Survivable Wireless ATM Network Architecture

Y. H. Wang, W. S. Soh, M. Y. Tsai, and H. S. Kim
Department of Electrical and Computer Engineering
Carnegie Mellon University
Pittsburgh, PA 15213-3890

***Abstract*-Amidst the rapid growth of wireless broadband networks, little attention has been paid to wireless network survivability issues. We propose a survivable wireless ATM network architecture that allows it to survive from a single base station failure condition, by redirecting a failure cell's traffic via its six neighboring cells. We present two failure-handling schemes for the proposed architecture. The first scheme is a bandwidth reservation (BR) scheme that is targeted at achieving high survivability by reserving bandwidths at appropriate locations in the network. The second scheme is a best-effort (BE) scheme that does not perform any reservation, and is targeted at achieving high bandwidth utilization. Simulation results show that the BR scheme achieves good survivability as expected. On the other hand, the BE scheme provides better utilization while having slightly lower average survivability. The decrease in average survivability for the BE scheme is not tremendous, largely due to its flexibility in the use of spare bandwidths from neighboring cells when a failure occurs. However, the BE scheme requires more frequent update messages between the switches in order to update each other about the amount of bandwidth that they could provide for failure-handling. We have also considered important issues such as switchover time, and data integrity, for our proposed schemes.**

I. Introduction

Due to a rapid growth of wireless broadband networks and the growing impact of a network failure, real-time failure recovery becomes a critical function for maintaining the robustness in future high-speed wireless networks. In order to provide the user with a transparent connectivity regardless of the state of the network, the survivable network should work under any failure pattern and be robust to a topological change, especially in wireless networks. The success of the restoration process, however, depends on failure location, traffic distribution, mobile behavior, and available spare capacity over the network when a failure occurs. Thus, it is necessary to dynamically distribute the traffic and spare capacity for better survivability. At the same time, the traffic flow should be assigned optimally in the network for efficient utilization. The mobility of mobile terminals (MTs) with different Quality-of-Service (QoS) requirements further complicates the design of survivable network architecture. The network needs to respond to preserve QoS upon a change in the network traffic pattern including the network failure.

While there has been much effort in survivable network design for wired networks [1–4], little attention has been paid to survivability issues in wireless networks. In this paper, we propose a survivable Wireless Asynchronous Transfer Mode (WATM) network architecture, and describe two failure-handling schemes that enable the wireless network to recover dynamically from a single link or base station (BS) failure.

The remaining of this paper is organized as follows. Section II describes the proposed survivable network architecture and the two failure-handling schemes in detail. In Section III, we discuss some of the important issues that need to be considered for the traffic redirection during failure. Section IV describes the simulation model and the assumptions made. In Section V, simulation results are presented to demonstrate the performance of the proposed schemes. Finally, we give our concluding remarks in Section VI.

II. Survivable Network Architecture and Failure-handling Schemes

The objective of our proposed schemes is to achieve high survivability for the existing connections in the failure cell, while achieving reasonably high bandwidth utilization. In the following context, we shall refer to a cell as a "failure cell" when either its BS or the immediate link connected to its BS fails. In our proposed failure-handling schemes, the key idea is to redirect all the failure cell's traffic via its six neighboring cells. Note that our schemes only cater for a single failure at any time, assuming that the time between successive failures is long enough for the problem to be fixed.

In order to provision for the traffic redirection when a cell failure occurs, we have the option of choosing either wired or wireless means. In our proposed scheme, we have chosen the latter, as it might be too costly, if not impossible, to interconnect all the neighboring BSs in the entire wireless network by wired means. By placing standby repeaters at appropriate positions in each cell, together with the reservation of an emergency spectrum to be used between these repeaters, the failure cell's traffic could be redirected through the neighboring cells without increasing the co-channel interference experienced by any cell. As illustrated in Fig. 1, a center-excited standby repeater is placed at the center of each cell, while three corner-excited repeaters are placed at three vertices of the cell to relay the traffic to neighboring BSs when a redirection is needed. The corner repeaters serve to convert the carrier frequencies between the failure cell and its neighboring cells, so as to avoid co-channel interference problems. The obvious disadvantage of using wireless means here, however, is the need to reserve part of the scarce radio resource available as emergency spectrum. In order to alleviate the effect of under-utilization, we could dynamically utilize the emergency spectrum for some calls during normal operation and possibly drop some calls according to their priorities when a cell failure occurs.

In WATM networks, when a MT in an active call moves from one BS's coverage area to another BS's coverage area, handoff-rerouting protocols [5–8] are necessary to reroute existing active connections to the new BS. When a handoff occurs, a new partial connection is established between the new cell and the "Crossover Switch" (COS) [9, 10]. A COS is an "End User Mobility Enabled ATM Switch" (EMAS) [9] on the original path of the connection between a MT and the other end-user, which acts as an anchor while rerouting the

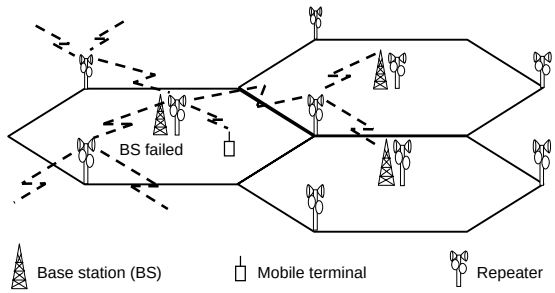

Fig. 1. Repeater Layout

connection to the new BS. Depending on the scheme chosen, different EMASs in the network could potentially become the COS for the rerouting process. For the case of connection rerouting during a cell failure, the connections from the failure cell are redirected in a manner similar to handoff rerouting. New partial connections are established between the neighboring cells providing the redirecting service and the respective COSs. In our work, we do not make any assumption on the COS discovery scheme that would be used. Interested readers are referred to [10] for extensive descriptions and performance evaluations of different COS discovery schemes.

In the following subsections, we shall describe two failure-handling schemes. The first scheme is a bandwidth reservation (BR) scheme that is targeted at achieving high survivability. The second scheme is a best-effort (BE) scheme that does not perform any reservation beforehand, and is targeted at achieving high bandwidth utilization.

A. Bandwidth Reservation Scheme (BR)

In order to achieve high survivability during a failure, bandwidths need to be reserved in advance. In this scheme, whenever a cell failure occurs, six neighboring cells with fixed preserved bandwidth would share the redistributed traffic from the failure cell. Since the failure cell has also reserved part of its bandwidth for emergency use, we only need to set aside 1/7 of the radio bandwidth available in each cell for emergency use. For example, assuming that each cell has a capacity of 70 units, but it only carries a maximum of 60 units of traffic while the remaining 10 units are reserved. When a cell failure occurs, only 60 units need to be redirected to the surrounding six neighbors, where each neighbor relays 10 units of traffic (1/7 of a cell's total capacity).

Besides reserving wireless bandwidths in the neighboring cells, bandwidth reservations in the backbone links are also needed. However, backbone reservations are only performed between the BSs in neighboring cells and their respective COSs. No reservations are required in those backbone links that are unaffected by the traffic redirection during failure-handling, which correspond to those links lying between the COS and the other ends of the connections.

We shall now describe the admission rule for a new call. Let B_{new} be the bandwidth requirement of the new call. Let R_l and E_l represent the remaining unused bandwidth and the emergency reserved bandwidth respectively in link $l \in L$, where L is the set of links comprising the entire end-to-end connection, inclusive of the wireless link. E_l is set to be 1/7 of the radio capacity if l is a wireless link. On the other hand, if l is a wired link, E_l would vary according to the connections that would be rerouted through l when a cell failure occurs. A new call request is accepted if

$$B_{\text{new}} \le R_l - E_l .$$

While 100% survivability is desired, it is only possible if sufficient bandwidths are reserved to accommodate the worst-case scenario in which the cell was fully loaded before the failure. From a user's point of view, forced termination of an ongoing call during handoff is more objectionable than the blocking of a new call. Therefore, handoff calls should be prioritized over new calls. Since cell failure does not occur frequently, it is reasonable for us to allow handoff calls to take advantage of the bandwidth reserved for failure-handling purpose. This would help to increase bandwidth utilization and decrease on-going handoff call dropping probability, at the expense of sacrificing 100% survivability. Thus, handoff calls are treated in a different manner; they are accepted if

$$B_{\text{handoff}} \le R_l ,$$

where B_{handoff} represents the bandwidth requirement of the handoff call. The links $l \in L'$ now represent the set of links between the target handoff cell and the COS, inclusive of the wireless link. Note that for new calls, the admission rules are performed on an end-to-end basis, but for handoff calls, the rules are performed only between the COS and the target handoff cell. This is because the route between the COS and the other end terminal remains unchanged after a handoff.

As handoff calls may occupy the bandwidth reserved for failure-handling purpose, we would need to terminate some calls when a cell failure occurs. We have the option of either terminating those calls that are currently utilizing the reserved spectrum, or we could terminate calls according to their priorities. For the latter case, if a higher priority call is currently utilizing the reserved spectrum, we may have to terminate a lower priority call outside the reserved spectrum, and reallocate its frequency for the higher priority call. Note that in either case, the same number of calls would be terminated.

When the reserved bandwidth in a cell is fully unused, there is a common understanding between this cell and its neighboring cells that it could provide 1/7 of its bandwidth in case of a cell failure. Thus there is no need to inform its neighbors every time it accepts a handoff request. On the other hand, whenever part of the reserved bandwidth is used by handoffs, the first-tier EMASs of its neighboring cells need to be informed of the reduction in reserved bandwidth. The EMAS could then pre-determine the connections that could survive in case of a cell failure, so as to speed up the actual failure-handling process when a real failure occurs. Note that a cell may not be attached to the same first-tier EMAS as its neighboring cells. As a result, the notification of available reserved bandwidth could happen across EMASs.

B. Best-effort Scheme (BE)

This scheme is targeted at achieving high resource utilization. There is no "reserved bandwidth", that is, the criterion for both new call and handoff call admission is

$$B_{\text{new/handoff}} \le R_l .$$

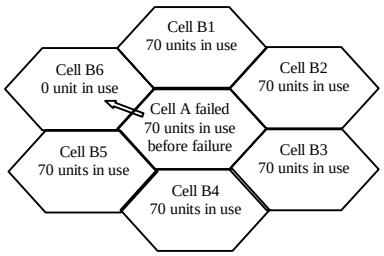

Fig. 2. Example scenario when BE scheme is applied.

In this scheme, system survivability is performed with best effort, depending on how much spare bandwidth is available from the neighboring cells to redirect the failure cell's traffic when a failure occurs. For example, as shown in the Fig. 2, suppose cells A, B1 to B5 are fully utilized, and cell B6 is fully unused. When cell A fails, its entire traffic could be redirected to cell B6 to achieve 100% survivability. From this point of view, the BE scheme provides a flexible backup strategy for the failure cell.

For every new or handoff call request that is accepted, or whenever a call is released from a cell, the first-tier EMASs of its neighboring cells need to be informed about the amount of bandwidth that is remaining for failure-handling purposes. Therefore, such notification messages occur much more frequently than in the BR scheme, which only requires notification upon the use or release of the reserved spectrum.

III. Important Issues for Traffic Redirection

In the following subsections, we shall look at two important issues that need to be considered when providing traffic redirection during a cell failure.

A. Switchover Time

When the traffic from the failure cell is redirected to its neighboring cells, new VPI/VCIs need to be assigned. As the failure occurs, the neighboring cells of the failure cell take some time to respond to the failure, including the time taken to get informed about the failure, as well as to get ready for the traffic redirection. In the following context, we shall refer to this period of response time as the "switchover time", defined as the time between the instant when a failure occurs and the instant when the neighboring cells are ready to redirect traffic for the failure cell.

In order to prevent data loss during the switchover time, we could equip each center-excited repeater with a data buffer for failure-handling use, where the size of the buffer is proportional to the length of the switchover time. Fig. 3 shows the reference model used for our switchover time calculation, while Fig. 4 shows the timing diagram of the corresponding switchover time.

In ATM networks, even when there is no data to transmit, ATM idle cells are still being transmitted in the links. Therefore, once the first-tier EMAS loses the signal from the corresponding BS for a certain amount of time, equal to a cell slot duration t_1, it could immediately detect the failure of either the BS or the link connected to the BS (ignoring the possibility of a false alarm).

As soon as the EMAS detects the failure, it begins to assign the alternative VPI/VCIs for the redirected traffic and notify

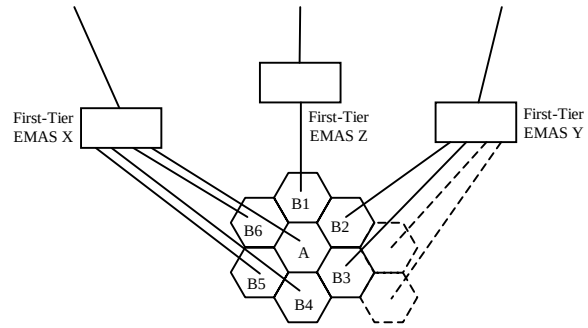

Fig. 3. Reference model for switchover time calculation.

the switch ports of the first-tier EMAS where the six surrounding cells are attached (processing time t_2). If the six surrounding cells are connected to different EMASs, an additional coordination time between EMASs (time t_3) would be introduced. When the EMAS(s) of the neighboring cells finish the assignment of alternative VPI/VCIs for the redirected connections, they need to signal the BSs of these cells (propagation delay t_4). The respective neighboring BSs would then perform some final processing (processing time t_5), before sending "READY" signals to the repeaters of the failure cell to activate them to redirect traffic (propagation delay t_6). The center-excited repeater can now begin to send out the packets stored in its buffer. Hence, we see that the size of the data buffer for failure-handling must be at least:

Radio bandwidth capacity $C \times$ switchover time ΔT,
where ΔT
= t_1 (cell slot duration) + t_2 (processing time in EMAS X)
+ t_3 (coordination time between EMASs X↔Y, and X↔Z)
+ t_4 (maximum propagation delay from X to B4-B6, Y to B2-B3, and Z to B1)
+ t_5 (maximum processing time in BSs B1-B6)
+ t_6 (maximum propagation delay of READY signals from BSs B1-B6 to A).

In order to estimate the magnitude of the required switchover time and the corresponding buffer size, we make a sample calculation based on the following assumptions:

$t_1 \sim 2.7\ \mu s$, t_2, t_4 & $t_5 \sim 30\ \mu s$, $t_3 \sim 200\ \mu s$, $t_6 \sim 10\ \mu s$.

This yields a switchover time of approximately 300 μs, and a buffer size of 6 KB. However, a practical choice for buffer size might be several times this value to give some tolerance.

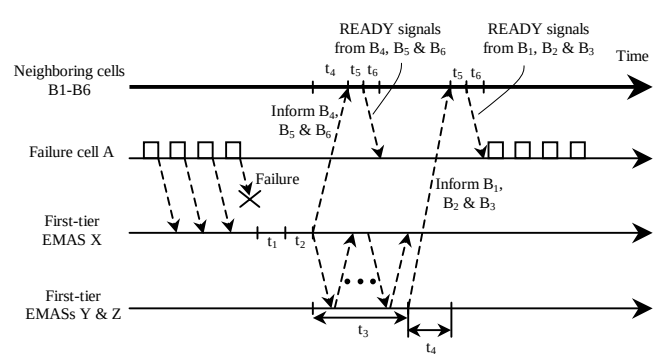

Fig. 4. Timing diagram of switchover time based on reference model in Fig. 3.

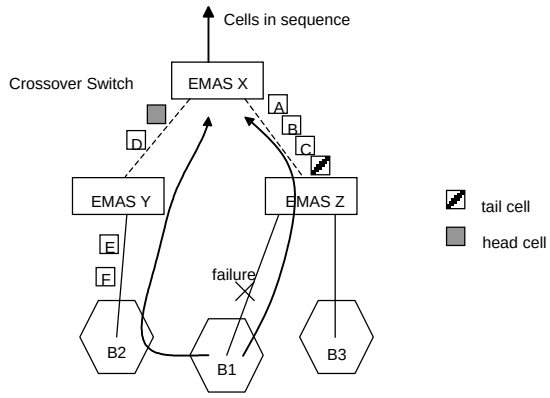

Fig. 5. Data integrity (out of sequence problem).

B. *Out Of Sequence Data*

When a failure occurs, the data redirected through the alternative path in the uplink might arrive at the destination earlier than the data sent out before the failure. This is because the same connection now takes a different path after a failure occurs. To ensure that cells arrive at the destination with the correct sequence, the COS should be responsible for reordering the out-of-sequence data when a failure occurs.

As illustrated in Fig. 5, suppose cell B1 fails after transmitting the ATM cells ABC. Cells DEF, which belong to the same connection as ABC, are transmitted after the failure. To ensure that cells ABC arrive at the COS (EMAS X) of BSs B1 and B2 before cells DEF do, the first-tier EMAS of the failure cell (EMAS Z) is responsible for transmitting a tail (ATM) cell for each connection after it detects the failure. BS B2, which is responsible for redirecting the traffic of the connection, needs to insert a head (ATM) cell for the connection before sending the cells DEF. With this mechanism of head/tail cell insertions, the COS could correct the order of the received packets when failure occurs. That is, if EMAS X receives the head cell before the tail cell, it needs to queue up all the subsequent ATM cells from that connection until it receives the tail cell of the same connection. As for the problem of downlink data integrity, we might have to rely on higher layer protocols to accomplish the task.

IV. SIMULATION MODEL

Our simulation model resembles the one adopted in [11]. A metropolitan area with 256 cells is assumed, as shown in Fig. 6. However, in contrast to [11], we have removed the effects of the backbone by assuming that it would not become bottlenecks. This is done to simplify the presentation and interpretation of results. Note however, that our proposed schemes are designed to accommodate situations where the backbone links could also become potential bottlenecks. We shall investigate the inclusion of potential backbone bottlenecks in an extended version of this paper.

Three kinds of traffic types are randomly generated in the system, with parameters as shown in Table I. The call arrivals in each cell are assumed to follow a Poisson distribution, with an average of λ calls/s/cell, and each call has equal probability of belonging to any of the three traffic types. We assume that each MT in an active connection communicates with a fixed wireless terminal, whose traffic should also be

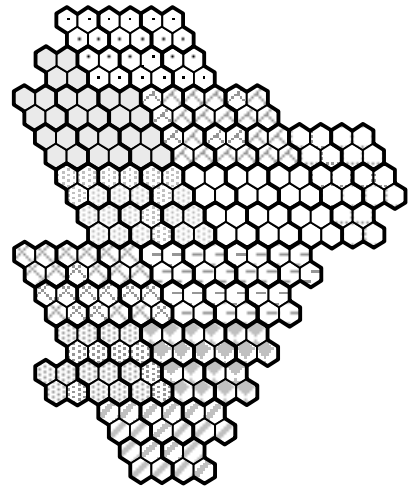

Fig. 6: Cell layout used in simulations.

redirected during a cell failure. Our MT's mobility model is similar to the one adopted by [11]. We assume that a MT's initial position within the cell during call establishment is uniformly distributed. Its initial speed is chosen from a Gaussian distribution, with mean S_{mean} and standard deviation σ_{speed}, but limited to the range [0, S_{max}]. We also assume that each MT will travel towards a random virtual destination cell (VDC) within the network with some degree of randomness, which will be explained shortly. A MT's initial direction of travel is assumed to be heading towards the center of the VDC. Its speed and direction may change according to the Bernoulli Process with probability P, at time intervals ΔT. The change in speed is chosen from a zero mean Gaussian distribution having a standard deviation of $\sigma_{speed\Delta}$. However, the new speed is always confined to the range [0, S_{max}]. The new heading of each MT, on the other hand, is chosen from a zero mean Gaussian distribution having a standard deviation of $\sigma_{heading\Delta}$, from the current heading towards the center of the VDC. Additional system parameters are shown in Table II.

We shall evaluate the following performance metrics with respect to the system's "normalized offered load", including survivability, bandwidth utilization of radio links, new call's blocking probability (CBP) and handoff dropping probability (HDP). For an average call arrival rate of λ calls/s/cell that is uniform over the entire system, the "normalized offered load" is defined as the offered load per cell from new calls, normalized by the radio bandwidth capacity C of a cell:

$$\lambda / 3C \sum_{i=1}^{3} E[\tau_i] \times B_i \ ,$$

where $E[\tau_i]$ represents the average type i call duration in seconds, B_i represents the average bandwidth for type i in Mb/s, and C is the radio bandwidth capacity per cell in Mb/s.

TABLE I
TRAFFIC TYPES USED IN SIMULATIONS

Traffic Class	Bandwidth Requirement	Connection Duration	Average Connection Duration
1	30 kb/s (CBR)	1 – 10 min	3 min
2	256 kb/s (CBR)	1 – 30 min	5 min
3	1-6 Mb/s (VBR) 3 Mb/s (mean)	5 – 60 min	10 min

TABLE II
MISCELLANEOUS SIMULATION PARAMETERS

Parameter	Value	Description
R	500 m	Cell radius
C	155 Mb/s	Radio bandwidth capacity per cell
P	0.1	Bernoulli probability for change in speed/direction
ΔT	1 s	Time interval for the above Bernoulli process
S_{max}	120 km/h	MT's maximum speed
S_{mean}	60 km/h	MT's mean initial speed
σ_{speed}	20 km/h	Standard deviation of MT's initial speed
$\sigma_{speed\Delta}$	10 km/h	Standard deviation for change in MT's speed
$\sigma_{heading\Delta}$	30°	Standard deviation for change in MT's direction

We evaluate the average, minimum and maximum survivability of the system. The way we calculate the average survivability at a sampling instant is by assuming single cell failure (without actually failing the cell) and observing how many connections in the failure cell could survive at that instant should a real failure occur. Obviously, we could repeat the procedure for all cells in the whole system, and then compute the average survivability values for each traffic class at the sampling instant. The minimum/maximum survivability estimate, on the other hand, is chosen from the cell that has the lowest/highest survivability value among the entire system. We collect these survivability estimates over many sampling instants such that their 95% confidence intervals are within ±5% of their mean values.

We have also defined two variants of survivability measures for the overall survivability of the three traffic classes, in terms of number of connections and bandwidth:

$$\text{overall connection survivability} = \frac{\sum_{i=1}^{i=3}(\text{No. of survivable type } i \text{ connections})}{\sum_{i=1}^{i=3}(\text{No. of type } i \text{ connections})}, \text{ and}$$

$$\text{overall bandwidth survivability} = \frac{\sum_{i=1}^{i=3}(\text{No. of survivable type } i \text{ connections} \times B_i)}{\sum_{i=1}^{i=3}(\text{No. of type } i \text{ connections} \times B_i)}.$$

V. SIMULATION RESULTS

In our first set of simulations, we do not implement any failure-handling scheme. The CBP and HDP values as well as the normalized radio bandwidth utilization obtained for different normalized offered load are shown in Fig. 7(a) and (b). These simulation results are to be used as references to track how they are affected when we implement our failure-handling schemes in subsequent simulations.

A. Bandwidth Reservation Scheme (BR)

In this simulation, we allow handoff call requests to utilize the reserved bandwidth, but not the new call requests. Fig. 8(a) and (b) show the corresponding average, minimum, and maximum survivability values obtained for varying normalized offered load. As shown in Fig. 8(a), the average bandwidth survivability is above 0.94. Its minimum value is also above 0.88, as shown in Fig. 8(b).

The tradeoff for the high survivability is an increase in CBP and a reduction of overall bandwidth utilization. Fig. 8(c) to (e) show the changes in the respective performance measures over the results shown earlier in Fig. 7. Note that by allowing handoff requests to take advantage of the reserved bandwidth, we have prioritized handoffs inherently, causing the HDP values to improve significantly over the case where no failure-handling is performed. Link utilization has also decreased from the case without failure-handling due to the reservation of bandwidths for failure-handling purpose.

B. Best-Effort Scheme (BE)

Fig. 9(a) and (b) show the simulation results for the BE scheme. As explained in Section II-B, the BE scheme does not make any reservation. Therefore, we do not show the results for CBP, HDP, and normalized radio bandwidth utilization because they are almost identical to those shown in Fig. 7(a) and (b). This implies that the BE scheme has better radio bandwidth utilization but poorer HDP performance compared with the BR scheme. As shown in Fig. 9(a), the average bandwidth survivability is approximately 0.925 when the normalized offered load is 1.0. High average survivability could still be maintained because each neighboring cell could now have the flexibility to provide any amount of spare bandwidth that is remaining for failure-handling purpose, without the restriction of 1/7 radio bandwidth capacity as in the BR scheme. Note however, that the minimum survivability values as shown in Fig. 9(b) have much lower values (approximately 0.73 at a normalized offered load of 1.0) compared to the BR scheme because no bandwidth is specially reserved for failure-handling purpose.

VI. CONCLUSION

In this paper, we have proposed a survivable wireless ATM network architecture that allows it to survive from a single cell failure condition, which could arise when either the BS, or the link directly connected to the BS, fails. Two failure-handling schemes were described. The first scheme – the BR scheme, aims to achieve high survivability performance by reserving bandwidths at appropriate locations in the network, at the expense of lowering bandwidth utilization and increasing the CBP for new calls. The reserved bandwidth could only be used by handoff call requests, but not by new call requests. The second scheme – the BE scheme, does not perform any bandwidth reservation, in which system survivability depends on the network load when a failure occurs. From the simulations, we find that the BR scheme achieves good survivability as expected. On the other hand, the BE scheme is able to achieve an average survivability that is close to that of the BR scheme, while having higher utilization, due to its flexible best-effort assignment of spare bandwidths from neighboring cells. However, the BE scheme requires more frequent update messages between the EMASs in order to update each other on the amount of bandwidth that they could provide for failure-handling.

In addition to performing simulations to verify the above expectations, we have also considered important issues such as switchover time computation, and data integrity maintenance, for our proposed schemes.

(a) (b)

Fig. 7. Simulation results without failure-handling. (a) CBP and HDP. (b) Normalized radio bandwidth utilization.

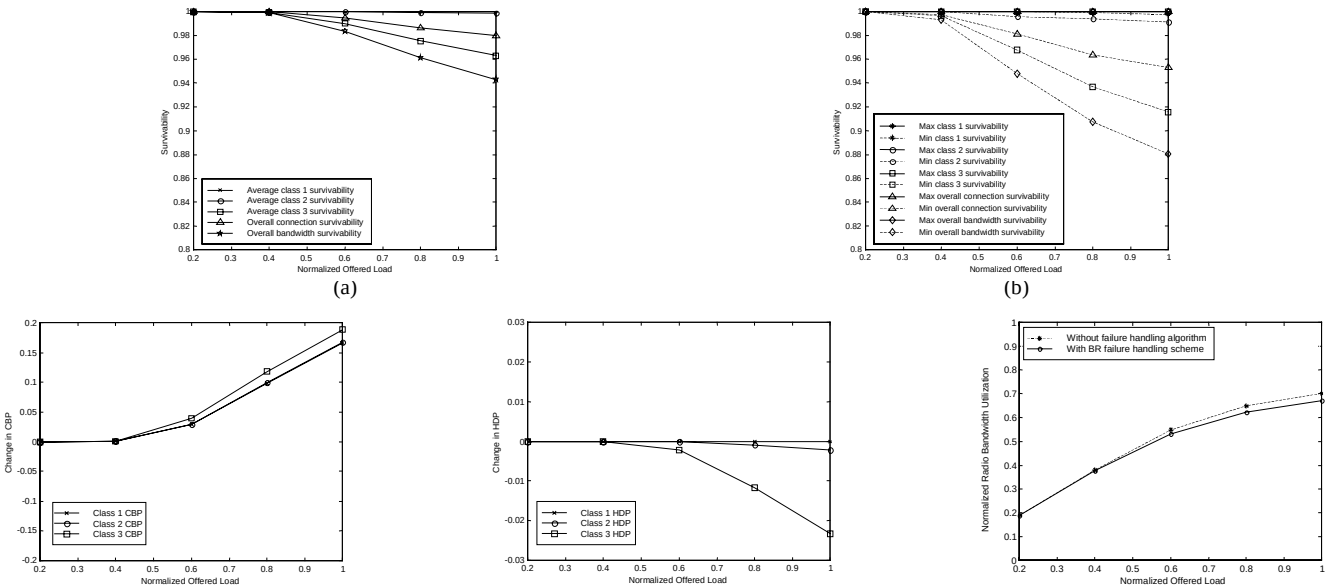

(a) (b) (c) (d) (e)

Fig. 8. Simulation results using BR failure-handling scheme. (a) Average system survivability. (b) Maximum and minimum system survivability. (c) Change in CBP with respect to Fig. 7(a). (d) Change in HDP with respect to Fig. 7(a). (e) Normalized radio bandwidth utilization.

(a) (b)

Fig. 9. Simulation results using BE failure-handling scheme. (a) Average system survivability. (b) Maximum and minimum system survivability.

REFERENCES

[1] K. Murakami and H.S. Kim, "Virtual path routing for survivable ATM networks," *IEEE/ACM Trans. Networking*, vol. 4, no. 1, Feb. 1996, pp. 22–39.

[2] W.D. Grover, "The self-healing network: A fast distributed restoration technique for networks using digital cross-connect machines," in *Proc. IEEE GLOBECOM*, Dec. 1987, pp. 1090–95.

[3] W.D. Grover, T.D. Bilodeau, and B.D. Venerables, "Near optimal spare capacity planning in a mesh restorable network," in *Proc. IEEE GLOBECOM*, Dec. 1991, pp. 2007–12.

[4] B.D. Venerables, W.D. Grover, and M.H. MacGregor, "Two strategies for spare capacity placement in mesh restorable networks," in *Proc. IEEE ICC*, May 1993, pp. 267–71.

[5] B.A. Akyol and D.C. Cox, "Handling mobility in a wireless ATM network," *Proc. IEEE INFOCOM*, Mar. 1996, pp. 1405–13.

[6] F.-C. Cheng and J.M. Holtzman, "Wireless intelligent ATM network and protocol design for future personal communication systems," *IEEE J. Select. Areas Commun.*, vol. 15, no. 7, Sep. 1997, pp. 1289–07.

[7] D. Raychaudhuri *et al.*, "WATMnet: a prototype wireless ATM system for multimedia personal communication," *IEEE J. Select. Areas Commun.*, vol. 15, no. 1, Jan. 1997, pp. 83–95.

[8] A. Acharya *et al.*, "Design and prototyping of location management and handoff protocols for WATM networks," *IEEE ICUPC*, Sep. 1996.

[9] ATM Forum, "Framework for dynamic COS discovery in wireless ATM," Feb. 1998.

[10] C.K. Toh, "Performance evaluation of crossover switch discovery for wireless ATM LANs," in *Proc. IEEE INFOCOM*, Mar. 1996, pp. 1380–87.

[11] W.-S. Soh and H.S. Kim, "Dynamic bandwidth reservation in hierarchical wireless ATM networks using GPS-based prediction," in *Proc. IEEE VTC 1999-Fall*, pp. 528-32, Sep. 1999.