stantially simplified the formulation of the delay bound computation (inequality (1)), and to Eberhard Mueller-Menrad, who discovered several bugs in an earlier version of Tables 1 and 2. He is also thankful to Ellen Boyle for her invaluable help with the preparation of the manuscript.

References

[BaMa91] A. Banerjea and B. Mah. "The Real-Time Channel Administration Protocol". *Proc. Second International Workshop on Network and Operating System Support for Digital Audio and Video*, Springer-Verlag, Heidelberg, Germany, 160–170, November 1991.

[FeBZ92] D. Ferrari, A. Banerjea, and H. Zhang. "Network Support for Multimedia - A Discussion of the Tenet Approach". Technical Report TR-92-072, International Computer Science Institute, Berkeley, October 1992.

[Ferr92] D. Ferrari. "Real-Time Communication in an Internetwork". *Journal of High Speed Networks*, 1(1):79–103, 1992.

[Ferr93] D. Ferrari. "Distributed Delay Jitter Control in Packet-Switching Internetworks". *Journal of Internetworking: Research and Experience*, 4(1):1–20, 1993.

[FeVe90] D. Ferrari and D. Verma. "A Scheme for Real-Time Channel Establishment in Wide-Area Networks". *IEEE Journal Selected Areas of Communication*, 8(4):368–379, April 1990.

[FKKM92] A. G. Fraser, C. R. Kalmanek, A. E. Kaplan, W. T. Marshall, and R. C. Restrick. "Xunet 2: A Nationwide Testbed in High-Speed Networking". *Proc. INFOCOM'92*, Florence, Italy, 1992.

[JSZC92] S. Jamin, S. Shenker, L. Zhang, and D. D. Clark. "An Admission Control Algorithm for Predictive Real-Time Service". *Proc. Third International Workshop on Network and Operating System Support for Digital Audio and Video*, Springer-Verlag, Heidelberg, Germany, 349–356, 1992.

[LaPa91] A. A. Lazar and G. Pacifici. "Control of Resources in Broadband Networks with Quality of Service Guarantees". *IEEE Communications Magazine*, 29(10):66–73, October 1991.

[ZhKe91] H. Zhang and S. Keshav. "Comparison of Rate-Based Service Disciplines". *Proc. of ACM SIGCOMM '91*, Zurich, Switzerland, 113–122, September 1991.

of course, makes the tasks of the destination heavier (and reduces those of the "intermediate destinations" to zero), as well as calling for much longer establishment messages, but simplifies the channel establishment code very substantially. It should also be easier to design and implement this code so that it is truly portable to a wide variety of internetworking environments.

All of these considerations should obviously (and will) be confirmed by experience with applying the new admission control method in several internetworking testbeds.

5 Conclusion

A more general and intellectually satisfying approach than the one previously proposed by the Tenet Group to admission control for real-time communication services in packet-switching internetworks has been presented. The new approach is based on a real-time channel model consisting of a cascade of single queueing servers, each one of which may be of one of two types (J or NJ) and contains given amounts of three types of resources; the connections between servers are characterized only by their (constant) delay. The admission tests and computations are designed and have to be executed for each server along the path of the channel to be established. The only aspect of a server that influences (some of) the tests and computations is the scheduling discipline implemented by the server. Servers with different scheduling disciplines can coexist along a channel's route, but the appropriate tests and computations will have to be run in each server during the establishment procedure. J-type and NJ-type servers can also coexist, of course. To show an example of tests and computations, we have presented those for the case of the D-EDD and J-EDD disciplines.

The advantages and the drawbacks of the new method have been briefly discussed. The method's generality is its most important characteristic: it is at the root of such advantages as the portability of the tests and computations as well as of the code, and the simpler one-level approach to admission control in internetworks; it is also responsible for such relative disadvantages as the higher establishment overhead and the longer establishment messages.

It seems clear, on the basis of these considerations, that the new method makes a positive and useful contribution to the technology of real-time communication. Only careful testbed implementation and experimentation, however, can confirm the validity of this conclusion.

Acknowledgements

The author is grateful to the members of the Tenet Group for being intelligent and creative sounding boards in several discussions of these topics. He is particularly indebted to Amit Gupta, who provided many suggestions and sub-

cases. The new method replaces it with a single server to be used as a building block for much more complicated and realistic node models. Since the new tests and computations pertain to this building block and not to the node model, the node model can be changed without having to modify the tests or the computations to be performed for admission control. Some of the tests and computations depend on the scheduling discipline implemented by the server; the correct ones will obviously have to be used during the channel establishment procedure in each server, the only condition (fortunately an easy one to satisfy) being that different tests and computations in different servers will interoperate, so that the destination will receive all the information it needs to make correct admission and allocation decisions.

It should be noted that most of the tests and computations devised for Schemes 0 and 1 are valid for each single server, as they were developed for nodes that were represented by single servers. This observation has allowed us to benefit greatly from our previous work; for instance, we did not have to prove inequalities (1), since the proof had already been given in [FeVe90]. The only difference, but a non-trivial one, is that tests and computations are now to be applied to each server, not to each node. Since it is to be expected that there will be one processor in each node running tests and computations whenever a channel establishment request is received, that processor will have to run as many sets of tests and computations as there are single servers in the model chosen for that node. One could argue that this will increase the overhead and the time of channel establishment; the counterargument, however, is that the new method is reliable, while the old one was rather shaky, either because it reduced the entire, possibly quite complicated, node to a single server, or because it relied on the assumption that the node had a severe bottleneck that could be correctly identified and would never change.

Another important consequence of the new approach is found in the implementation of admission control within an internetwork. With the old method, heterogeneous internetworks were treated in a hierarchical way: each subnetwork had its own characterization, and was dealt with, to a large extent, independently of the rest of the internetwork. The channel's destination allocated resources (i.e., assigned local performance bounds) only to highest-level subnetworks and links; within each subnetwork, such assignment was done by the "local destination" (the endpoint of the channel's path within the subnetwork), and so on recursively [Ferr92, BaMa91]. All of this can be done without violating the single roundtrip requirement of the Tenet channel establishment procedure. However, while the hierarchical procedure is elegant and efficient, its implementation, debugging, and testing are quite complicated. The new method does not require resorting to such a procedure, as long as each internetwork component along a channel's path can be modeled as a cascade of single servers as in Figure 2. If all heterogeneity can be absorbed by the modeling generality of the new method, a condition that is likely to be satisfied in practical cases, we can deal with one-level paths even when the internetwork is complex. This,

Table 2: Destination Tests for Admission Control

<table>
<tr><th>Name</th><th>Computation</th><th>Test</th><th>Forward</th></tr>
<tr><td>$D\text{-}Test$</td><td>$D_{prop,i} \leftarrow \sum_{n=1}^{N} D_{prop,n}$;
$D_{min,i} \leftarrow \sum_{n=1}^{N} D_{min,n}$</td><td>$D_{min,i} \le D_{max,i} \le D_{prop,i}$</td><td>–</td></tr>
<tr><td>$W\text{-}Comp$</td><td>if $W_{min,i} < 1,\ w_{n,i} \leftarrow W_{min,i}{}^{1/N}\ (n = 1, ..N)$
if $W_{min,i} = 1,\ w_{n,i} \leftarrow 1\ (n = 1, ..N)$</td><td>–</td><td>$w_{n,i}(n = 1, ..N)$</td></tr>
<tr><td>$Z\text{-}Test$</td><td>if $Z_{min,i} < 1,\ Z_{max,i} \leftarrow \prod_{n=1}^{N} zmax_{n,i}$,
$z_{n,i} = zmax_{n,i} \left(\frac{Z_{min,i}}{\prod_{n=1}^{N} zmax_{n,i}} \right)^{1/N} (n = 1, ..N)$</td><td>$Z_{min,i} \le Z_{max,i}$</td><td>$z_{n,i}(n = 1, ..N)$</td></tr>
<tr><td>$d\text{-}Comp$</td><td>if $\sum_{n=1}^{N} dmax_{n,i} \le D_{max,i}$, then
$D'_{max,i} \leftarrow \sum_{n=1}^{N} dmax_{n,i}$,
$d_{n,i} \leftarrow dmax_{n,i}(n = 1, ..N)$
if $\sum_{n=1}^{N} dmax_{n,i} > D_{max,i}$, then
$D'_{max,i} \leftarrow D_{max,i}$
$d_{n,i} \leftarrow dmin_{n,i} + \frac{1}{N}(D_{max,i} - D_{min,i})$
$(n = 1, ..N)$</td><td>–</td><td>$d_{n,i}(n = 1, ..N)$
$D'_{max,i}$</td></tr>
<tr><td>$J\text{-}Test$</td><td>if $J_{max,i}$ is defined, then m = position of the last J-type server
and $J_{N,i} \leftarrow \sum_{k=m}^{N} (d_{k,i} - t_{k,i})$;
$D_{prop,n,i} \leftarrow \sum_{k=1}^{n} d_{prop,k}\ (n = 1, ..N)$
$D^*_{n,i} \leftarrow D_{prop,n,i} + \sum_{k=1}^{n-1} d_{k,i}\ (n = 1, ..N)$</td><td>$J_{N,i} \le J_{max,i}$</td><td>$D^*_{n,i}(n = 1, ..N)$</td></tr>
<tr><td>$b\text{-}Comp$</td><td>$\Delta_{n,i} \leftarrow \sum_{k=m}^{n} (d_{k,i} - t_{k,i})$
if $J_{max,i}$ is undefined, then $m = 1$
if $J_{max,i}$ is defined, then m = position of the previous J-type server
$b_{n,i} \leftarrow s_{max,i} \left\lceil \frac{d_{n,i} + \Delta_{n,i}}{x_{min,i}} w_{n,i} \right\rceil\ (n = 1, ..N)$</td><td>–</td><td>$b_{n,i}(n = 1, ..N)$</td></tr>
</table>