Optimal Routing in the Worst-Case-Error Metric

Edwin Soedarmadji
California Institute of Technology, Pasadena, CA 91125, USA,
edwin@systems.caltech.edu

***Abstract*— This paper considers the problem of finding the path with minimum (zero) worst possible number of errors in a network with V nodes where (1) some nodes are capable of correcting up to a maximum number of x_{max} errors, (2) the nodes are connected by q-ary Symmetric Channels e_i parametrized by their Bit Error Ratios (BER) p_i. We introduce (1) the BER and Worst-Case Error (WCE) metrics and (2) an algebra that allows us to compute the path BER length from its edge lengths, and use them to measure network QoS. The WCE and BER metrics can be used with a generalized Dijkstra's Algorithm to compute the path of minimum WCE length. Finally, we present an algorithm that solves the above problem in the worst-case time complexity of $O(V^3)$.**

I. Introduction

In mission critical systems, the penalty for *any* error could be extremely high. In such systems, improving the network Quality of Service (QoS) involves comparing network paths (and edges) and selecting the optimal path that minimizes the worst possible number of errors (Worst-Case Error, or WCE) received — not a trivial task, especially in wireless, sensor and ad-hoc networks where the Bit Error Ratio (BER) is really high.

This paper introduces a new QoS metric, which we call the WCE metric. With this metric, network path "lengths" (the WCE lengths) are measured and compared to select the optimal path. A path's WCE length depends on its edges' WCE lengths. If each edge represents a q-ary Symmetric Channel (q-SC), then its WCE length can be determined from its BER. Thus, the BER can be used to measure edge and path lengths (the BER lengths), and is a metric — the BER metric. We prove that a path's WCE length is a non-decreasing function of its BER length, and hence, the minimum BER path is also the minimum WCE path. To calculate a path's BER length from its edges' BER lengths, we introduce a new BER algebra and prove that it can be used with the Generalized Dijkstra's Algorithm [1] to compute the shortest BER path.

We show that if *some* network nodes are capable of correcting edge errors, then it is possible to find a path (or paths) of zero WCE length. In wireless and ad-hoc networks where compute power and energy are constrained, available resource for complex mathematical operations used by high-performance FEC (such as finite fields arithmetic and iterative algorithms) is limited. It is therefore desirable to have only as many such "overhead" nodes as needed for reliability. In addition, compared to repeater nodes, error-correcting nodes incur delay and consume bandwidth resources. In this paper, we present a $O(V^3)$ algorithm that routes information with zero WCE through the appropriate error-correcting nodes.

This work was supported by the Caltech Lee Center for Advanced Networking and NSF Grant No. CCF-0514881

How relevant are error-correcting nodes in modern networks? For long trailing the Automated Repeat reQuest (ARQ) as the method of choice for improving network reliability, Forward Error Correction (FEC) methods have recently been shown to improve the QoS of wireless multimedia applications [2][3][4][5] beyond what could be achieved using only ARQ.

Unlike TCP and ARQ, FEC does not use return requests and thus consumes less bandwidth [6], especially in large multicast networks [7][8]. Even in peer-to-peer networks, FEC methods that use strategically positioned erasure-correcting nodes proved to be superior compared to replication-based strategies [9]. Other multicast algorithms such as *Digital Fountain* [10] and *Bullet* [11] employ FEC-based source erasure codes. Most recently, a network-coding FEC algorithm has also been proposed [12]. In response to criticisms that FEC consumes bandwidth and computation to process its redundant data, adaptive [4][5] and hybrid (ARQ-FEC) QoS-driven algorithms have been proposed [13][14][15] to dynamically adjust FEC activity to network conditions.

To the best of our knowledge, the problem we pose in the abstract has not been answered in the literature. In addition, the BER metric, WCE metric, and the minimum WCE routing algorithm presented in this paper are new and applicable to many contemporary network problems. Although in this paper the WCE metric is used to calculate the zero WCE path, it can be used in other QoS network optimizations that involve measuring (or minimizing) the occurence probability of non-typical, but possibly highly catastrophic, events.

This paper is organized as follows: the notations are introduced in section II. The main results and theorems are given in section III. Finally, the algorithm and a theorem of correctness are presented in section IV.

II. Formulation and Notation

We model a network as a digraph $G = (V, E)$, where V, E, and $\mathbf{\Pi}$ are the node, edge and path sets of G, respectively. Q is the q-ary alphabet $\{0, \dots, q-1\}$. The nodes $s, d \in V$ are the source and destination nodes, and $\Pi \subset \mathbf{\Pi}$ is the set of all paths from s to d.

A path π whose nodes V_π are connected by E_π is denoted by $\langle v_0, \dots, v_J \rangle$, $\langle e_1, \dots, e_J \rangle$, or $\langle v_0, e_1, \dots, e_J, v_J \rangle$. Denote by $|\pi|_v$ (or $|\pi|_e$) the number of nodes (or edges) in π from v_i to v_{i+1}. For (non-) adjacent nodes, $\langle v_i, v_{i+1} \rangle$ denotes the edge (path) connecting the two nodes. A partial path π_j is $\langle v_0, \dots, v_j \rangle$, with $0 < j \leq J$, and a truncated path $\bar{\pi}_j$ is $\langle v_0, e_1, \dots, v_{j-1}, e_j \rangle$.

Denote a data block by $B \in Q^n$, with $B_l \in Q$, and $l = 1 \dots n$. B can also be defined as a block with n packets of m symbols each. Instead of an alphabet, Q is the set of possible packet states encoding QoS metrics such as delay, loss, jitter, etc. Using these generalizations, our symbol- and alphabet-based results can be applied to other packet- and state-based QoS problems.

Let B_i (and $b_i = B_{li}$) denote B (and B_l) as it departs from v_i; and let $\bar{B}_i$ (and $\bar{b}_i = \bar{B}_{li}$) denote B (and B_l) as it leaves e_i. Both v_i and e_i are part of a path $\langle v_0, e_1, \ldots, e_J, v_J \rangle$, along which B evolves as follows:

$$B_0 \xrightarrow{e_1} \bar{B}_1 \xrightarrow{v_1} B_1 \xrightarrow{e_2} \bar{B}_2 \xrightarrow{v_2} \cdots \xrightarrow{e_J} \bar{B}_J \xrightarrow{v_J} B_J$$

where v_i and e_i corresponds to the operators $\boldsymbol{v}_i, \boldsymbol{e}_i \in \mathcal{E} : Q^n \to Q^n$ given by $B_i = \boldsymbol{v}_i(\bar{B}_i)$ and $\bar{B}_{i+1} = \boldsymbol{e}_i(B_i)$. The operator for π is $\boldsymbol{\pi} = \boldsymbol{v}_J \circ \boldsymbol{e}_J \circ \cdots \circ \boldsymbol{e}_1 \circ \boldsymbol{v}_0$. For π_j, it is $\boldsymbol{\pi}_j = \boldsymbol{v}_j \circ \boldsymbol{e}_j \circ \cdots \circ \boldsymbol{e}_1 \circ \boldsymbol{v}_0$, and for $\bar{\pi}_j$, it is $\bar{\boldsymbol{\pi}}_j = \boldsymbol{e}_j \circ \boldsymbol{v}_{j-1} \cdots \circ \boldsymbol{e}_1 \circ \boldsymbol{v}_0$. Thus, $\boldsymbol{\pi}_j(B_0) = B_j$ and $\bar{\boldsymbol{\pi}}_j(B_0) = \bar{B}_j$. The number of errors in B_i relative to $B_{i'}$ is denoted by $X_{i,\,i'} = X(B_i, B_{i'}) = |\{ l \,|\, b_i \neq b_{i'} \}|$. Also, $X_{i,\bar{j}}$ is a shorthand for $X(B_i, \bar{B}_j)$.

The nodes $\{u_i \in U \subseteq V\}$ implement (possibly different) q-ary, error-correcting codes of length n that can correct up to x_{max} errors in B_i relative to B_0. By default, $d \in U$. For packet-based systems, U are network nodes that can restore up to x_{max} packets in unfavorable states back to favorable states. and for a partial path π_j:

$$\begin{aligned} X_{0,j} &= 0 && ,\ v_j \in U \text{ and } X_{0,\bar{j}} \leq x_{max} \\ &\geq X_{0,\bar{j}} && ,\ v_j \in U \text{ and } X_{0,\bar{j}} > x_{max} \\ &= X_{0,\bar{j}} && ,\ v_j \in V \backslash U \end{aligned} \quad (1)$$

If $v_j \in U$ and $\bar{B}_j$ does not have more than x_{max} errors, then v_j transforms $\bar{B}_j$ back to B_0. Otherwise, if x_{max} is exceeded, v_j may increase the number of error in B. If $v_j \notin U$, it simply repeats the content of $\bar{B}_j$ into B_j.

Each edge $e = \langle v_i, v_{i+1} \rangle$ is a q-ary Symmetric Channel with a BER of $p \in \mathcal{P} = [0,1] \cup \infty$, or q-SC($p$). The value $p = \infty$ means v_i is not connected to v_{i+1}.

For each symbol in B, the transition probability is:

$$P(\, b_{i+1} \,|\, b_i \,) = \begin{cases} 1 - p & ,\ b_i = b_{i+1} \\ p \,/(q-1) & ,\ b_i \neq b_{i+1}. \end{cases} \quad (2)$$

The probability of $B_{i+1} = \boldsymbol{e}_i(B_i)$ having x errors relative to B_i is $P(X_{i,\,i+1} = x) = P(x,p)$ given by:

$$P(x,p) = \begin{cases} \binom{n}{x} p^x \, (1-p)^{n-x} & ,\ p \in (0,1) \\ \delta(x) & ,\ p = 0 \\ \delta(x-n) & ,\ p = 1, \infty \end{cases} \quad (3)$$

The BER metric uses p, which determines $P(x,p)$, to measure the QoS of e. The WCE metric measures the worst "possible" (defined as $P(x,p)$ above a threshold ϵ) number of errors (the Worst-Case Error, or WCE) $\bar{x}$ in B_{i+1} relative to B_i, denoted by $\bar{x} \in \mathcal{X} = [0,n]$:

$$\bar{x}(p,\epsilon) = \max_x \{\, x \mid P(x,p) \geq \epsilon \,,\, 0 \leq x \leq n \,\} \quad (4)$$

From (3), $\bar{x}(1,\epsilon) = \bar{x}(\infty,\epsilon) = n$, $\bar{x}(0,\epsilon) = 0$, and $\bar{x}(p,\epsilon)$ is not continuous. The functions $\beta : \boldsymbol{\Pi} \to \mathcal{P}$ and $\omega : \boldsymbol{\Pi} \to \mathcal{X}$ measure the BER and WCE path lengths.

Consider a path $\pi = \langle e_1, \ldots, e_J \rangle \in \boldsymbol{\Pi}$ and its partial path $\pi_j = \langle e_1, \ldots, e_j \rangle$ with $1 \leq j \leq J$. Let $p_j = \beta(e_j)$ and $\bar{x}_j = \omega(e_j)$. The path QoS length p_π (or $\bar{x}_\pi$) depends on p_j (or $\bar{x}_j$), but in general, $p_\pi \neq \sum p_j$ and $\bar{x}_\pi \neq \sum \bar{x}_j$.

Denote the BER and WCE addition operators by $\oplus$. If $p_1 = \beta(e_1)$, $p_2 = \beta(e_2)$, $\bar{x}_1 = \omega(e_1)$, $\bar{x}_2 = \omega(e_2)$, and $\pi = \langle e_1, e_2 \rangle$, then $p_\pi = p_1 \oplus p_2$ (or $\bar{x}_\pi = \bar{x}_1 \oplus \bar{x}_2$). For q-SC($p$), the $\oplus$ operators are defined below.

$$\begin{aligned} p_1 \oplus p_2 &= 1 - (1-p_1)(1-p_2) - (p_1 p_2) \,/\, (q-1) \\ \bar{x}_1 \oplus \bar{x}_2 &= \max\{\, x \mid P(x, p_1 \oplus p_2) \geq \epsilon \,\} \end{aligned} \quad (5)$$

We can now define p_π and $\bar{x}_\pi$ in terms of p_j and $\bar{x}_j$ using a generalized summation: $p_\pi = \bigoplus p_e$ and $\bar{x}_\pi = \bigoplus \bar{x}_e$. The pairing of $\mathcal{P}$ and $\oplus$ forms an algebraic structure (a *magma*) which we call the BER algebra **B**. Likewise, the pairing of $\mathcal{X}$ and $\oplus$ forms the WCE algebra **W**.

Between two points, the BER (or WCE) optimal path π^* is the path with the "least" BER (or WCE) path length. To calculate π^*, we need to compare path lengths. Therefore, we need a total order $\preceq$ on $\mathcal{P}$ (or $\mathcal{X}$) to evaluate expressions like $p_\pi \preceq p_{\pi'}$ (or $\bar{x}_\pi \preceq \bar{x}_{\pi'}$).

$$\begin{aligned} \bar{x}^* &= \min_\pi \{\, \bar{x}_\pi = \bar{x}(p_\pi, \epsilon) \mid \pi \in \Pi \,\} \\ p^* &= \min_\pi \{\, p_\pi \mid \pi \in \Pi \,\} \end{aligned} \quad (6)$$

The lengths p_{π_j} and $\bar{x}_{\pi_j}$ are non-decreasing functions of j — an assertion we prove later in this paper. This means that on any path, the worst possible $X_{0,j}$ also increases as B traverses more edges. If G cannot correct any errors, reliable communication (as far as the worst-case scenario is concerned) becomes very difficult to achieve except for very small values of $p_j \ll 1$.

III. Main Results

The problem of finding the optimal path of minimum WCE in G bears many similarities to the Shortest Path Problem (SPP) that can be solved using Dijkstra's Algorithm (DA). The Generalized Dijkstra's Algorithm (GDA) below [1] is the perfect tool for the problem:

Algorithm 1 (GDA): The GDA is practically identical to DA except for the relaxation step, where $\oplus$ and $\preceq$ operators act on a general metric space $\mathcal{M}$ (instead of the step in DA, where $+$ and $\leq$ operators act on $\mathbb{R}$) :

```
procedure DIJKSTRA (G, p, s)
    for all v ∈ V do
        l[v] ← ∞
        π[v] ← NIL
    Q ← V
    l[s] ← 0
    while Q ≠ ∅ do
        u ← MIN(Q)
        for all node v ∈ N(u) do
            if l[v] ≻ l[u] ⊕ p(u, v) then
                l[v] ← l[u] ⊕ p(u, v)
                π[v] ← u
```

On line 9, $N(u)$ denotes the set of all nodes adjacent to u. The argument $\boldsymbol{p}$ is the BER lengths of the edges in G, each of which is an element in $\mathcal{M}$, and $\boldsymbol{p}\,(u,v)$ is the BER of $\langle u,v\rangle$. Lines 10–12 perform the relaxation step of the GDA. This step depends on the definitions of $\mathcal{M}$, $\oplus$, and $\preceq$. If the definitions are such that the GDA (in)correctly returns the path in G with minimum length measured in $\mathcal{M}$, then $(\mathcal{M},\oplus)$ and $\preceq$ are said to be *(in)compatible* with the GDA. □

Proposition 1: An algebra $\mathbf{A}=(\mathcal{M},\oplus)$ and a total order $\preceq$ is compatible with the GDA if and only if it satisfies *all* the properties in the set denoted by **P** below:

P1 is a commutative *monoid*, that is, for $a,b,c\in\mathcal{M}$:
- $\mathcal{M}$ is *closed* under $\oplus$: $a\oplus b\in\mathcal{M}$;
- $\oplus$ is *associative* : $a\oplus(b\oplus c)=(a\oplus b)\oplus c$;
- 0 is the *identity* : $a\oplus 0=0\oplus a=a$;
- $\oplus$ is *commutative* : $a\oplus b=b\oplus a$.

P2 There exists $\infty\in\mathcal{M}\mid a\oplus\infty=\infty\oplus a=\infty$.
P3 $\preceq$ is a *total order* on $\mathcal{M}$, i.e., $\preceq$ is :
- *reflexive*: $a\preceq a$;
- *anti-symmetric*: if $a\preceq b$ and $b\preceq a$ then $a=b$;
- *transitive*: if $a\preceq b$ and $b\preceq c$ then $a\preceq c$;
- *total*: for every $a,b\in\mathcal{M}$ either $a\preceq b$ or $b\preceq a$.

P4 There exists the least element 0 that satisfies $0\preceq a$.
P5 $a\oplus c\prec b\oplus c$ if $a\prec b$ and $c\in\mathcal{M}-\{\infty\}$.

PROOF: Refer to [1] for a complete proof. □

Theorem 2: The algebra $\mathbf{B}=(\mathcal{P},\oplus)$ and the total order $\preceq$ defined in section II satisfy all the properties in **P**, and thus compatible with the GDA.

PROOF: Recall the definition of $\oplus$ from equation (5), which is repeated here for convenience:

$$a\oplus b=1-(1-a)(1-b)-(ab)\,/\,(q-1) \tag{5}$$

P1 Except for closure, the monoid properties can be proven with algebraic manipulation of equation (5). Recall that $\mathcal{P}=[0,1]\cup\infty$. Closure is obvious if $a,b\in[0,1]$. If $a=\infty$ or $b=\infty$, then by the definition of ∞ as a special designation for the absence of connection between two nodes, then we must have $a\oplus b=\infty$.
P2 The proof is derived from closure on ∞.
P3 The proof follows the definition of $\mathcal{P}$.
P4 Same as above.
P5 The proof is obvious if $b=\infty$. However, if $b\neq\infty$, then by substituting $a,b,$ and c into (5) we obtain the two values $a\oplus c$ and $b\oplus c\in[0,1]$ given by:

$$a\oplus c=1-(1-a)\,(1-c)-a\,c/\,(q-1)$$
$$b\oplus c=1-(1-b)\,(1-c)-b\,c/\,(q-1).$$

Since both are in $[0,1]$, the order $\prec$ is just $<$, and the expression $a\oplus c\prec b\oplus c$ is equivalent to the inequality $(a\oplus c)-(b\oplus c)<0$, which can be simplified into:

$$(1-c\,q/\,(q-1))\,(a-b)<0$$

If $c\leq\frac{q-1}{q}$, then $0<(1-c\,q/\,(q-1))\leq 1$, and $(a-b)<0$ (because $a\prec b$ and $b\neq\infty$). Thus, the above inequality is true, and we have proven all the properties in **P**. □

Given two nodes $v_0,v_J\in V$ and the set of all paths Π_J connecting them, the path length of $\pi\in\Pi_J$ is given by $p_\pi=\beta(\pi)$. Using the GDA with $(\mathbf{B},\preceq)$ can compute the *shortest* path $\pi^*\in\Pi_J$ where $p_{\pi^*}\preceq p_\pi,\ \forall\pi\in\Pi_J$. Again, if π_j denotes the partial path of $\pi\in\Pi_J$,

Proposition 3: If $V_{\pi_J}\cap U=\varnothing$, then $\beta(\pi_j)$ is a non-decreasing function of j. The minimum $\beta(\pi_j)=0$ is possible if the edge lengths $p_j=0$ for all $j=0\ldots J$.

PROOF: **P5** with $0=a\prec b=p_j$, and $c=p_{\pi_{j-1}}$ gives us $\beta(\pi_{j-1})=p_{\pi_{j-1}}\prec p_{\pi_{j-1}}\oplus\ p_j=\beta(\pi_j)$, proving that $\beta(\pi_j)$ is a non-decreasing function of j. The second part of the proof can be derived directly from **P1**. □

The preceding proposition shows the practical benefit of including error-correcting nodes in the network, and of routing information through them. If the path includes $u\in U$, then for some $j\in[0,J]$, we can have $\beta(\pi_j)=0$ even with $p_j\neq 0$ for all j. Next, we prove that within an admissible range of $\epsilon\in[0,\bar{\epsilon}]$, the WCE function $\bar{x}(p,\epsilon)$ is a non-decreasing (albeit discontinuous) function of p.

Lemma 4: For a given n and a fixed x, the probability function $P(x,p)$ is maximized at $p=\frac{x}{n}$. Furthermore, $P_*=P(\lfloor\frac{n}{2}\rfloor,\frac{1}{2})$ minimizes $P(x,\frac{x}{n})$ over all $x\in\mathcal{X}$.

PROOF: From the definition of $P(x,p)$ in equation (3), the lemma is true for $x=0$ and $x=n$. Consider :

$$\tfrac{\partial P(x,p)}{\partial p}=P(x,p)\,(\tfrac{x}{p}-\tfrac{n-x}{1-p})=0$$

Solving $(\frac{x}{p}-\frac{n-x}{1-p})=0$ for p gives us $p=\frac{x}{n}$, which maximizes the function $P(x,p)$ for $x\in(0,1)$.

The next question is, for $0<x<n$, which x minimizes $P(x,\frac{x}{n})$? Unlike with p, we cannot differentiate $P(x,p)$ with respect to x because it is a discrete variable. Instead of approximating $P(x,p)$ with a Gaussian distribution, which is valid only for certain n and p, we resort to the upper and lower bounds for $\binom{n}{x}$ in [16] :

$$\xi(A)=\frac{e^A\ n^{n+\frac{1}{2}}}{(2\pi)^{\frac{1}{2}}\ (n-x)^{n-x+\frac{1}{2}}\ x^{x+\frac{1}{2}}} \tag{7}$$
$$\lambda_{nx}=\lambda(n,x)=\xi\big(\tfrac{1}{12n}-\tfrac{1}{12x}-\tfrac{1}{12(n-x)}\big)$$
$$\mu_{nx}=\mu(n,x)=\xi\big(\tfrac{1}{12n}-\tfrac{1}{12x+1}-\tfrac{1}{12(n-x)+1}\big)$$
$$\lambda_{nx}<\tbinom{n}{x}<\mu_{nx}$$

Using (7), $P(x,p)$ is now bounded by two continuous and differentiable functions in x.

$$\lambda_{nx}\,p^x\,(1-p)^{n-x}<P(x,p\,)<\mu_{nx}\,p^x\,(1-p)^{n-x}$$

Since $p=\frac{x}{n}$, we substitute $x=np$ into the equation below and solve the p roots of the p derivatives of both the lower and upper bounds of $P(x,p\,)$ to find the minima with respect to p.

The lower and upper bounds are minimized at $p = \frac{1}{2}$. Since $p = \frac{x}{n}$, then $x = \frac{n}{2}$:

$$\sqrt{\frac{2}{n\pi}}\, e^{-\frac{18n-1}{12n(6n+1)}} < P(\lfloor \tfrac{n}{2} \rfloor, \tfrac{1}{2}) < \sqrt{\frac{2}{n\pi}}\, e^{-\frac{1}{4n}} \qquad (8)$$

in which for large values of n the lower and upper bounds converge. In fact, in (7), λ_{nx} converge to μ_{nx} for all x, including at the discrete points $0 < x < n \in \mathbb{N}$. Thus the minima for λ_{nx} and μ_{nx} over the continuous x must also be the minimum for $P(x, x/n)$ over the discrete x, where it is denoted by $P_* = P(\lfloor \frac{n}{2} \rfloor, \frac{1}{2})$ □

A function $f(x)$ is *unimodal* over $x \in [a, b]$ if there exists an x_0 such that $f(x)$ is monotonically increasing for $x < x_0$ and monotonically decreasing for $x > x_0$.

Lemma 5: $P(x, p)$ is unimodal over x and p.

PROOF: To prove unimodality over x, solve $P(x, p) < P(x+1, p)$ for x, giving $x < np-(1-p)$ or $x < x_0 = \lfloor np \rfloor$. Similarly, $P(x, p) > P(x+1, p)$ or $x > x_0$. Unimodality over p is proven with calculus:

$$P'(x, p) = \frac{\partial}{\partial x} P(x, p) = P(x, p)(\tfrac{x}{p} - \tfrac{n-x}{1-p})$$

$P'(x, p_0) = 0$ at $p_0 = \frac{x}{n}$. For $p < p_0$, we have $P'(x, p) > 0$ and for $p > p_0$, we have $P'(x, p) < 0$. □

Corrolary 6: The set $\boldsymbol{\epsilon}$ of admissible ϵ is $[0, \bar{\epsilon} = P_*]$.

PROOF: Recall that $\bar{x}(p, \epsilon) = \max\{\, x \mid P(x, p) \geq \epsilon \}$. To ensure $\bar{x}(p, \epsilon)$ is valid for all $p \in \mathcal{P}$, for each value of p we must have $P(x, p) \geq \epsilon$ for some x. This condition is trivially met if $\epsilon \leq 0$ because $P(x, p) \geq 0$. If $\epsilon > P_*$, then $\{\, x \mid P(x, \frac{1}{2}) \geq \epsilon \} = \varnothing$, and $\bar{x}(p, \frac{1}{2})$ is invalid. Thus, the set of admissible ϵ is given by $[0, P_*]$. □

Denote by $p_i(x, \epsilon) = p_i(x) = \{\, p \in \mathcal{P} \mid P(x, p) = \epsilon \}$ the i-th roots of $P(x, p) = \epsilon$ given n and ϵ. From lemma 5, if $\epsilon \in \boldsymbol{\epsilon}$, then there is at least one root. At $x = 0$, the single root is $p_0(0) = 0$ and at $x = n$, it is $p_1(n) = 1$. Except for another special case when $\epsilon = P_*$, where at the midpoint $x_m = \lfloor \frac{n}{2} \rfloor$ there is only one root to $p_0(x_m) = p_1(x_m) = \frac{1}{2}$, in general, there are two distinct roots $p_0(x), p_1(x) \in [0, 1]$, with $p_0(x) < \frac{x}{n} < p_1(x)$. If ϵ goes toward 0, then $p_0(x) \to 0$ and $p_1(x) \to 1$, except $p_1(0) = 0$ and $p_0(1) = 1$. If ϵ goes toward P_*, then $p_0(x)$ increases, while $p_1(x)$ decreases.

Define the sets P_0 and P_1, each having the $n + 1$ values of $p_0(x)$ and $p_1(x)$, for $0 \leq x \leq n$. These values are also referred to as $x_0(p)$ and $x_1(p)$, for $0 \leq p = \frac{x}{n} \leq n$.

Lemma 7: The roots $p_0(x)$ and $p_1(x)$ are non-decreasing functions of x with $p_0(x) = p_1(x)$ only at $x = 0$ and $x = n$ (or $x = x_m$ for $\epsilon = P_*$).

PROOF: First, observe that $P(x, p) = P(x + 1, p)$ only has one root at $p = p^\times = \frac{x+1}{n+1}$ between the maxima of $P(x, p)$ and $P(x + 1, p)$, i.e., $\frac{x}{n} < p^\times < \frac{x+1}{n}$. From lemma 5, this implies that if $p < p^\times$ then $P(x, p) > P(x + 1, p)$, and if $p > p^\times$, $P(x, p) < P(x + 1, p)$. Hence, $p_0(x) \leq p_0(x + 1)$ and $p_1(x) \leq p_1(x + 1)$, that is, $p_0(x)$ and $p_1(x)$ are both non-decreasing functions of x. □

Theorem 8: The values $\bar{x}(p, \epsilon) = \max_x x_0(p)$ and $\max_x x_1(p)$ are non-decreasing functions of p.

PROOF: First, the $\max_x$ function is used because it is possible to have $p_0(x) = p_0(x') \in P_0$ and $x \neq x'$. For example, if $\epsilon = 0$, $p_0(0) = \cdots = p_0(n - 1) = 0$ (the same argument applies to P_1). Just as $\max_x$ in $\bar{x}(p, \epsilon)$ isolates the largest x satisfying $P(x, p) \geq \epsilon$, the function $\max_x x_0(p)$ isolates the largest x satisfying $P(x, p) = \epsilon$. After establishing a unique x for each p, the proof follows from the monotonicity of $p(x)$. □

Theorem 8 proves that for any path π, the WCE length $\omega(\pi)$ is a non-decreasing function of its BER $\beta(\pi)$. Therefore, the minimum BER path between two nodes (computed by the GDA) is also the minimum WCE path. Therefore, from this point on, the terms "minimum BER" and "minimum WCE" are interchangeable.

In the next section, we finally present the algorithm MIN-WCE-PATH that computes the minimum (zero) WCE path in the presence of error-correcting nodes U.

IV. ALGORITHM

A path $\phi = \langle v_1, v_2 \rangle$ is feasible iff $\omega(\langle s, v_1 \rangle) \oplus \omega(\phi) \leq x_{max}$ — the WCE at v_1 plus the WCE length of ϕ must be less than x_{max}. Denote by $\Phi \subseteq \mathbf{\Pi}$ the feasible paths in $\mathbf{\Pi}$, and by $\Phi(v_1, v_2)$ the feasible paths between v_1, v_2.

Theorem 9: A path π^* is the minimum WCE path iff it solves the SPP given by $G' = (V', E')$, where $V' = \{s\} \cup U$. An edge connecting two nodes $v_1, v_2 \in V'$ represents the shortest path in $\Phi(v_1, v_2)$. Therefore

$$E' = \{\, \mathrm{argmin}_\phi \{\, \omega(\phi) \mid \phi \in \Phi(v_1, v_2) \,\} \,\big|\, v_1, v_2 \in V' \,\}.$$

PROOF: Suppose π^* contains $n+1$ segments ϕ_i connecting the nodes in $V'' = \{s, U_{\pi^*}, d\}$, where $U_{\pi^*} = U \cap V_{\pi^*}$. In segment notation, π^* is denoted by $s \rightsquigarrow u_1 \rightsquigarrow \cdots \rightsquigarrow u_j \rightsquigarrow d$, with $\{u_i\} = U_{\pi^*}$, and $0 \leq j \leq |U|$.

Then ϕ_i must be the shortest feasible paths between adjacent nodes in V'', and U_{π^*} must be the set that minimizes $\sum \beta(\phi_i)$. Otherwise, a better path ξ^* can be obtained by modifying ϕ_i or U_{π^*}, contradicting the claim that π^* is optimal. For the forward proof, note that $V'' = \{s, d, U_{\pi^*}\} \subseteq V'$. Further, since each ϕ_i is a shortest path between nodes in V', then $\phi_i \in E'$. Therefore π^* solves the SPP given by (V', E').

For the reverse proof, suppose ξ^* is the SPP solution but is not the minimum WCE path π^*. From the forward proof, if π^* minimizes WCE, then it also solves the SPP, thus $\omega(\pi^*) \leq \omega(\xi^*)$. However, if $\omega(\pi^*) < \omega(\xi^*)$, then ξ^* is not the SPP solution – a contradiction.

Thus, $\omega(\pi^*) = \omega(\xi^*)$, and if path lengths are unique, $\pi^* = \xi^*$. Hence, the SPP solution π^* is the minimum WCE path. □

Hence, to find the minimum WCE path we first compute the minimum WCE path for each pair of nodes in V'. These minimum paths are then converted into edges in E', connecting the nodes in V'. The overall minimum WCE path is then computed from these edges using the GDA.

Algorithm 2: Theorem 9 proves the correctness of the optimization algorithm listed below :

```
1: procedure MIN-WCE-PATH (G , p , s)
2:     E \= { e ∈ E | ω(e) > x_max }
3:     V \= { v ∈ V | deg(v) = 0 }
4:     for v_1 ∈ V′ do
5:         SP_1 = GDA ( G , p , v_1 )
6:         for v_2 ∈ V′ ≠ v_1 do
7:             E′ = E′ ∪ ⟨v_1, v_2⟩
8:     E′ \= { e′ = (v_1, v_2) ∈ E′ | e′ ∉ Φ(v_1, v_2) }
9:     V′ \= { v′ ∈ V′ | deg(v′) = 0 }
10:    SP_2 = GDA ( G′ , p , s )
```

On lines 2 and 3, the algorithm prunes all the infeasible edges. Then, on line 5, it runs the GDA on all $v_1 \in V'$, every time producing a shortest path tree SP_1 rooted at v_1. On line 7, the edges connecting v_1 and $v_2 \in V'$ are added into E' based on SP_1. The first stage is finished and the second stage begins. On lines 8 and 9, infeasible edges in E' are pruned and any isolated nodes in V' are removed. Finally, π^* is obtained. □

Denoting $|V'|$ by α, the first stage produces a complete graph with α nodes and $\alpha(\alpha-1)$ directed edges by executing the GDA α times on line 5, and thus has a time complexity of $O(\alpha V^2)$ (if the GDA is implemented using Fibonacci heap, then its complexity could reach $O(V \log V + E)$ [17]). Lines 8 and 9 search linearly over them with $O(\alpha^2)$ time complexity. The GDA on line 10 has a time complexity $O(\alpha^2)$. Hence, overall time complexity is $O(\alpha V^2 + \alpha^2) < O(V^3)$.

V. Conclusion

In this paper, we considered the problem of finding the path with minimum (or zero) worst possible number of errors in mission-critical communication networks. We made the assumptions that some network nodes are capable of correcting up to a maximum number of x_{max} errors in a block of n symbols, and that the nodes are connected by q-ary Symmetric Channels parametrized by the Bit Error Ratios (BER).

We introduced (1) the BER and Worst-Case Error (WCE) metrics and (2) a BER algebra that allows for computation of the path BER length from its edge lengths. The metric-algebra pair was then used to optimize the network QoS in the BER and WCE metrics using a generalized Dijkstra's Algorithm in the worst-case time complexity of $O(V^3)$, where V is the number of nodes in the network.

Future research can explore the issues of whether the approach outlined in this paper can be generalized to other types of distributions and scenarios, the incorporation of the algorithm presented here into existing networking protocols, as well as simulation and experimental verification of the algorithm on standard network models.

Acknowledgment

The author would like to acknowledge Prof. Robert J. McEliece for valuable discussions and guidance, as well as useful suggestions that have improved this manuscript, and the anonymous referees for their insightful comments.

References

[1] Sobrinho, J.L., "Algebra and algorithms for QoS path computation and hop-by-hop routing in the Internet," *IEEE/ACM Transaction on Networking.*, vol. 10, pp. 541–550, August 2002.

[2] Kurceren, B. et al., "A joint source-channel coding approach to network transport on digital video," *Proceeding of INFOCOM 2000*, vol. 2, pp. 717–726, 2000.

[3] Michelson, A.M. et al., "Reliable ATM transmission of multimedia in wireless networks," *Universal Personal Communications, 1998*, vol. 2, pp. 1021–1027, 1998.

[4] Nikaein, N. et al., "MA-FEC: a QoS-based adaptive FEC for multicast communication in wireless networks," *Proceeding of IEEE ICC, 2000*, vol. 2, pp. 954-958, 2000.

[5] Wang G.J. et al., "Channel-adaptive error protection for scalable video over channels with bit errors and packet erasures," *Proceeding of ISCAS 2002*, vol. 2, 712–715, 2002.

[6] Barsocchi, P. et al., "Experimental Results with Forward Erasure Correction and Real Video Streaming in Hybrid Wireless Networks," *2nd International Symposium on Wireless Communication Systems, 2005*, pp. 662–666, 2005.

[7] Huitema, C., The case for packet level FEC, in *Proceedings of IFIP PfHSN 1996*.

[8] Nonnenmacher, J., et al., Parity-based loss recovery for reliable multicast transmission, *IEEE/ACM Transactions on Networking*, vol. 6, no. 4, pp. 349–361, 1998.

[9] Lancerica, L., et al. "Evaluation of content-access QoS for various dissemination strategies in peer to peer networks," *Proceeding of IEEE ICON 2003*, pp. 337–342, 2003.

[10] Byers, J., et al., Informed Content Delivery Across Adaptive Overlay Networks, *Proceedings of ACM SIGCOMM*, 2002.

[11] Kostic, D., et al., Bullet: High Bandwidth Data Dissemination Using an Overlay Mesh, *Proc. of ACM SOSP, 2003*.

[12] Wang, M., et al., A High-Throughput Overlay Multicast Infrastructure with Network Coding, *Proceedings of IFIP IWQOS 2005* pp. 37-53.

[13] Du Q., et al., "Adaptive Low-Complexity Erasure-Correcting Code-Based Protocols for QoS-Driven Mobile Multicast Services," *Proceeding of QShine 2005*, pp. 29–29, 2005.

[14] Li Q. et al., "Providing adaptive QoS to layered video over wireless local area networks ...," *IEEE Transactions on Multimedia*, vol. 6, no. 2 pp. 278–290, 2004.

[15] Li J., et al., "A practical proactive hybrid FEC/ARQ technique for reliable multi-session multicast communication," *Proceeding of IEEE ICECS 2003*, vol. 2, pp. 519–522, 2003.

[16] Stanica, P., "Good Lower and Upper Bounds on Binomial Coefficients," *Journal of Inequalities in Pure and Applied Mathematics.*, vol. 2, no. 3, article 30, 2001.

[17] Cormen, T.H.; Leiserson, C.E.; Rivest, R.L., *Introduction to Algorithms*, MIT Press, MA and McGraw-Hill, NY, 1990.