

DRA: A Dependable Architecture for High-Performance Routers

Malcolm Mandviwalla and Nian-Feng Tzeng

Center for Advanced Computer Studies

University of Louisiana at Lafayette

Lafayette, LA 70504, USA

Email: <mhm4231,tzeng>@cacs.louisiana.edu

Abstract

Due to a relentless increase in the amount of mission-critical real-time traffic over a revenue-generating network, dependability has come forth as a vital measure for such a network. To prevent extended downtimes and loss of critical data, most commercial routers achieve fault-tolerance by adding substantial redundancies for critical components. However, they often fail to utilize existing resources efficiently, unduly limiting their achievable dependability. To significantly improve the dependability of existing routers, we propose an efficient dependable architecture for high-performance routers, called dependable router architecture (DRA). DRA augments existing resources to create an additional level of interconnection across linecards for channeling resources from non-faulty linecards to linecards with faulty components. We analyze DRA using Markov models to assess its dependability improvement.

1. Introduction

The Internet consists of a large number of routers; thus, dependable service over the Internet naturally calls upon high router dependability. This fact is widely acknowledged by router manufacturers, who include high levels of explicit component redundancy to make their devices dependable. For example, the Cisco 12000 series routers [1] and the Juniper T640 distributed routers [2] include not only duplicated power supplies and fans, but also redundant switching fabrics and route processors. Additionally, most service providers use two or more routers within a point-of-presence to increase the network availability. Explicit redundancy, though costly, cannot be avoided for certain critical single-unit components, such as the power supply, the fabric, and the route processor. However, key functionality of a high-performance router lies in its linecards (LC's). Since a router consists of multiple LC's (which are functionally similar to each other), significant cost-savings as well as higher dependability measures can be achieved with a provision to make healthy LC's cover a faulty one, giving rise to fault-tolerance. The capacity of the Internet is colossal, but most

of this capacity is unused, with the mean link utilization being somewhere in-between 10-15% [3]. Hence LC's usually have workloads much lower than 100%, making it possible for them to serve routing requests of a failed LC. None of the commercially available routers take advantage of this fact, nor do they have any hardware redundancy at the LC's, rendering LC's non-dependable.

A single LC component failure brings down all its interfaces (i.e. ports), which continue to remain offline either until the LC with the faulty component is replaced (note that LC's are generally hot-swappable) or a switch is made to a standby router. Those routers are thus not fully fault-tolerant, despite that non-faulty LC's have lots of unutilized capacity, which could otherwise be exploited to cover a faulty LC, if an appropriate provision is introduced to permit this fault coverage.

To this end, we propose a simple architecture for fault-tolerant high-performance routers, called dependable router architecture (DRA), which allows multiple LC's to cover a faulty LC through an interconnect derived from augmenting a bus present in any commercial router, yielding a highly available system. Our proposed DRA channels unutilized resources (which in current systems are often abundant [3]) from healthy to faulty LC's, without resorting to any redundant LC. It enhances router dependability in a distributed manner by dealing with units along the routing path (such as LC components and switching fabric ports) and central to the basic functionality of the router, while ignoring other units which are usually duplicated with sufficient redundancies (such as fans, power supplies, etc.).

Current commercial implementation of LC's uses different ASIC's for different protocol types, leading to not only inflexible and expensive LC's [4], but also potentially low reliability. Given that all LC's include both common functionality independent of the protocol implemented therein and protocol-dependent functionality, DRA proposes to move all the protocol-dependent functionality of an LC into a separate unit at the LC; the unit can then be realized by an ASIC or an FPGA (which is) programmed to the designated protocol implementation. Every LC now consists of common components plus a programmable (or an ASIC) unit to implement the given protocol. This way of separating protocol-independent functionality from

This work was supported in part by the National Science Foundation under Grant CCR-0105529

protocol-dependent one lowers the design and deployment costs, while making it easier to accommodate future extensions and protocol updates. Above all, this functional separation enables DRA to channel resources among LC's even if they implement different protocols. Our dependability analysis validates that DRA enhances the dependability of a router significantly at a low cost.

2. Related Work

Recently, most manufacturers have adopted distributed architecture [5] for their products, including Cisco 12000 series routers [1] and Juniper T-series routers [2]. A basic distributed router (BDR) architecture comprises LC's, a route processor and a main routing engine, a switching fabric, and an internal bus (for maintenance), as shown in Figure 1. The route processor (RP) is generally a powerful processor [5], which runs the applications and protocols supported by the router to realize its functions. Copies of the routing table are distributed by the RP to the local forwarding engine (LFE) in each LC.




Figure 1. Distributed router architecture.

A switching fabric provides connectivity among LC's in the router, for data transfer from an incoming LC to an outgoing LC. Since the switching fabric is central to the functionality of a router, manufacturers often add sufficient redundancy to make it fault-tolerant. For example, a Cisco 12000 series router includes five switching fabric cards, where four are active and the fifth provides 1:4 redundancy.

Router components communicate among themselves using the switching fabric and also via an internal bus [2], [6], which performs the following functions: (1) Discovery of system cards at startup by the RP, (2) Collection and dissemination of maintenance information (like voltages, temperature, etc.), and (3) Possible dissemination of route updates and other control information to linecards.

Much router functionality is implemented in the linecards (LC's), each of which contains the following functional units: (1) physical interface units (PIU's), (2) the segmentation and reassembly unit (SRU), and (3) a local forwarding engine (LFE). A router sends and receives traffic over its LC ports. An LC may have one or multiple ports for terminating external links. Each port at an LC communicates with other functional units of the LC through a PIU. A PIU accepts input traffic over a media-

specific interface (such as Ethernet, SONET, etc.) and converts it into an electronic form (if necessary). It then detects the packet/frame boundaries of the protocol for that interface type, extracts the data payload (IP/ATM), and forwards it to the SRU.

The SRU receives the data payload from the PIU. It extracts header and control information before segmenting the packet into fixed-length cells for transfer over the switching fabric. The SRU also sends the packet's destination address to the LFE for packet lookup, classification, and filtering. The lookup results are returned to the SRU, which on receiving these results sends the cells constituting the packet over the switching fabric to the destination LC, where they are reassembled into the original packet by the SRU and then transferred to the output PIU. The SRU and LFE are protocol-independent, identical for every LC type.

3. Dependable Router Architecture (DRA)

As mentioned earlier, commercial routers do not tolerate LC failures, despite that LC's carry out most of the critical functions equipped in the routers. The only way to provide fault tolerance at the LC's in existing systems is to add at least one redundant LC for each protocol type supported by the router (which is clearly an expensive proposition). It is thus essential to devise a cost-effective mechanism for LC fault-tolerance so that high overall router dependability can be achieved affordably. To this end, DRA's LC structure is introduced, followed by a description of its fault model.

3.1. LC Structure under DRA

DRA redefines two aspects of current distributed routers in order to provide fault-tolerance for LC's without adding any redundant LC. First, it redistributes the functionality of an LC so as to differentiate the protocol-dependent and protocol-independent functions, with the former clustered at the "protocol-dependent logic unit", as depicted in Figure 2. Secondly, the internal bus (for maintenance originally) is modified and upgraded so that it can channel resources from non-faulty to faulty LC's.

The functional units of an LC under DRA include (1) Physical Interface Units (PIU's), (2) Protocol-Dependent Logic Unit (PDLU), (3) Segmentation and Reassembly Unit (SRU), and (4) Local Forwarding Engine (LFE) as illustrated in Figure 2.




Figure 2. DRA's LC structure.

All protocol-dependent (PD) functionality (except IP) is moved from the PIU and SRU of BDR (see Figure 1) to the PDLU, such that resources from non-faulty LC's can thus be channeled to faulty LC's even if they implement

different protocols. This leads to the following advantages: (1) providing fault-tolerance without the need of adding any redundant LC and (2) employing identical units in every LC, except for the PDLU realized by an FPGA (that is programmed) or an ASIC to implement the given protocol for possibly lowering the development and deployment costs. Under DRA, failures at the SRU, LFE, and over the switching fabric can be covered by any LC (even if it implements a different protocol). If a failure occurs at the PDLU of an LC, however, only a healthy LC with the same protocol may cover this failed LC.

As shown in Figure 2, an input traffic stream received by the PIU is forwarded to the PDLU, which detects packet/frame boundaries of the Layer 2 protocol and extracts the IP payload from it. PDLU also extracts the L2 and L3 headers, and sends them along with the IP payload to the SRU. The process is reversed for output traffic, where the PDLU receives the payload from the SRU, attaches the appropriate headers, encloses it in the L2 specific envelope, converts it to the media-specific protocol specifications, and sends it to the PIU. The remaining *basic* functionality of PIU, SRU and LFE is identical to that of BDR counterpart.

DRA requires a high bandwidth connection across all LC's, and such a connection can be obtained by upgrading the existing internal bus and modifying its functionality to enable its coverage of a faulty LC, referred to as the *enhanced internal bus* (EIB), as depicted in Figure 2. When no failures exist, LC's in DRA route packets through the regular switching fabric. If an LC component fails, communication through the switching fabric may be hampered, rendering an LC unable to route data to and from all its port interfaces. In such a case, other LC's communicate with the faulty LC through the EIB. The communication between source and destination LC's (over the EIB) may require stepping through an *intermediate* LC (say, LC_{inter}). This LC_{inter} is selected based on the location of the fault (occurring at the PDLU, SRU or LFE). For efficient functioning, all LC's store information about the location of faults existing in the router. This is easily achieved through the exchange of control packets over the EIB (explained later).

3.2. DRA Fault Model

DRA considers failures at components only along the routing path, and the fault model described is functional in nature. It ignores hardware failures at units outside the routing path (such as RP, fans, and power supplies). We focus on permanent faults, which manifest themselves as hardware failures of individual router components, leading to a loss in packet delivery service to one or more LC ports. They are rectified only if the faulty component is replaced (or a swap is made to a redundant unit). Our goal is to arrive at a router able to continuously deliver packets in a timely fashion, even in the presence of faults. The fault model of DRA emphasizes the location of a fault and how

to avoid service disruption (unavailability) due to the fault. Based on the fault location, DRA has a provision to make healthy LC components available to cover a faulty LC.

To facilitate subsequent discussion, a packet stream in transit is said to be along the forward (or reverse) path if it originates at (or is destined for) a faulty LC. Consider a traffic flow from an incoming LC (say, LC_{in}) to an outgoing LC (say, LC_{out}). If a failure is encountered by this traffic flow, the failure may reside (1) over the switching fabric between LC_{in} and LC_{out} , (2) at LC_{in} , or (3) at LC_{out} . In any case, the EIB is employed to transfer data and lookup requests among LC's (see Figure 3). The EIB is never invoked if no traffic flow encounters a failure.

(Case 1) A failure over the switching fabric. Given adequate redundancy for the fabric, the failure poses no service disruption and is tolerable.

(Case 2) A failure at LC_{in} , namely, at the PIU, PDLU, SRU, or LFE of LC_{in} . For a failure at the PIU, packet transfer is stalled. If it is at the PDLU (or SRU), the PIU (or PDLU) of LC_{in} transfers the incoming packets over the data lines to the PDLU (or SRU) of a healthy LC with sufficient resources available. For the failure at the PDLU, the healthy and faulty LC's must support the same protocol. If the failure occurs at the LFE, the SRU of LC_{in} transfers the lookup request of each incoming packet over the EIB to the LFE of any healthy LC for carrying out the table lookup, with the lookup result sent back over the EIB to the SRU of LC_{in} . The incoming packets are then delivered over the switching fabric to their outgoing linecards, guided by the lookup results.

(Case 3) A failure at LC_{out} , namely, at the PIU, PDLU, SRU, or LFE of LC_{out} . If it occurs at the PIU, data transfer ceases. If the failure is at the PDLU, two alternatives exist. First, if both LC_{in} and LC_{out} implement the same protocol, the PDLU of LC_{in} , subsequent to packet processing by the SRU and LFE, sends packets over the EIB directly to the PIU of LC_{out} . Second, if the protocol implementations are different, cells are sent over the switching fabric to an LC_{inter} (i.e., an intermediate LC) whose PDLU then sends the reassembled packets to the PIU of LC_{out} . Finally, if the failure is at the SRU, LC_{in} sends the reassembled data through its SRU to the PDLU of LC_{out} .

4. DRA Implementation Details

One implementation of our DRA is detailed in Figure 3, where the EIB is split into control lines and data lines separately for high performance. The control lines are used mainly to arbitrate access to the data lines in a distributed manner, and to exchange lookup requests and replies. Each LC houses a simple bus controller. A distributed bus offers several advantages: (1) it contains a set of passive lines, which have extremely low failure rates, (2) the failure of a single bus controller (at any LC) does not affect transmission over the bus, and (3) it can handle packets of variable lengths directly without segmenting them into cells of a fixed length before delivery (unlike over the crossbar

The bus controller is an essential component for providing fault-tolerance at an LC, and it is crucial to be made very simple and highly reliable. The bus controller supports: (1) CSMA/CD over the control lines, (2) simple round-robin based time-division multiplexing to permit bus bandwidth sharing among multiple LC's, and (3) switching information between the EIB and the units inside an LC.



EIB Protocol

Addressing Tier:

either an LC_{in} or an LC_{inter} , depending on the type of communication (specified by the communication tier).

After the LC's have been identified, the type of communication within a control packet needs to be recognized. This tier initiates different control packets and distinguishes them based on the type of communication in progress, essential in deciding the kind of processing required. The different types of control packets transferred over the control lines are listed below.

- Processing Tier:*

- **Data rate:** Included in all data communications, this parameter gives the rate of transmission requested by LC^{init} . This parameter is used by an LC^{rec} to check if it has sufficient resources to accept the stream at that rate.
- **Protocol type:** This parameter is used to distribute information about the protocol implementations of the LC's, necessary in identifying a proper LC^{inter} .
- **Faulty component:** The faulty component of an LC is identified by this parameter. This information is used by healthy LC's to determine whether to send data directly to the faulty LC as packets (to PDLU, possibly via an LC^{inter}) or as cells (to SRU).
- **Lookup address/result:** In communications involving a lookup request, this parameter denotes the IP address for the lookup operation specified by LC^{init} , and the lookup result returned by LC^{rec} .

Detailed description of transfer of control and data carried out over the EIB is explained next.

- 

and the fault location. All candidate LC_{inter} 's then examine the stream's data rate to determine if they can accept it. For a failure of the PDLU at LC_{in} , an additional check is made to see if LC_{inter} and LC_{in} implement the same protocol. All LC_{inter} 's able and qualified to accept the stream then initiate REP_D broadcast, indicating their intention to accept the stream. The first LC_{inter} to successfully acquire the control lines for this purpose is the one which will process the packet stream, and all other LC's on hearing the REP_D , know that the stream has been accepted for processing and terminate their own REP_D broadcasts, if any. It may so happen that multiple LC_{inter} 's can accept the stream, in which case there may be a collision of replies over the control lines (collisions are handled by the CSMA/CD protocol). On receiving an REP_D packet, LC_{in} starts transferring data over the data lines for a duration dictated by its bandwidth requirement and the other concurrent data transfers (via time multiplexing) over the EIB. An LC may additionally initiate an REL_D broadcast under conditions explained earlier. On receiving an REL_D , LC_{in} reinitiates an REQ_D and the whole process repeats.

- (b) *Reverse path*: The communication sequence along the reverse path is similar to that for a stream along the forward path. Along the reverse path, LC_{init} (LC_{in} or LC_{inter}) initiates an REQ_D , sent to the faulty destination LC (LC_{out}). An LC on receiving an REQ_D control packet from LC_{init} and identifying it as LC_{out} , initiates an REP_D over the control lines to LC_{init} , which then starts the data transfer. LC_{init} initiates an REL_D when it detects that it no longer has data destined for LC_{out} .
- (c) *Lookup*: Servicing of lookup requests from a failed LFE occurs entirely over the control lines. The lookup address is included within an REQ_L control packet and sent over the control lines. The REP_L packet containing the lookup result is often smaller in size than the lookup address itself. Hence, to reserve the data lines exclusively for larger (and faster) data transfers, the smaller lookup replies are conveniently enclosed within a control packet.

EIB Scheduling and Arbitration

The bus controllers implement a distributed protocol to access EIB. The exchange of REQ_D and REP_D sets up a logical path (LP) over the data lines between LC_{init} and LC_{rec} (as depicted in Figure 4). Due to the parameters specified by the processing tier, each LC has a global view of the faulty component locations and all the data streams being transferred. The allocation of the data lines for use by concurrent multiple LP's follows a simple round-robin time-division multiplexing scheme. The bandwidth taken by an LC is proportional to its requirement posted and accepted during its LP setup process, and is limited by the total bandwidth capacity of the EIB. The bandwidth amount given to an LC can be calculated as follows:

Let B_{LC} be the bus bandwidth asked by an LC, B_{LCT} be the total bus bandwidth requested by all LC's, B_{BUS} be the

bandwidth capacity of the data bus, and B_{prom} be the bandwidth amount promised for the LC. If $B_{LCT} \leq B_{BUS}$, then $B_{prom} = B_{LC}$. If $B_{LCT} > B_{BUS}$, however, all the requesting LC's scale back their transmission rates accordingly by dropping packets, to arrive at, $B_{prom} = (B_{LC} / B_{LCT}) \times B_{BUS}$.




Figure 4. Scheduling of data transfers.

To facilitate data line sharing, each LC maintains three counters (1) Ctr_β , (2) Ctr_{LC} , and (3) Ctr_c . Let β , which is stored in Ctr_β , be the total number of requesting LC's sharing the data lines at any time. This value of β is known to all LC's in the router and is incremented (or decremented) each time an LP is established (or released). Each LC_{init} is assigned with a unique ID (which is stored in Ctr_{LC}), depending upon the order when they finish establishing their logic paths. In Figure 4, for example, LC_{init} 1 requests and completes an LP establishment first so it is assigned with ID = 1. Subsequently LC_{init} 2 finishes another LP establishment and is assigned with ID = 2. These two LP's take turns to utilize the data lines, in succession as demonstrated in Figure 4. As soon as an LP has been established by an LC_{init} , the following assignments are made (in their order) to LC_{init} 's counters. (1) Ctr_{LC} is set to $Ctr_\beta + 1$ (viz. the assigned ID for LC_{init}), (2) Ctr_c is initialized to $4 \times Ctr_\beta + 1$, and (3) Ctr_β is incremented by one indicating the successful addition of a new LP.

The reason for initializing Ctr_c with $4 \times Ctr_\beta + 1$ is to ensure correct functioning after the first LP has been established and also each subsequent completion of LP establishment. Ctr_c is decremented by one, when a requesting LC finishes its turn of utilizing the data lines for transmitting its data existing in its buffer, realized by lowering the control line L_t (which is seen by all LC's, and thus causes every Ctr_c to be decremented simultaneously). If Ctr_c of an LC_{init} equals its assigned ID (stored in Ctr_{LC}), LC_{init} starts to transmit its data held in its buffer. Upon completing its transmission LC_{init} lowers L_t to signal the next requesting LC for transmission. When any Ctr_c reaches zero, another control line L_β will be raised, forcing every LC_{init} to load Ctr_c with value β . Now, the most recently added requesting LC has its first turn to utilize the data lines for transmission. This process repeats, and every requesting LC gets its turn to use the data lines in sequence.

When an LP is released (by broadcasting an REL_D), the ID of the requesting LC (LC_{init}) which started that LP, say ID_r , is announced over the control lines by enclosing it within the REL_D . This announcement makes (1) Ctr_β decremented by 1, and (2) the ID of requesting LC (stored in Ctr_{LC}) decremented by 1, provided that the ID is larger than ID_r . This arbitration on accesses to the data lines is simple and efficient.

5. Dependability and Performance Analysis

In order to assess router dependability enhancement resulting from DRA, we analyze the dependability aspects of LC's under DRA using Markov models. As dependability is a collective measure of availability and reliability, models for both availability and reliability are developed, with the results of DRA compared with those of its BDR counterpart. To simplify our analysis of DRA, the switching fabric is assumed to always function properly due to its adequate redundancy incorporated. A PIU is assumed to be fault-free since its failure disconnects an external link from the router. Our analysis takes the following constant, exponentially distributed component failure rates (in *hours*).

- (1) $\lambda_{LC} = 0.00002$, failure rate of an LC, which can be further split into:
 - $\lambda_{LPD} = 0.000006$, failure rate of the *local* PDLU (where local implies the LC under analysis (LC_{UA})),
 - $\lambda_{LPI} = 0.000014$, failure rate of the *local* protocol independent (PI) units.
- (2) $\lambda_{BC} = 0.000001$, failure rate of LC_{UA} 's bus controller.
- (3) $\lambda_{BUS} = 0.0000001$, failure rate of the EIB.

The LC failure rate assumed above is akin to that of a Cisco 7000 series OC-48 LC [8]. The EIB is a set of passive lines, which are simpler and have a lower failure rate than the bus controller. Naturally, the bus controller is much simpler than an LC. The following additional definitions and assumptions are applicable to the Markov models presented next:

- (1) For a data stream originating at LC_{UA} , LC_{out} is assumed to be fault-free and hence is not considered in the analysis. This assumption avoids situations where a data stream may need to pass through more than one LC_{inter} to reach LC_{out} .
- (2) The router has N LC's, comprising an LC under analysis (LC_{UA}), an LC_{out} , and $(N-2)$ LC_{inter} 's. Out of these N LC's, $(M-1)$ LC's implement the same protocol as LC_{UA} , and hence they have identical PDLU's. For a given LC_{UA} we have $(N-2)$ intermediate PI units and $(M-1)$ intermediate PDLU's in the router.
- (3) While an LC_{inter} may fail both at the PDLU and at the PI units, LC_{UA} can fail either at the PI units or at the PDLU, but not at both.
- (4) λ_{IPD} is the combined failure rate of an LC_{inter} 's PDLU and its bus controller ($\lambda_{IPD} = 0.000007$). Similarly, λ_{IPI} is the combined failure rate of an LC_{inter} 's PI units and its bus controller ($\lambda_{IPI} = 0.000015$). These failure rates indicate that if either the bus controller or the PDLU fails (or PI units fail), that LC_{inter} is unable to cover a faulty PDLU (or PI units).

5.1. Reliability

The Markov models illustrated in Figure 5 are used to analyze the reliability of an LC, which is the probability that packets can be transferred successfully to and from that LC at any time t in $(0, t)$. For the DRA model given in

Figure 5 (b), there are up to $(N-2)$ LC_{inter} 's available to cover a faulty LC.

The states of the Markov model represent the state of LC_{UA} , as defined below:

- State F is the state where data transfer through LC_{UA} has stopped due to the failure of its bus controller, or the EIB, or the failure of all $(N-2)$ LC_{inter} PI units or $(M-1)$ LC_{inter} PDLU's.
- State $(0, 0)$ is the state of the LC where no failures occur at the LC's, the EIB, and the bus controller.
- T' is the state where only the EIB or the LC_{UA} 's bus controller has failed (and packet transfer continues via the switching fabric).
- States i_{PD} (where $0 \leq i \leq M-2$) are states where i of $(M-1)$ LC_{inter} PDLU's have failed, after the failure of the PDLU of LC_{UA} . Similarly, i_{PI} (where $0 \leq i \leq N-3$) are states where i of $(N-2)$ LC_{inter} PI units have failed, after the failure of the PI units of LC_{UA} .
- States (i, j) (where $1 \leq i \leq N-3$ and $1 \leq j \leq M-2$) are the states where i of $(M-1)$ LC_{inter} PDLU's, and i of $(N-2)$ LC_{inter} PI units have failed, but there are no failures at LC_{UA} .

The transitions between states are governed by the failure rates of the LC units, the EIB, and the bus controllers. The states are divided into two zones, where the states in Zone- LC_{inter} signify failures that may occur at LC_{inter} 's before a failure at LC_{UA} (or at the EIB), while the states in Zone- LC_{UA} depict the failures that occur after a failure at LC_{UA} (or at the EIB). A transition from Zone- LC_{inter} to Zone- LC_{UA} arises when a failure occurs at LC_{UA} or at the EIB. Given State (i, j) (where $0 \leq i \leq N-3$ and $0 \leq j \leq M-2$), if the PI units (or the PDLU) of LC_{UA} fail(s) the LC moves to State i_{PI} (or State i_{PD}). Transitions between individual states in Zone- LC_{UA} (or Zone- LC_{inter}) are governed by the failures of LC_{inter} components. All states (except F) move to State T' if the EIB or LC_{UA} 's bus controller fails.




Figure 5. Markov models for reliability analysis of (a) BDR, (b) DRA. (Note: For clarity, not all transitions are shown.)

In Figure 5(b), $\pi_{i_PI}(t)$, $\pi_{j_PD}(t)$, $\pi_{i,j}(t)$ (with, $0 \leq i \leq N-3$ and $0 \leq j \leq M-2$), $\pi_T(t)$, and $\pi_F(t)$ are the system's probabilities to be in States $i_PI, j_PD, (i, j), T'$, and F , respectively. All states, except State F , are referred to as operational states. The individual state probabilities were obtained by solving the markov models. The reliability, $R(t)$, is the probability of the LC to be in any of the operational states, depicted in Figure 6 as a function of time (in hours) for varying values of M and N . Also shown in the plot is the reliability of BDR obtained from the Markov model of Figure 5(a).




Figure 6. Reliability of an LC under (a) BDR, (b) DRA.

The reliability values are obtained by first fixing $M=2$ and varying $3 \leq N \leq 9$, and second by fixing $N=9$ and varying $4 \leq M \leq 8$. Due to relatively lower failure rates of the EIB and the bus controller, an LC is more likely to enter State F caused by failure of all intermediate PI units (or PDLU's) than by the failed EIB or the failed bus controller. As seen from the plot above, the number of PI units has a greater impact on $R(t)$ than the number of PDLU's, where the values of $R(t)$ for $M > 4$ are very close to each other. By splitting the protocol dependent functionality into a separate unit with a relatively lower failure rate, DRA enables a larger number of PI units to cover failures at LC_{UA} 's PI units. As seen from Figure 6, significant improvement in $R(t)$ is achieved for DRA, where the reliability for $N=9$ (and $M \geq 4$) remains close to 1.0 for the first 40,000 hours; this is in sharp contrast to BDR whose reliability drops down to less than 0.5. Even for $M=2$ and $N=3$, DRA offers reasonably large improvement in reliability over a comparable BDR, indicating that a single covering LC is sufficient for improving reliability considerably. However, gains in $R(t)$ tend to shrink over successively increasing values of M and N . Therefore, DRA needs only a few LC_{inter} 's to significantly improve its reliability. In practical scenarios, it is extremely unlikely that multiple LC's would go down without any repair/replacement being undertaken. Hence, an LC's availability measure which includes a repair process seems to be more appropriate and is presented next.

5.2. Availability

For BDR, a failure at any of its LC component leads to LC_{UA} unable to route packets, rendering it unavailable. On the other hand, LC_{UA} is available under DRA, provided that

the bus controller and the EIB are operational and there is a healthy LC to cover it. To evaluate availability, we modify the Markov models of Figure 5 to include a repair process, which brings the system back to state $(0, 0)$ from any state i (with, $i \neq (0, 0)$). The repair process involves the replacement of the failed units with healthy ones, and it is assumed to take a fixed amount of time, irrespective of the type and the number of such units.

The availability of LC_{UA} is the probability that it is in any of the operational states (not in State F) at any given time t . We assume the same LC and EIB failure rates given earlier, and a repair rate of $\mu = 1/3$ (or $\mu = 1/12$) indicates that it may take three hours (or half a day) to repair/replace one or multiple faulty units. Figure 7 gives values of availability (A) obtained for BDR and DRA under various values of M and N , where $9^a x$ denotes x consecutive 9s after the decimal point (e.g., $9^a 4 = 0.9999$).




Figure 7. Availability of (a) BDR, (b) DRA.

The availability of DRA increases with increases in M and N , but it saturates at $9^a 9$ (or $9^a 8$) with $\mu = 1/3$ (or $\mu = 1/12$) for all $M \geq 4$. This is because the probability of LC_{UA} entering a state with more than three LC's failed and without the repair process being completed for the prior failed states, is minimal. As can be observed from Figure 7, a single covering LC_{inter} ($M=2, N=3$) gives an availability figure of $9^a 8$ for $\mu = 1/3$ (or $9^a 7$ for $\mu = 1/12$), which is significantly higher than $9^a 4$ (or $9^a 3$) for BDR. Therefore, akin to the reliability analysis, DRA does not require many LC_{inter} 's to give significant improvements in LC availability.

5.3. Performance Degradation in DRA

LC's in a router are hot-swappable, and a system restart is usually not needed to replace a faulty LC; hence, the performance of healthy LC's is not affected. Unlike BDR whose performance drops down to zero if LC_{UA} fails, in a DRA-based router non-faulty LC's are able to cover LC_{UA} (possibly with some performance degradation). In this sub-section, we analyze the performance of DRA under multiple failures and identify the extent of and conditions under which performance degradation occurs. Bandwidth available to a faulty LC at any given time (under one or multiple failures) is a performance measure of DRA with faults [9]. Our analysis utilizes the fact that the average link utilization of a router is usually less than 15% [3], and it makes the following assumptions and definitions.

- The variations in M and N make the performance analysis exceedingly complex. Hence, we do not treat the PI units and PDLU separately, but evaluate the performance of LC_{UA} as a single unit where any of the $N - 2$ LC's may cover LC_{UA} (irrespective of the fault location). A failure at LC_{inter} does not allow it to cover a faulty LC_{UA} unit, even if the corresponding LC_{inter} unit is functional. This assumption gives us the lower bound on performance (for $M = N$), since an LC_{inter} now has to be fully functional to be able to cover LC_{UA} . We assume a router consisting of N LC's each with a maximum capacity $c_{LC} = 10$ Gbps.
- Uniform traffic loads, $0.15 \leq L \leq 0.7$, are assumed at all LC's. These loads are based on the average link utilizations of the LC ports, which are assumed to vary between 15% and 70% [3].
- $\psi = c_{LC} - (L \times c_{LC})$, is the maximum bandwidth offered by a non-faulty LC to faulty LC's.
- $X_{nonfaulty}$ = number of non-faulty covering LC's, and X_{faulty} = number of faulty LC's ($X_{nonfaulty} + X_{faulty} = N$, where LC_{out} is assumed to be fault-free).

For each new failure, the aggregate bandwidth capacity of the router decreases by c_{LC} . We define B_{faulty} as the bandwidth available to a faulty LC at any time, under zero or more LC failures. B_{faulty} is the performance measure of LC's with faults. In case of multiple failures, $\sum B_{faulty}$ cannot exceed the bandwidth capacity of the EIB (B_{BUS}).

Figure 8 shows variations in B_{faulty} for increasing values of X_{faulty} , under $N = 6$ and with various LC loads (L). B_{faulty} is normalized to L and expressed as a percentage. The y -axis hence indicates the percentage of the required bandwidth available to a faulty LC, for an increasing number of LC failures, under varying load conditions.




Figure 8. Performance degradation for $N = 6$.

As can be observed, for $L = 15\%$ (which is the current average link utilizations of routers) DRA does not suffer from any performance degradation and is able to completely support up to $N - 1$ faulty LC's at the required capacity (for $N \leq 6$). Performance starts to degrade as L increases for a bigger N . In the worst case, for $X_{faulty} = 5$ and a load of 70%, less than 10% of the required capacity is available to faulty LC's. On the other hand, for lower loads, DRA is able to maintain a healthy performance level.

However, it is unlikely that more than two LC's would fail without a replacement procedure being completed.

A larger N results in higher values for B_{faulty} as long as the number of failed LC's (X_{faulty}) is small. As X_{faulty} approaches N , the reduced available capacity has to be shared amongst a larger number of faulty LC's, resulting in lower performance levels. Compared with BDR, DRA offers significantly better graceful fault-tolerant degradation.

6. Conclusion

We have presented Dependable Router Architecture (DRA) for providing fault-tolerance in high-performance routers. Current routers lack the dependability level demanded by today's mission-critical networks, due to absence of fault-tolerance in their linecards. DRA offers an efficient scheme to significantly increase the dependability (availability, reliability) and performance (under failure conditions) of current high-performance routers. The DRA design can also be applied to large-scale metro switches, which have a router-like LC-based architecture. The dependability of DRA was analyzed using Markov chains, and it was found to be significantly higher than that of a corresponding BDR. DRA's performance under single and multiple failures was also analyzed and compared favorably. A simple and highly dependable configuration indeed results from DRA, readily applicable to high-performance routers with multiple LC's.

References

- [1] Cisco Systems, Inc., *Cisco 12016 Gigabit Switch Router, Data Sheet*. 1999, URL - <http://www.cisco.com>.
- [2] C. Semeria, *T-series Routing Platforms: System and Packet Forwarding Architecture*, white paper, April 2002, URL - <http://www.juniper.net>.
- [3] A. Odlyzko, "The Current State and Likely Evolution of the Internet," *Proc. Globecom '99*, Dec. 1999, pp. 1869-1875.
- [4] Y. Luo, L. Bhuyan, and X. Chen, "Shared Memory Multiprocessor Architectures for Software IP Routers," *IEEE Trans. Parallel and Distributed Systems*, vol. 14, pp. 1240-1249, December 2003.
- [5] H. Chan, H. Alnuweiri, and V. Leung, "A Framework for Optimizing the Cost and Performance of Next-Generation IP Routers," *IEEE J. Selected Areas in Communications*, vol. 17, pp. 1013-1029, June 1999.
- [6] Cisco Systems, Inc., *Cisco 12000 Series Internet Router Architecture: Maintenance Bus, Power Supplies and Blowers and Alarm Cards*, Tech. Notes, URL - <http://www.cisco.com>.
- [7] PMC-Sierra, Inc., *Designing Multi-Gigabit Serial Backplanes with High Speed SERDES Solutions*, white paper, Nov. 2002, URL - <http://www.pmc-sierra.com/pressRoom/whitePapers.html>.
- [8] Cisco Systems, Inc., *Cisco 7300 1-Port OC-48c/STM-16 Packet over SONET/SDH Linecard, Data Sheet*. 1999, URL - <http://www.cisco.com>.
- [9] C. Das and L. Bhuyan, "Bandwidth Availability of Multiple-Bus Multiprocessors," *IEEE Trans. Computers*, vol. c-34, Oct. 1985, pp. 918-926.