip time is complicated by the noisy and nonlinear nature of the sampled values. Currently, TCP predicts future round trip times by sampling the behaviour of packets sent over a connection and averaging those samples into a "smoothed" round-trip time estimate value using a digital filter. The digital filter used in TCP belongs to the family of Fading Memory Polynomial (FMP) filters of degree zero. FMP filters produce an estimated value based on a weighted sequence of prior values and prior estimated values. FMP filters are not self-initialising, and accordingly the values predicted at the start-up of TCP connections are drastically affected by the initial input to the filter. The start-up transient behaviour of the TCP filter makes the initial estimates produced by the TCP filter inaccurate.

20

In K. Jacobsson, H. Hjalmarsson, N. Moller and K. H. Johanson, "Round-Trip Time Estimation in Communication Networks Using Adaptive Kalman Filtering," KTH - Royal Institute of Technology., 2004, an adaptive Kalman filter is introduced with a Cumulative Sum (CUSUM) change detection scheme to detect sudden changes in the average RTT. Although the filter corrects the start-up problem of the TCP filter it introduces other challenges, such as the online determination of drift and the possibility of misdetection when a change on the average RTT occurs.

30

OBJECT OF THE INVENTION

It is an object of this invention to provide a system and a method for estimating round-trip time in telecommunication networks which is relatively simple and which at least partially alleviates some of the aforementioned disadvantages.

In this specification, the term "fading memory polynomial (FMP) filter" means any filter which takes as its input a series of samples with a constant sampling interval, and has an output that is a function of the most recent input sample weighted by a constant, and the samples preceding it weighted by successively smaller values. The term "expanding memory polynomial (EMP) filter" means any filter which takes as its input a series of samples with a constant sampling interval, and has an output that is based on fitting a polynomial of predetermined degree by least-square error to a certain number of the input samples. The term "composite polynomial filter" means a filter that includes both FMP and EMP filters; and the term "round trip time (RTT)" is the time interval that elapses between the sending of a data packet and the receipt of its acknowledgement

20

SUMMARY OF THE INVENTION

In accordance with the invention there is provided a method for estimating round-trip time (RTT) values for data packets travelling in a telecommunication network in which a sequence of RTT samples are obtained at periodic intervals, the method comprising: during a first time interval which begins with the start of a data transmission, generating the estimated RTT values by filtering the RTT samples with a self-initialising expanding memory polynomial (EMP) filter; and, after a first switch-over time at the end of the first time interval, generating the estimated RTT values by filtering the RTT samples with a fading memory polynomial (FMP) filter.

Further features of the invention provide for the FMP filter to run in the background during the first time interval and to produce an FMP estimation error which is the difference between an output of the FMP filter and the RTT samples; for the FMP estimation error to be continually compared with an EMP estimation error which is the difference between an output of the EMP filter and the RTT samples; and for the first switch-over time to be triggered when the FMP estimation error becomes equal to or less than the EMP estimation error.

10

Still further features of the invention provide for the EMP and FMP filters to be implemented in a single composite polynomial filter which has an EMP and FMP filter mode; for the filter modes to be implemented with a similar filter structure but different sets of filter weights; and for the composite polynomial filter to switch from the EMP mode to the FMP mode by changing the filter weights.

15

In accordance with a first aspect of the invention, after the switch-over time the RTT samples are filtered with the FMP filter for the remainder of the data transmission.

20

In accordance with a second aspect of the invention, the method includes continually comparing the FMP estimation error and the EMP estimation error during the remainder of the data transmission that follows the switch-over time, and adaptively switching from one filter to the other depending on which filter generates the lowest error.

25

In TCP applications an EMP and FMP filter of degree zero may be used. In other applications, for example where a network controller desires the estimated RTT values to more closely follow the RTT samples, EMP and FMP filters up to degree four may be used.

30

The invention extends to a system for estimating round-trip time (RTT) values for data packets travelling in a telecommunication network in which a sequence of RTT samples are obtained at periodic intervals, the system comprising a composite polynomial filter for generating the estimated RTT values by filtering the RTT samples, the composite polynomial filter having an expanding memory polynomial (EMP) filter mode and a fading memory polynomial (FMP) filter mode which can be alternately selected, wherein the EMP mode is selected during a first time interval which begins with the start of a data transmission and ends at a first switch-over time, and the FMP mode is selected at the first switch-over time.

Further features of the invention provide for the filter modes to be implemented by the same filter structure but with different sets of filter weights, so that the composite polynomial filter is operable to switch from the EMP filter mode to the FMP filter mode by changing the filter weights.

Still further features of the invention provide for the FMP filter mode to run in the background during the first time interval and be used to produce an FMP estimation error which is the difference between an output of the FMP filter mode and the RTT samples; for the FMP estimation error to be continually compared with an EMP estimation error which is the difference between an output of the EMP filter mode and the RTT samples; and for the first switch-over time to be triggered when the FMP estimation error becomes equal to or less than the EMP estimation error.

In one embodiment, after the first switch-over time, the estimated RTT values are generated by filtering the RTT samples with the FMP filter for the remainder of the data transmission. In a different embodiment, after the first switch-over time, the FMP estimation error and EMP estimation error are continually compared for the remainder of the data transmission, and the

system is configured to adaptively switch between the FMP filter mode and EMP filter mode depending on which filter mode generates the lowest estimation error.

- 5 The system of the invention may have a composite polynomial filter of order zero and be used for RTT estimation in Transmission Control Protocol (TCP) applications.

BRIEF DESCRIPTION OF THE DRAWINGS

10

In the drawings:-

Figure 1 is a schematic diagram of a composite polynomial filter used in the method and system of the invention;

15

Figure 2 is a schematic diagram of an experimental setup used in evaluating the composite polynomial filter of the invention;

20

Figure 3 is a graph showing measured RTT samples and estimated RTT values produced by a prior art TCP filter and the composite polynomial filter of the invention of order-0;

25

Figure 4A is a graph showing measured RTT samples and estimated RTT values produced by a composite EMP/FMP filter of order-3; and

Figure 4B is a graph showing the error in replacing the measured RTT samples by the estimated RTT values in Figure 4A.

30

DETAILED DESCRIPTION WITH REFERENCE TO THE DRAWINGS

In packet switched telecommunication networks such as the Internet, estimated round-trip times (RTTs) for data packets sent in the network are derived from a sequence of RTT samples obtained for data packets that are sent during a particular data transmission interval. The estimated RTT values are derived by digitally filtering the RTT samples.

The digital filters currently used for generating the estimated RTT values in the Transmission Control Protocol (TCP) used in the Internet are Fading Memory Polynomial (FMP) filters of degree zero, which will also be termed TCP filters. These TCP filters work by sampling the behaviour of packets sent over a connection and averaging those sequence of RTT samples into a "smoothed" RTT estimate, \hat{x} . When a packet is sent over a TCP connection the sender measures the time it takes for the packet to be acknowledged, producing a sequence $(x_n, n \geq 0)$ of round trip time samples. Each new sample x_n is used to produce a new estimated value \hat{x} using the formula

$$\hat{x}_{n+1} = \theta \hat{x}_n + (1 - \theta)x_n \quad (1)$$

where \hat{x}_n is the current estimate of the round-trip time, \hat{x}_{n+1} is the new computed value and θ is a constant between 0 and 1 that controls how rapidly \hat{x} adapts to changes in the RTT sample values, x_n .

At the start of a TCP connection the TCP filter must begin to generate estimated RTT values from the first RTT sample onwards. FMP filters such as the TCP filter have the drawback that they are not self-initialising. Because of this, the estimated RTT value \hat{x} is drastically affected by the initial input to the filter, making the initial estimate produced at the start of TCP connection inaccurate. The result is a slow start-up time observed with the existing TCP filter.

Expanding Memory Polynomial (EMP) filters operate by fitting a polynomial to sampled data by a least-squares process. For EMP filters of order zero, for example, a straight line is fitted to the sampled data. In EMP filters of higher order various other polynomials may be fitted to the sampled data. In EMP filters, all of the incorporated samples have an equal role in the determination of the estimate which makes EMP filters self-initialising.

Although EMP filters are self-initialising they are not suited to following arbitrary processes such as a sequence of RTT samples. If an m^{th} -degree EMP filter is smoothing data that comes from a process that is a polynomial of degree m or less, then the filter model will be perfectly matched to the process. However if the process is not a polynomial of degree m or less, then the EMP will initially track the process with acceptable error but later will develop unacceptably large errors. Because RTT samples follow an arbitrary process rather than polynomial curves, EMP filters develop unacceptably large errors when tracking the RTT samples.

With FMP filters, on the other hand, the least square procedure uses a fading memory which means that the most recent sample is weighted by one and the samples preceding it by successively smaller, exponentially decaying values. This makes FMP filters good at tracking any arbitrary process, but as previously mentioned they have the drawback of not being self-initialising.

In this invention, a composite polynomial filter is provided which combines the features of the FMP and EMP filters. The composite polynomial filter provides better initial estimates than the fading memory polynomial TCP filters, and also provides good tracking of the dynamic and non-linear behaviour of the RTT samples, leading to more accurate RTT estimation throughout the data transmission interval and consequent improvements in network congestion control.

The composite polynomial filter of the invention is a combination of an EMP and FMP filter both of the same degree. During a first time interval which begins at the start of a data transmission interval, an EMP filter mode is used to produce the estimated RTT values. In the background, an FMP filter mode is also monitored and the virtual estimation error produced by the FMP filter during the first time interval is compared with the actual estimation error produced by the EMP filter. During the first time interval, the error of the FMP filter will gradually decrease as more samples become available and the transient effects of initialisation diminish. As soon as the FMP filter shows the same error in its estimation as the EMP filter, the system switches over to use the FMP filter for further estimation of the RTT values.

It is possible to model the EMP and FMP filter modes as a mathematical model with the same structure differing only in their filter weights, so that the composite polynomial filter can switch between the EMP and FMP filters by simply changing the weights. Figure 1 illustrates this in schematic form, in which the composite filter (10) filters a series of RTT samples (12) to produce estimated RTT values (14) as its output. At a data transmission start time, the composite filter (10) is set to use EMP weights (16) so that its properties are those of a self-initialising EMP filter. At some point the composite filter is then switched to an FMP mode in which it uses FMP weights (18), and its properties then become those of the FMP filter. In this way the composite filter gets the best performance from both filters, namely the self-initialising property of an EMP filter and the ability of an FMP filter to track any arbitrary process.

N_s denotes the sample number at which the switchover from the EMP to FMP weights takes place. If N_s is correctly selected then the move from EMP to FMP takes place seamlessly and without transients. To determine N_s , it is assumed that the EMP filter has been running from $n = 0$, and also that a

transient-free FMP filter has been running in the background. N_s is selected such that the outputs of the foreground EMP and the background FMP contain similar errors.

- 5 It is possible to determine the value of N_s based on the value of θ . Table 1 below gives a summary of N_s for polynomial filters up to degree 4.

| Filter Degree | 1 | 2 | 3 | 4 |
|---------------|---------------|-----------------|------------------|------------------|
| N_s | $2(1-\theta)$ | $3.2(1-\theta)$ | $5.51(1-\theta)$ | $6.63(1-\theta)$ |

Table 1: Polynomial Degree vs Switching Number N_s

10

The methodology for deriving Table 1 is described in N. Morrison, *Filter Engineering for Radar and Tracking: Gauss-Newton, Swerling, Kalman and Polynomial*, <http://www.rrsq.uct.ac.za/tracksmooth.tgz> and will not be explained in detail herein.

15

RTT and Filter Models

As mentioned, it is possible to model the EMP and FMP filters as a mathematical model with the same structure, differing only in their filter weights.

20

The round-trip time $x(t)$ is measured from a Packet Acknowledgement (ACK's) or Real-Time Control Protocol (RTCP) receiver report, resulting in a sequence y_n , where y_n is the sequence of raw RTT samples that include delays caused by transient network effects such as sporadic cross-traffic.

25

The short-lived duration of these transient effects means that their contribution to the RTT can be considered as noise from the point of view of congestion control and it is necessary to filter them out. The model of the sequence of RTT samples is then

$$y_n = x_n + v_n \quad (2)$$

in which v_n is the measurement noise. The model y_n is then input to a 1-step
5 polynomial filter, resulting in the following transposed vector

$$\hat{Z}_{n+1,n} = (\hat{x}, \tau \dot{\hat{x}}, \tau^2/2! \ddot{\hat{x}}, \dots, \tau^m/m! D^m \hat{x})_{n+1,n}^T \quad (3)$$

where the first subscript in $\hat{Z}_{n+1,n}$ signifies that the validity instant of this state
10 vector is t_{n+1} and the second subscript signifies that \hat{Z} is an output that was
based on observations that were made up to and including time t_n . In
equation 3 and the equations that follow, n indicates the most recent time
(now), t is the time, t_n the most recent time instant, m the degree of the
fitting polynomial, and τ the constant sampling interval.

15

The vector \hat{Z} having this structure is called the normalized vector and must
be de-normalized in order to obtain the desired estimate $\hat{x}(t)$ of the process
 $x(t)$. If $\hat{X}_{n+1,n} = (\hat{x}_{n+1,n}, \dot{\hat{x}}_{n+1,n}, \ddot{\hat{x}}_{n+1,n}, \dots, D^m \hat{x}_{n+1,n})$ is the state vector of the
estimate, then the general de-normalization equation for the 1-step predictor
20 is given by

$$\hat{X}_{n+1,n} = D(\zeta) \hat{Z}_{n+1,n} \quad (4)$$

in which $D(\zeta)$ is a diagonal matrix whose elements are given by
25

$$D(\zeta)_{i,i} = i! / \zeta^i \quad (5)$$

The basic structure of the 1-step-predictor EMP is illustrated using a third
degree expanding memory polynomial

$$e_n = y_n - \hat{z}_{0,n,n-1} \quad (6)$$

$$\hat{z}_{3,n+1,n} = \hat{z}_{3,n,n-1} + \delta e_n \quad (7)$$

$$\hat{z}_{2,n+1,n} = \hat{z}_{2,n,n-1} + 3\hat{z}_{3,n+1,n} + \gamma e_n \quad (8)$$

$$5 \quad \hat{z}_{1,n+1,n} = \hat{z}_{1,n,n-1} + 2\hat{z}_{2,n+1,n} - 3\hat{z}_{3,n+1,n} + \beta e_n \quad (9)$$

$$\hat{z}_{0,n+1,n} = \hat{z}_{0,n,n-1} + \hat{z}_{1,n+1,n} - \hat{z}_{2,n+1,n} + \hat{z}_{3,n+1,n} + \alpha e_n \quad (10)$$

where

$$10 \quad \delta = \frac{140}{(n+4)(n+3)(n+2)(n+1)} \quad (11)$$

$$\gamma = \frac{120(2n+1)}{(n+4)(n+3)(n+2)(n+1)} \quad (12)$$

$$\beta = \frac{20(6n^2 + 6n + 5)}{(n+4)(n+3)(n+2)(n+1)} \quad (13)$$

$$\alpha = \frac{8(2n^3 + 3n^2 + 7n + 3)}{(n+4)(n+3)(n+2)(n+1)} \quad (14)$$

15 For maximum speed and to save on the required RAM, the algorithm is implemented in computer code as follows:

$$e = y - \hat{z}_0 \quad (15)$$

$$\hat{z}_3 = \hat{z}_3 + \delta e \quad (16)$$

$$20 \quad \hat{z}_2 = \hat{z}_2 + 3\hat{z}_3 + \gamma e \quad (17)$$

$$\hat{z}_1 = \hat{z}_1 + 2\hat{z}_2 - 3\hat{z}_3 + \beta e \quad (18)$$

$$\hat{z}_0 = \hat{z}_0 + \hat{z}_1 - \hat{z}_2 + \hat{z}_3 + \alpha e \quad (19)$$

The 1-step-predictor FMP mode has precisely the same structure and there is no need to repeat its equations. The only difference between the EMP and FMP algorithms is the expressions that are used for the filter weights $\delta, \gamma, \beta, \alpha$; which for an FMP of degree 3 are set as follows:

5

$$\delta = \frac{1}{6}(1-\theta) \quad (20)$$

$$\gamma = (1-\theta)^3(1+\theta) \quad (21)$$

$$\beta = \frac{1}{6}(1-\theta)^2(11+14\theta+11\theta^2) \quad (22)$$

$$\alpha = 1-\theta^4 \quad (23)$$

10

Thus, it is possible to switch seamlessly from EMP to FMP mode by simply changing the weights, to create the composite EMP/FMP filter. The variable n does not appear in the FMP filter weights α, β, γ as it does in the EMP filter weights. Instead the FMP filter contains what is called the fading memory parameter θ . This is a real number selected by the user that lies in the range $0 < \theta < 1$ and controls the properties of the filter by controlling the exponential rate of decay of its memory.

15

Experimental Results

20

Figure 2 is a schematic diagram of an experimental setup (100) used in evaluating the composite filter of the invention. A LAN consisting of a multimedia server (102) and a multimedia client (104) was set up. Video content was stored at the server and transferred to the multimedia client through a router (106). The router also routed other kinds of traffic during busy times.

25

A simple experiment was conducted using the open source Live555 multimedia streaming library. To predict round-trip time, Live555 was

extended to support the following three filters: a standard TCP (FMP) filter, a composite filter of the invention of degree-0 and a composite filter of the invention of degree-3. The value of θ was set equal to 0.9 for the TCP filter and the FMP part of the composite filter of degree-0, while $\theta = 1/(n+1)$ for the
5 EMP-component of the composite filter of degree-0. The composite filter of degree-3 was implemented as described previously with $\theta = 0.9$.

To facilitate extraction of the receiver report contents, Live555 was set to the debugging mode. Each time that a receiver report was to be generated, the
10 receiver first computed the RTT and other network metrics as described in RFC3550. The computed value of RTT was then input into the three filters, to predict the next possible value of the RTT. The objective of the TCP and the composite filter of degree-0 was to predict the future value of the average RTT, whereas the purpose of the composite filter of degree-3 was to predict
15 the actual value of the RTT.

Figure 3 is a graph showing the measured RTT samples (200) and the estimated RTT values produced by the TCP filter (202) and the composite filter of order-0 (204). The graph is shown over a data transmission interval of
20 300 seconds with a transmission start time at $t = 0s$ and a start-up time interval (206) between $t = 0s$ and $t = 35s$.

As seen from the broken line, the TCP filter (202) takes about 35s to predict an acceptable value of the average RTT. The estimates during the start-up
25 time interval are completely wrong; after $t = 35s$ the TCP filter (202) shows an improvement on its predicted value. The composite EMP/FMP filter (204) of the invention, by contrast, shows a much better estimation during the start-up time interval (206). At $t = 35s$, the switch-over time at which the error of the FMP filter becomes equal to the error of the EMP filter, the composite
30 EMP/FMP filter switches seamlessly from the EMP mode to the FMP mode. Its estimates are therefore identical to the TCP filter after $t = 35s$.

The following table compares some of the initial values of the composite filter and the TCP filter.

| k | x_k | $x_{k,TCP}^*$ | $x_{k,EMP/FMP}^*$ | $error_{TCP}$ | $error_{EMP/FMP}$ |
|-----|---------|---------------|-------------------|---------------|-------------------|
| 1 | 0.00000 | 0.0000000 | 0.0000000 | 0 | 0 |
| 2 | 0.00052 | 0.0000519 | 0.0002594 | 0.0004671 | 0.000260 |
| 3 | 0.00037 | 0.0000833 | 0.0002950 | 0.0002827 | 7.1e-05 |
| 4 | 0.00076 | 0.0001513 | 0.0004120 | 0.0006117 | 0.0003510 |
| 5 | 0.00067 | 0.0002033 | 0.0004639 | 0.0004677 | 0.0002071 |
| 6 | 0.00078 | 0.0002608 | 0.0005163 | 0.0005172 | 0.0002617 |
| 7 | 0.00073 | 0.0003079 | 0.0005471 | 0.0004241 | 0.0001849 |
| 8 | 0.00073 | 0.0003504 | 0.0005703 | 0.0003816 | 0.000162 |
| 9 | 0.00061 | 0.0003764 | 0.0005748 | 0.0002336 | 3.525e-05 |

5

Table 2: Initial Values of the Composite EMP/FMP and TCP Filters when Estimating RTT

As can be seen in Table 2, the composite filter ($x_{k,EMP/FMP}^*$) provides much lower error rates during the initial samples as compared to the TCP filter.

In one embodiment of the invention, after the switch-over time the composite EMP/FMP filter remains in the FMP mode for the remainder of the data transmission interval and the EMP mode is only activated at the start of a new data transmission interval.

However, in an alternative embodiment of the invention, the composite filter continually compares the error of the FMP filter and the error of the EMP filter during the remainder of the data transmission interval and adaptively switches from one filter to the other depending on which generates the lowest error. The composite filter is thus able to switch not only from EMP to FMP but also back from FMP to EMP during the transmission interval. This embodiment is also illustrated in Figure 3. As can be seen in Figure 3, at about $t = 200s$ a large spike (208) in the RTT samples is observed.

Immediately after the spike, the error produced by the EMP mode of the composite filter (which was running in the background prior to $t = 200s$) becomes smaller than the error produced by the FMP mode. The composite filter (204) then switches back to the EMP mode between $t = 200s$ and $t = 250s$ after which the error produced by the passive FMP mode (now running in the background between $t = 200s$ and $t = 250s$) becomes smaller than the error of the active EMP mode and the composite filter switches back to the FMP mode again.

10 It is clear from Figure 3 that providing the ability to switch between filter modes during the data transmission interval provides the composite filter with the ability to recover quickly from sudden changes in the RTT sample values, whereas the TCP filter takes longer to recover, as demonstrated by the long tail (210) in the TCP filter output following the spike.

15

Higher order composite filters were also tested. Figure 4A is a graph showing measured RTT samples (300) and the estimated RTT values (302) produced by a composite EMP/FMP filter of order-3. In some applications a network controller may require the estimated RTT values to more closely follow the RTT samples than is the case with the composite filter of degree-0 and in such an instance a composite filter of degree-3 may be appropriate.

20 As seen in Figure 4A, the higher-order composite filter follows the RTT samples more closely between $t = 50s$ and $t = 150s$) than the composite filter of order-0 of Figure 3. However, this is only the case in the absence of spikes. When spikes appear, the filter develops large errors. These errors are shown most clearly in Figure 4B, which is a graph showing the error in replacing the measured RTT samples (300) by the estimated RTT values (302) in Figure 4A. Figure 4B clearly shows that large errors develop around the spikes at $t = 160s$ and $t = 200s$. For most applications such as TCP it is envisaged that the composite polynomial filter of order-0 will be most

25

30

appropriate, as it is able to cope well with large spikes. Both composite filters, however, display much better start-up performance than prior art TCP filters based on FMP only.

- 5 The invention therefore provides a method and a system for estimating round-trip time values in telecommunication networks which relies on a simple composite polynomial filter that provides improved start-up performance compared with existing TCP filters but does not suffer from the drawbacks associated with more complex adaptive Kalman filters. Improved
10 round-trip time estimation allows efficient development of congestion control techniques for multimedia applications, leading to noticeable improvement in network congestion control.

It will be appreciated that the invention is not limited to the specific composite
15 EMP/FMP filters disclosed, and that EMP and FMP filters of any order may be used in the composite filter of the invention. In particular, any means of switching between the EMP and FMP filters may be employed, and the filters need not be of equal order.

CLAIMS:

1. A method for estimating round-trip time (RTT) values for data packets travelling in a telecommunication network in which a sequence of RTT samples (12) are obtained at periodic intervals, the method comprising: during a first time interval (206) which begins with the start of a data transmission, generating the estimated RTT values (14) by filtering the RTT samples with a self-initialising expanding memory polynomial (EMP) filter (16); and, after a first switch-over time at the end of the first time interval, generating the estimated RTT values by filtering the RTT samples with a fading memory polynomial (FMP) filter (18).

5

10
2. The method as claimed in claim 1 in which, during the first time interval: the FMP filter is run in the background and is used to produce an FMP estimation error which is the difference between an output of the FMP filter and the RTT samples; the FMP estimation error is continually compared with an EMP estimation error which is the difference between an output of the EMP filter and the RTT samples; and the first switch-over time is triggered when the FMP estimation error becomes equal to or less than the EMP estimation error.

15

20
3. The method as claimed in claim 1 or claim 2 in which the EMP and FMP filters are implemented in a single composite polynomial filter (10) which has an EMP filter mode and an FMP filter mode.

25
4. The method as claimed in claim 3 in which the filter modes are implemented by the same filter structure but with different sets of filter weights, so that the composite polynomial filter can be switched from the EMP filter mode to the FMP filter mode by changing the filter weights.

30

5. The method as claimed in any one of claims 1 to 4 in which, after the first switch-over time, the estimated RTT values are generated by filtering the RTT samples with the FMP filter for the remainder of the data transmission.
- 5
6. The method as claimed in any one of claims 2 to 4 in which, after the first switch-over time, the method includes continually comparing the FMP estimation error and EMP estimation error for the remainder of the data transmission, and adaptively switching between the two depending on which filter generates the lowest estimation error.
- 10
7. The method as claimed in any one of claims 1 to 6 in which the EMP filter and the FMP filter are selected to have the same polynomial order.
- 15
8. The method as claimed in claim 7 in which the polynomial order is less than or equal to four.
- 20
9. The method as claimed in claim 8 in which the polynomial order is zero and the method is used for RTT estimation in Transmission Control Protocol (TCP) applications.
- 25
10. A system for estimating round-trip time (RTT) values for data packets travelling in a telecommunication network in which a sequence of RTT samples (12) are obtained at periodic intervals, the system comprising a composite polynomial filter (10) for generating the estimated RTT values by filtering the RTT samples, the composite polynomial filter having an expanding memory polynomial (EMP) filter mode (16) and a fading memory polynomial (FMP) filter mode (18) which can be alternately selected, wherein the EMP mode is selected during a first
- 30

time interval (206) which begins with the start of a data transmission and ends at a first switch-over time, and the FMP mode is selected at the first switch-over time.

- 5 11. The system as claimed in claim 10 in which the filter modes are implemented by the same filter structure but with different sets of filter weights, so that the composite polynomial filter is operable to switch from the EMP filter mode to the FMP filter mode by changing the filter weights.
- 10
12. The system as claimed in claim 10 or claim 11 in which, during the first time interval: the FMP filter mode is run in the background and is used to produce an FMP estimation error which is the difference between an output of the FMP filter mode and the RTT samples; the FMP estimation error is continually compared with an EMP estimation error which is the difference between an output of the EMP filter mode and the RTT samples; and the first switch-over time is triggered when the FMP estimation error becomes equal to or less than the EMP estimation error.
- 15
- 20
13. The system as claimed in any one of claims 10 to 12 in which, after the first switch-over time, the estimated RTT values are generated by filtering the RTT samples with the FMP filter for the remainder of the data transmission.
- 25
14. The system as claimed in claim 12 in which, after the first switch-over time, the FMP estimation error and EMP estimation error are continually compared for the remainder of the data transmission, and the system is configured to adaptively switch between the FMP filter mode and EMP filter mode depending on which filter mode generates the lowest estimation error.
- 30

15. The system as claimed in any one of claims 10 to 14 in which the order of the composite polynomial filter is zero and the method is used for RTT estimation in Transmission Control Protocol (TCP) applications.
- 5

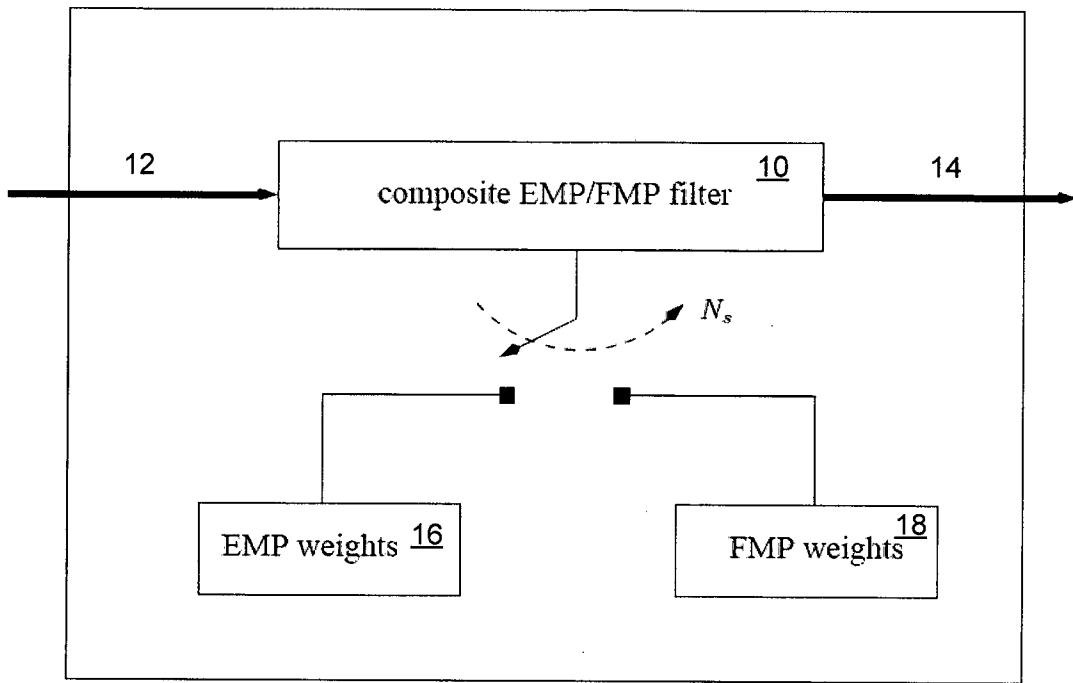


Figure 1




Figure 2




Figure 3

3/3




Figure 4A




Figure 4B

INTERNATIONAL SEARCH REPORT

International application No.

PCT/IB2011/000393

| A. CLASSIFICATION OF SUBJECT MATTER | | | | | | | | | | | | |
|---|--|--|--|---|---|---|---|--|--|---|--|--|
| Int. Cl. | | | | | | | | | | | | |
| <i>H04L 12/56</i> (2006.01) <i>H04L 12/28</i> (2006.01) | | | | | | | | | | | | |
| According to International Patent Classification (IPC) or to both national classification and IPC | | | | | | | | | | | | |
| B. FIELDS SEARCHED | | | | | | | | | | | | |
| Minimum documentation searched (classification system followed by classification symbols) | | | | | | | | | | | | |
| Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched | | | | | | | | | | | | |
| Electronic data base consulted during the international search (name of data base and, where practicable, search terms used) | | | | | | | | | | | | |
| EPODOC, WPI, INSPEC, TXTUS1, TXTUS2, TXTUS3, TXTUS4, TXTEP1, TXTGB1, TXTWO1, Google Patents, Google Scholar: with round trip time, RTT, estimate, fading memory, expanding memory, polynomial, and similar key words | | | | | | | | | | | | |
| C. DOCUMENTS CONSIDERED TO BE RELEVANT | | | | | | | | | | | | |
| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. | | | | | | | | | | |
| A | JACOBSON, Congestion Avoidance and Control, Proc. ACM SIGCOMM '88, Vol. 18, No. 4, August 1988, pages 314-329 Section 2 | | | | | | | | | | | |
| A | KARN et al., Improving Round-Trip Time Estimates in Reliable Transport Protocols, ACM Transactions on Computer Systems, Vol. 9, No. 4, November 1991, Pages 364-373 Sections 2 and 3 | | | | | | | | | | | |
| <input checked="" type="checkbox"/> Further documents are listed in the continuation of Box C <input type="checkbox"/> See patent family annex | | | | | | | | | | | | |
| <p>* Special categories of cited documents:</p> <table border="0"> <tr> <td>"A" document defining the general state of the art which is not considered to be of particular relevance</td> <td>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</td> </tr> <tr> <td>"E" earlier application or patent but published on or after the international filing date</td> <td>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone.</td> </tr> <tr> <td>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</td> <td>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art</td> </tr> <tr> <td>"O" document referring to an oral disclosure, use, exhibition or other means</td> <td>"&" document member of the same patent family</td> </tr> <tr> <td>"P" document published prior to the international filing date but later than the priority date claimed</td> <td></td> </tr> </table> | | | "A" document defining the general state of the art which is not considered to be of particular relevance | "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention | "E" earlier application or patent but published on or after the international filing date | "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone. | "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) | "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art | "O" document referring to an oral disclosure, use, exhibition or other means | "&" document member of the same patent family | "P" document published prior to the international filing date but later than the priority date claimed | |
| "A" document defining the general state of the art which is not considered to be of particular relevance | "T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention | | | | | | | | | | | |
| "E" earlier application or patent but published on or after the international filing date | "X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone. | | | | | | | | | | | |
| "L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified) | "Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art | | | | | | | | | | | |
| "O" document referring to an oral disclosure, use, exhibition or other means | "&" document member of the same patent family | | | | | | | | | | | |
| "P" document published prior to the international filing date but later than the priority date claimed | | | | | | | | | | | | |
| Date of the actual completion of the international search 01 July 2011 | | Date of mailing of the international search report 6 - JUL 2011 | | | | | | | | | | |
| Name and mailing address of the ISA/AU AUSTRALIAN PATENT OFFICE PO BOX 200, WODEN ACT 2606, AUSTRALIA E-mail address: pct@ipaaustralia.gov.au Facsimile No. +61 2 6283 7999 | | Authorized officer DR. THAYAPARAN THANABALASINGHAM AUSTRALIAN PATENT OFFICE (ISO 9001 Quality Certified Service) Telephone No : 61262832238 | | | | | | | | | | |

INTERNATIONAL SEARCH REPORT

International application No.

PCT/IB2011/000393

| C (Continuation). DOCUMENTS CONSIDERED TO BE RELEVANT | | |
|---|---|-----------------------|
| Category* | Citation of document, with indication, where appropriate, of the relevant passages | Relevant to claim No. |
| A | FILLON et al., Multi-objective Genetic Programming for Improving the Performance of TCP, Proceedings of 10th European Conference on Genetic Programming, April 2007, pages 170 180 Section 2 | |