Emergency Calls in Flow-Aware Networks

Andrzej Jajszczyk, *Fellow, IEEE,* and Robert Wójcik

***Abstract*— In this paper it is shown that Flow-Aware Networks (FAN), although providing superior transmission quality, may force us to wait for the network resources. Such a situation is inconvenient for the realization of emergency VoIP connections. In order to overcome the presented problem, differentiated blocking along with the Static Router Configuration approach is presented, and it is believed to be a simple, adequate, and above all, feasible solution.**

***Index Terms*— Flow-Aware Networking, FAN, service differentiation, QoS.**

I. Introduction

THE concept of Flow-Aware Networking (FAN) as a novel approach to assure quality of service in packet networks was introduced in 2004 [1]. The goal of FAN is to enhance the current IP network, by improving its performance under heavy congestion. To achieve that, certain traffic management mechanisms to control link sharing are introduced, namely: measurement-based admission control ([2]) and priority fair queuing ([1], [3]). The former is used to keep the flow rates sufficiently high, to provide a minimal level of performance for each flow in case of overload. The latter realizes fair sharing of link bandwidth, while ensuring negligible packet latency for flows emitting at lower rates.

In FAN, admission control and service differentiation are implicit. There is no need for a priori traffic specification, as well as there is no class of service distinction. Both streaming and elastic flows achieve a necessary quality of service without any mutual detrimental effect. Nevertheless, streaming and elastic flows are implicitly identified inside the FAN network. This classification, however, is based solely on the current flow peak rate. All flows emitting at lower rates than the current *fair rate* are referred to as streaming flows, and packets of those flows are prioritized. The remaining flows are referred to as elastic flows.

A classic FAN thinking includes a general rule to limit the number of active flows, so that the transmissions currently in progress could always obtain at least a decent quality of service level. Such an approach is considered beneficial, however, for certain usages (e.g., the Voice over IP technology) it may be hazardous. The purpose of this work is to expose and document that negative aspect of FAN, i.e., blocking the incoming connections upon congestion, and to present a viable solution to enhance the architecture. We believe that by simple means, i.e., the Static Router Configuration approach, we can improve the perceivability of the FAN networks.

Manuscript received April 19, 2007. The associate editor coordinating the review of this letter and approving it for publication was Prof. Iakovos Venieris.

The authors are with the Department of Telecommunications, AGH University of Science and Technology, Krakow, Poland (email: {jajszczyk, robert.wojcik}@kt.agh.edu.pl).

Digital Object Identifier 10.1109/LCOMM.2007.070609.

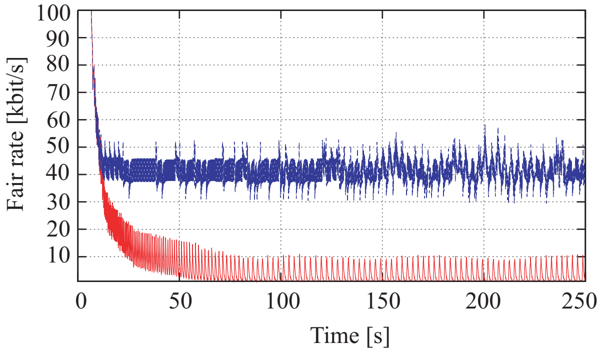

Fig. 1. Performance under congestion of a classic IP link (lower line) and a FAN link (upper line).

II. Performance under Congestion

Flow-Aware Networking is meant to impose fairness to the network. In terms of congestion, each user should be able to utilize the same, fair amount of link bandwidth. One of the most important parameters describing the current link state is the *fair rate* (FR). FR indicates the amount of resources that each flow could utilize at the moment. Therefore, it is vital to the FAN router to maintain the value of FR above a certain threshold, so that each flow in progress could achieve a reasonable quality of service level.

Under congestion, FAN performance may be considered superior to the behavior of the classic IP network. It is due to the fact, that only a limited number of flows may be simultaneously admitted on a link. Such an approach virtually guarantees that once a flow is admitted, it will perceive at least a decent quality of service. To demonstrate the difference between the behavior of classic IP and FAN links, a simple simulation was performed. The scenario in which 300 TCP-based elastic flows and 25 UDP-based VoIP flows compete for resources of a 1 Mbit/s link was identical for both cases. Figure 1 compares the results by showing the measured fair rate values over time. As can be observed, these values constantly fluctuate and the oscillations are caused by the high frequency of the measurements. If the interval between the consecutive measurements is longer, the character of the obtained values is more averaged and the oscillations fade. It should be noted that choosing the proper measurements interval is a very important issue, however, it is outside of the scope of this paper.

On the classic IP link (lower line), all flows are admitted, once they appear. Since their number is significant, the rate at which they can transmit quickly drops down below 10 kbit/s.

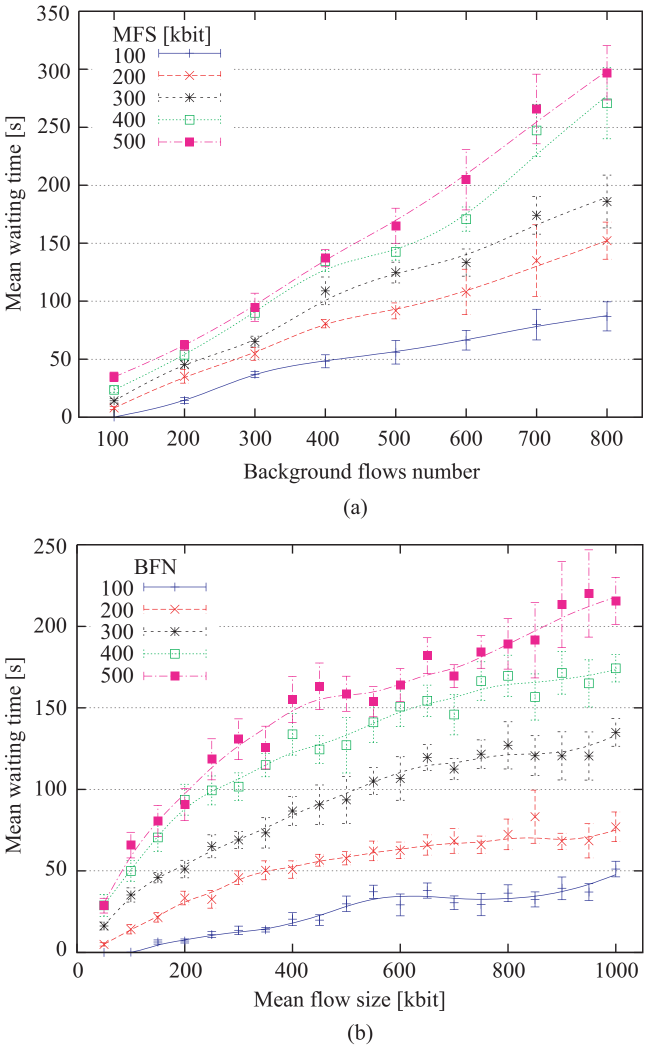

Fig. 2. Mean VoIP flow waiting time with respect to (a) background flows number (BFN) and (b) mean background flow size (MFS).

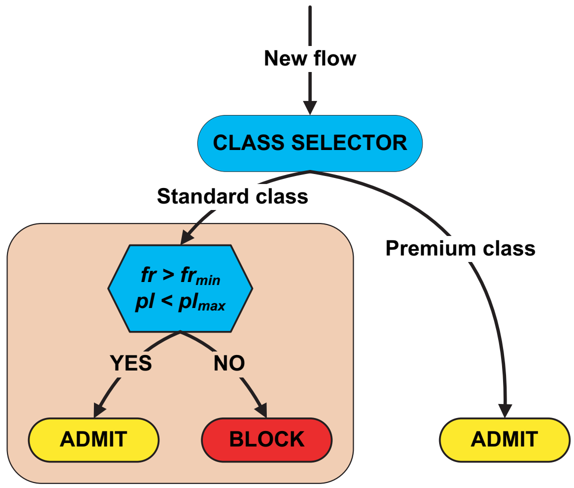

Fig. 3. Admission control routine of FAN with premium class of flows. The marked area presents the original FAN routine.

On the other hand, FAN (upper line) preserves the fair rate on a level of approximately 40-50 kbit/s. Unfortunately, in order to achieve this goal, some flows must be temporarily blocked. However, as the Internet becomes part of everyday's life, more and more customers use the Voice over IP technology, instead of the ordinary PSTN telephone service. This means that in case of emergency, the ability to contact with emergency services depends on the current congestion status in the network. For these customers, the availability to make a phone call is much more important than its quality.

Figure 2 presents the mean waiting times for VoIP flows while they compete for network resources with other TCP flows during the scenario, when the link is FAN-aware. At least 10 simulations were performed to obtain each value: from them the average as well as the 95% confidence intervals (considering a Student's t distribution) were calculated.

In Figure 2a the link congestion rises along with the number of background flows, while in Figure 2b the increased congestion is caused by varying the mean background flow size. As seen in both parts of the figure, the mean waiting time for transmission grows along with the offered load, which is expected. The greater the flow number, the lower the chance of a particular VoIP flow to be admitted, and therefore, the longer waiting time. On the other hand, when the flow number is constant, but their mean size grows, the more rarely a flow ends, hence, a new one may be admitted with lower frequency, which also increases the average waiting time.

The values presented in Figure 2 are averaged. In fact, during the simulations, certain amount of VoIP flows observed very short and completely acceptable waiting times. However, for the rest of them that period was excessively long and simply much too long for life-saving emergency connections.

It is worth mentioning that in case of classic IP networks the emergency connections are also endangered by congestion, as the low transmission rates may render voice imperceivable. FAN networks, although providing superior transmission quality, may force us to wait for the network resources. Fortunately, both these disadvantages may be overcome by introducing differentiated blocking into FAN networks.

III. Differentiated Blocking Approach

The notion of differentiated (selective) blocking, which aims at applying different blocking criteria to newly arriving flows, was envisaged in [2]. The standard FAN routine (marked area in Figure 3) causes the admission control block to make the decision based on currently measured values of the fair rate and priority load. Therefore, a new flow is admitted if the current fair rate (fr) is greater than the minimal fair rate (fr_{min}), and the current priority load (pl) does not exceed the maximum priority load (pl_{max}) threshold. Otherwise, the incoming flow is blocked.

In the simplest example, the differentiated blocking scenario includes two classes of service, namely: the standard class and the premium class. The admission control procedure in such a situation is presented in Figure 3. The role of the class selector is to recognize which blocking criteria should be applied to the incoming flow. Flows belonging to the standard class are subject to admission control under the rules of the original classless FAN. The premium class flows are always admitted. It is also possible to introduce additional classes of service, however, for the purpose of realizing the emergency calls, the premium class is sufficient.

Differentiated blocking operates only when congestion occurs, as in the other cases, there is obviously no need for blocking the arriving flows. Additionally, this mechanism does not interfere with protected flows. Any flow that is already placed in the protected flow list is always forwarded. Furthermore, differentiated blocking does not have any influence on flows that are currently in progress. In other words, all flows receive the same treatment from the scheduling algorithm once they are admitted.

The procedure presented in Figure 3 is well suited for emergency VoIP connections. All flows related to the VoIP emergency call would belong to the premium class, i.e., they would never be blocked by admission control in a FAN router. Consequently, since these flows are not subject to admission control, they would not observe any transmission waiting time.

This scheme, however, introduces a certain drawback. As we interfere with the admission control mechanism, we may observe the performance degradation, because prioritized flows are admitted on the link, even under the circumstances in which they normally would not be. Fortunately, this behavior is believed to be insignificant to the overall link performance for two reasons. Firstly, the required bit rate of a single internet telephony connection is relatively low, especially compared to the core link capacities, and therefore, admitting even a few additional flows should not degrade the quality of the remaining transmissions significantly. Secondly, the fair rate degradation is a temporal process. It is temporal due to the fact that while active flows terminate naturally, new ones are not admitted until the fair rate returns to its desired value.

Although introducing differentiation mechanisms to FAN routers is very simple, the signalling issue remains. As the experience of IntServ and DiffServ has shown, every method of introducing to the network the knowledge about the treatment of particular flows, is inevitably associated with a major increase of complexity or severe scalability reduction. Therefore, each explicit service differentiation mechanism should not rely on any signalling or packet marking procedure, as the IP's and FAN's original simplicity and scalability must be preserved. That is the reason why, in the next section, a Static Router Configuration approach is proposed.

IV. Static Router Configuration

The proposed mechanism of explicit service differentiation is easy to implement, does not require any new functionalities and hardly complicate the existing ones. However, the signalling remains a great issue. It is very difficult to inform the nodes which flows should be discriminated, without reducing the scalability of the architecture. Implicit service differentiation works well in FAN because it does not rely on any network signalling. Flows are prioritized or discriminated based on their performance which is internally measured by proper cross-protect mechanisms. However, to implement differentiated blocking, routers must be somehow informed which flows should be treated differently.

The IntServ and DiffServ experiences have shown that introducing explicit service differentiation is difficult, due to the signalling problems and the required inter-domain agreements. Therefore, it seems that it is impossible to introduce differentiated blocking into FAN networks globally. However, for a limited scope, the explicit service differentiation procedures may be used in FAN, along with the Static Router Configuration approach.

Static Router Configuration (SRC) is a strategy of manually defining classes of flows and their treatment by network administrators. This approach, obviously, cannot be used globally, yet it is the easiest way to provide explicit service differentiation without any network complication or modification. SRC seems to be an adequate and simplest solution for introducing differentiated blocking to FAN networks.

Because emergency calling is a local matter (always to the nearest emergency center), the SRC approach may be used. An emergency center is responsible for a certain geographical region. For the differentiated blocking scheme to be used, all nodes in the region must recognize and prioritize flows with the source or destination IP address equal to the address of the proper emergency center. Provided that the emergency center's IP address is static (does not change over time), all routers in the region must be configured only once.

The SRC strategy is the only solution that does not interfere with FAN's superior scalability. Obviously, this approach is not sufficient for many services, however, it is perfectly suited for VoIP emergency connections. Moreover, with SRC, the differentiated blocking scheme may be used for any other local scope service.

V. Conclusion

Admission control and scheduling blocks of a FAN's cross-protect router are the key components responsible for improving network performance in case of overload. In order for the active flows to perceive a good enough quality of service, only a certain amount of flows may be simultaneously admitted on a link. Unfortunately, this mechanism may be dangerous for the Internet telephony, especially for the emergency connections.

To overcome the described negative behavior, we proposed introducing the differentiated blocking scheme, and make all flows related to realizing emergency connections unblockable by admission control blocks. In order to achieve this goal, the Static Router Configuration, as a way to inform all the nodes which flows should be prioritized, is also proposed. Considering significant benefits, along with a reasonably low cost associated with the proposition, we believe that introducing differentiated blocking along with the Static Router Configuration approach will greatly improve the end-user perception of the FAN architecture.

References

[1] A. Kortebi, S. Oueslati, and J. Roberts, "Cross-protect: implicit service differentiation and admission control," in *Proc. IEEE HPSR 2004*.

[2] S. Oueslati and J. Roberts, "A new direction for quality of service: Flow-aware networking," in *Proc. NGI 2005*.

[3] A. Kortebi, S. Oueslati, and J. Roberts, "Implicit Service Differentiation using Deficit Round Robin," in *Proc. ITC19*, Aug./Sept. 2005.