A Decision-Theoretic Approach to Reliable Message Delivery*

Francis C. Chu   Joseph Y. Halpern

Department of Computer Science
Upson Hall, Cornell University
Ithaca, NY 14853-7501, USA

{fcc,halpern}@cs.cornell.edu

To be, or not to be: that is the question:
Whether 'tis nobler in the mind to suffer
The slings and arrows of outrageous fortune,
Or to take arms against a sea of troubles,
And by opposing end them?

Hamlet (III, i)

Abstract

We argue that the tools of decision theory should be taken more seriously in the specification and analysis of systems. We illustrate this by considering a simple problem involving reliable communication, showing how considerations of utility and probability can be used to decide when it is worth sending heartbeat messages and, if they are sent, how often they should be sent.

Keywords: decision theory, specifications, design and analysis of distributed systems

1 Introduction

In designing and implementing systems, choices must always be made: When should we garbage collect? Which transactions should be aborted (to remove a deadlock)? How big should the page table be? How often should we resend a message that is not acknowledged? Currently, these decisions seem to be made based on intuition and experience. However, studies suggest that decisions made in this way are prone to inconsistencies and other pitfalls [RS89]. Just as we would like to formally verify critical programs in order to avoid bugs, we would like to apply formal methods when making important decisions in order to avoid making suboptimal decisions. Mathematical logic has given us the tools to verify programs, among other things. There are also standard mathematical tools for making decisions, which come from *decision theory* [Res87]. We believe that these tools need to be taken more seriously in systems design. We view this paper as a first step towards showing how this can be done and the benefits of so doing.

Before we delve into the technical details, let us consider a motivating example. Suppose Alice made an appointment with Bob and the two are supposed to meet at five. Alice shows up at five on the dot but Bob is nowhere in sight. At 5:20, Alice is getting restless. The question is "To stay or not to stay?" The answer, of course, is "It depends." Clearly, if Bob is an important business

*This work was supported in part by NSF grant IRI-96-25901 and by the Air Force Office of Scientific Research grant F49620-96-1-0323.

client and they are about to close a deal, she might be willing to wait longer. On the other hand, if Bob is an in-law she never liked, she might be happy to have an excuse to leave. At a more abstract level, the *utility* of actually having the meeting is (or, at least, should be) an important ingredient in Alice's calculations. But there is another important ingredient: likelihood. If Alice and Bob meet frequently, she may know something about how prompt he is. Does he typically arrive more or less on time (in which case the fact that he is twenty minutes late might indicate that he is unlikely to come at all) or is he someone who quite often shows up half an hour late? Not surprisingly, utilities and probabilities (as measures of likelihood) are the two key ingredients in decision theory.

While this example may seem far removed from computer systems, it can actually be viewed as capturing part of *atomic commitment* [SKS97]. To see this, suppose there is a coordinator p_c and two other processes p_a and p_b working on a transaction. To commit the transaction, the coordinator must get a YES vote from both p_a and p_b. Suppose the coordinator gets a YES from p_a, but hears nothing from p_b. Should it continue to wait or should it abort the transaction? The types of information we need to make this decision are precisely those considered in the Alice-Bob example above: probabilities and utilities. While it is obvious that the amount of time Alice should wait depends on the situation, atomic commit protocols typically have a context-independent timeout period. If p_c has not heard from all the processes by the end of the timeout period, then the transaction is aborted. Since the importance of the transaction and the cost of waiting are context-*dependent*, the timeout period would not be appropriate in every case.

Although it is not done in atomic commit protocols, there certainly is an awareness that we need to take utilities or costs into account elsewhere in the database literature.[1] For example, when a deadlock is detected in a database system, some transaction(s) must be rolled back to break the deadlock. How do we decide which ones? The textbook response [SKS97, p. 497] is that "[we] should roll back those transactions that will incur the minimum cost. Unfortunately, the term minimum cost is not a precise one." Typically, costs have been quantified in this context by considering things like how long the transaction has been running and how much longer it is likely to run, how many data items it has used, and how many transactions will be involved in a rollback. This is precisely the type of analysis to which the tools of decision theory can be applied. Ultimately we are interested in when each transaction of interest will complete its task. However, some transactions may be more important than others. Thus, ideally, we would like to attach a utility to each vector of completion times. Of course, we may be uncertain about the exact outcome (e.g., the exact running time of a transaction). This is one place where likelihood enters the picture. Thus, in general, we will need both probabilities and utilities to decide which are the most appropriate transactions to abort. Of course, obtaining the probabilities and utilities may in practice be difficult. Nevertheless, we may often be able to get reasonable estimates of them (see Section 6 for further discussion of this issue), and use them to guide our actions.

In this paper, we illustrate how decision theory can be used and some of the subtleties that arise in using it. We focus on one simple problem involving reliable communication. For ease of exposition, we make numerous simplifying assumption in our analysis. Despite these simplifying assumptions, we believe our results show that decision theory can be used in the specification and design of systems.

We are not the first to attempt to apply decision theory in computer science. Shenker and his colleagues [BBS98, BS98], for example, have used ideas from decision theory to analyze various

[1]Awareness of cost is by no means limited to the database community. For example, a sampling of the papers at a recent DISC (Distributed Computing) Conference, showed that cost was mentioned in at least seven of them [BMPP98, CM98, EHWG98, FMS98, MIB98, TRAR98, YAGW98]. Cost and utility are also discussed, for example, in [Kes97] and [KL95, LS98].

network protocols; Microsoft has a Decision Theory and Adaptive Systems group that has successfully used decision theory in a number of applications, including troubleshooting problems with printers and intelligent user interfaces in Office '97. (See `http://research.microsoft.com/dtas/` for further details.) Mikler et al. [MHW96] have looked at network routing from a utility-theoretic perspective. One important difference between our paper and theirs is that they do not treat the utility function as a given: Their aim is to find a good utility function so that the routing algorithm would exhibit the desired behavior (of avoiding the hot spot). More generally, our focus on writing specifications in terms of utility, and the subtleties involved with the particular application we consider here—reliable communication—make the thrust of this paper quite different from others in the literature.

The rest of this paper is organized as follows. We briefly review some decision-theoretic concepts in Section 2. In Section 3 we describe the basic model and introduce the communication problem that serves as our running example. We show that the expected cost of even a single attempt at reliable communication is infinite if there is uncertainty about process failures. We then show in Section 4 how we can achieve reliable communication with finite expected cost by augmenting our system with *heartbeat* messages, in the spirit of Aguilera, Chen, and Toueg [ACT97]. However, the heartbeat messages themselves come at a cost; this cost is investigated in Section 5. We offer some conclusions in Section 6. Some proofs are relegated to the appendix.

2 A Brief Decision Theory Primer

The aim of decision theory is to help agents make rational decisions. There are a number of equivalent ways of formalizing the decision process. In this paper, we assume that (a) we have a set $\mathcal{O}$ of possible states of the world or *outcomes*, (b) the agent can assign a *utility* from $\mathbf{R} \cup \{\infty, -\infty\}$ (denoted $\mathbf{R}^*$) to each outcome in $\mathcal{O}$, and (c) each action or choice a of the agent can be associated with a subset $\mathcal{O}_\mathsf{a}$ of $\mathcal{O}$ and a probability measure $\Pr_\mathsf{a}$ on $\mathcal{O}_\mathsf{a}$. (This is essentially equivalent to viewing $\Pr_\mathsf{a}$ as a probability measure on $\mathcal{O}$ which assigns probability 0 to the outcomes in $\mathcal{O} - \mathcal{O}_\mathsf{a}$.)

Roughly speaking, the utility associated with an outcome measures how happy the agent would be if that outcome occurred. Thus, utilities quantify the preferences of the agent. The agent prefers outcome o_1 to outcome o_2 iff the utility of o_1 is higher than that of o_2. The set $\mathcal{O}_\mathsf{a}$ of outcomes associated with an action or choice a are the outcomes that might arise if a is performed or chosen; the probability measure on $\mathcal{O}_\mathsf{a}$ represents how likely each outcome is if a is performed. These are highly nontrivial assumptions, particularly the last two. We discuss them (and to what extent they are attainable in practice) in Section 6. For now, though, we just focus on their consequences.

Recall that a *random variable* on the set $\mathcal{O}$ of outcomes is a function from $\mathcal{O}$ to $\mathbf{R}^*$. Given a random variable X and a probability measure Pr on the outcomes, the *expected value* of X with respect to Pr, denoted $\mathbf{E}^{\Pr}(X)$, is $\sum_{v \in X(\mathcal{O})} v \Pr(X = v)$, where $X(\mathcal{O})$ is the range of X and $X = v$ denotes the set $\{o \in \mathcal{O} : X(o) = v\}$. We drop the superscript Pr if it is clear from the context. Note that utility is just a random variable on outcomes. Thus, with each action or choice, we have an associated *expected utility*, where the expectation is taken with respect to $\mathcal{O}_\mathsf{a}$ and $\Pr_\mathsf{a}$. Since utilities can be infinite, we need some conventions to handle infinities in arithmetic expressions. If $x > 0$, we let $x \cdot \pm\infty = \pm\infty$; if $x < 0$, we let $x \cdot \pm\infty = \mp\infty$. For all $x \in \mathbf{R}$, we let $x + \pm\infty = \pm\infty$. Finally, we let $0 \cdot \infty = 0$. We assume that $+$ and $\cdot$ remain commutative on $\mathbf{R}^*$, so this covers all the cases but $\infty + (-\infty)$, which we take to be undefined.

The "rational choice" is typically taken to be the one that maximizes expected utility. While other notions of rationality are clearly possible, for the purposes of this paper, we focus on expected utility maximization. Again, see Section 6 for further discussion of this issue.

We can now apply these notions to the Alice-Bob example from the introduction. One way of characterizing the possible outcomes is as pairs (m_a, m_b), where m_a is the number of minutes that Alice is prepared to wait, and m_b is the time that Bob actually arrives. (If Bob does not arrive at all, we take $m_b = \infty$.) Thus, if $m_a \geq m_b$, then Alice and Bob meet at time m_b in the outcome (m_a, m_b). If $m_a < m_b$, then Alice leaves before Bob arrives. What is the utility of the outcome (m_a, m_b)? Alice and Bob may well assign different utilities to these outcomes. Since we are interested in Alice's decision, we consider Alice's utilities. A very simple assumption is that there is a fixed positive benefit **meet-Bob** to Alice if she actually meets Bob and a cost of **c-wait** for each minute she waits, and that these utilities are additive. We assume here that **c-wait** ≤ 0. (In general, *costs* are described by non-positive utilities.) Under this assumption, the utility of the outcome (m_a, m_b) is **meet-Bob** $+ m_b$**c-wait** if $m_a \geq m_b$ and m_a**c-wait** if $m_a < m_b$.

Of course, in practice, the utilities might be much more complicated and need not be additive. For example, if Alice has a magazine to read, waiting for the first fifteen minutes might be relatively painless, but after that, she might get increasingly frustrated and the cost of waiting might increase exponentially, not linearly. The benefit to meeting Bob may also depend on the time they meet, independent of Alice's frustration. For example, if they have a dinner reservation for 6 p.m. at a restaurant half an hour away, the utility of meeting Bob may drop drastically after 5:30. Finally, the utility of (m_a, m_b) might depend on m_b even if $m_a < m_b$. For example, Alice might feel happier leaving at 5:15 if she knew that Bob would arrive at 6:30 than if she knew he would arrive at 5:16.

Once Alice has decided on a utility function, she has to decide what action to take. The only choice that Alice has is how long to wait. With each choice m_a, the set of possible outcomes consists of those of the form (m_a, m_b), for all possible choices of m_b. Thus, to compute the expected utility of the choice m_a, she needs a probability measure over this set of outcomes, which effectively means a probability measure over Bob's possible arrival times.

This approach of deciding at the beginning how long to wait may seem far removed from actual practice, but suppose instead Alice sent her assistant Cindy to meet Bob. Knowing something about Bob's timeliness (or lack thereof), she may well want to give Cindy instructions for how long to wait. Taking the cost of waiting to be linear in the amount of time that Cindy waits is now not so unreasonable, since while Cindy is tied up waiting for Bob, she is not able to help Alice in other ways. If Cindy goes to meet Bob frequently for Alice, it may make more sense for Alice just to tell Cindy her utility function, and let Cindy decide how long to wait based on the information she acquires regarding Bob's punctuality. Of course, once we think in terms of Alice sending an assistant, it is but a small step to think of Alice running an application, and giving the application instructions to help it decide how to act.

3 Reliable Communication

We now consider a problem that will serve as a running example throughout the rest of the paper. Consider a system consisting of a sender p and a receiver q connected by an unreliable bidirectional link. We assume that the link satisfies the following properties:

- The transmission delay of the link is τ.
- The link can only fail by losing (whole) messages and the probability of a message loss is γ.

We assume that the transmission delay and the probability of message loss are independent of the state of the system.[2] A process is *correct* if it never crashes. For $x \in \{p, q\}$, let α_x be the probability

[2]The results of this paper hold even if these quantities do depend on the state of the link. For example, γ may be a function of the number of messages in transit. We stick to the simpler model for ease of exposition.

that x is correct (more precisely, the probability of the set of runs in which x is correct). In runs in which x is not correct, x crashes in each time unit with probability $\beta_x > 0$, independent of all other events in the system (such as the events that occurred during the previous time unit).

The assumptions that seems most reasonable to us is that $\alpha_p = \alpha_q = 0$: in practice, there is always a positive probability that a process will crash in any given round.[3] We allow the possibility that $\alpha_x \neq 0$ to facilitate comparison to most of the literature, which does not make probabilistic assumptions about failure. It also may be a useful way of modeling the scenario in which processes stay up forever "for all practical purposes" (for example, if the system is scheduled to be taken off-line before the processes crash).

We want to implement a reliable link on top of the unreliable link provided by the system. That is, we want to implement a reliable send-receive protocol SR using the (unreliable) sends and receives provided by the link, denoted send and receive. SR is a *joint* protocol, consisting of a SEND protocol for the sender and a RECEIVE protocol for the receiver. SR can be initiated by either p or q. A send-receive protocol is said to be *sender-driven* if it is initiated by p and *receiver-driven* if it is initiated by q. (Web browsing can be viewed as an instance of a receiver-driven activity. The web browser queries the web server for the content of the page.) We assume that sends and receives take place at a time t, while SENDs and RECEIVEs take place over an interval of time (since, in general, they may involve a sequence of sends and receives).

We assume that send and receive satisfy the following two properties:

- If q receives m at time t, then p sent m at time $t - \tau$ and m was not lost (since the link cannot create messages or duplicate messages and the transmission delay is known to be τ).
- If p sends m at time t, then with probability $1 - \gamma$, q will receive m at time $t + \tau$; if q does not receive m at time $t + \tau$, q will never receive it.

What specification should SR satisfy? Clearly we do not want the processes to create messages out of whole cloth. Thus, we certainly want the following requirement:

$\mathbf{S}_0$. If q finishes RECEIVing m at time t, then p must have started SENDing m at some time $t' \leq t$ and q must have received m at some time $t'' \leq t$.

We shall implicitly assume $\mathbf{S}_0$ without further comment throughout the paper.

The more interesting question is what liveness requirements SR should satisfy. Perhaps the most obvious requirement is:

$\mathbf{S}_1$. If p and q are correct and SR is started with m as the message, then q eventually finishes RECEIVing m.

Although $\mathbf{S}_1$ is very much in the spirit of typical specifications, which focus only on what happens if processes are correct, we would argue that it is rather uninteresting, for two reasons (which apply equally well to many other similar specifications). The first shows that it is too weak: If $\alpha_p = \alpha_q = 0$, then p and q are correct (i.e., never crash) with probability 0. Thus, specification $\mathbf{S}_1$ is rather uninteresting in this case: It is saying something about a set of runs with vanishingly small likelihood. The second problem shows that $\mathbf{S}_1$ is too strong: In runs where p and q are correct, there is a chance (albeit a small one) that the link may lose all messages. In this case, q cannot finish RECEIVing m, since it cannot receive m (as all the messages are lost). Thus $\mathbf{S}_1$ is not satisfied.

[3]We assume that *round* k takes place between time $k - 1$ and k.

Of course, both of these problems are well known. The standard way to strengthen $\mathbf{S}_1$ to deal with the first problem is to require only that p and q be correct for "sufficiently long", but then we need to quantify this; it is far from clear how to do so. The standard way to deal with the second problem is to restrict attention to *fair* runs, according to some notion of fairness [Fra86], and require only that q finishes RECEIVing m in fair runs. Fairness is a useful abstraction for helping us characterize conditions necessary to prove certain properties. However, what makes fairness of practical interest is that, under reasonable probabilistic assumptions, it holds with probability 1.

Our interest here, as should be evident from the introduction, is to make more explicit use of probability in writing a specification. For example, we can write a probabilistic specification like the following:

$\mathbf{S}_2$. $\lim_{t\to\infty} \Pr(q$ finishes RECEIVing m no later than t time units after the start of SR | p and q are up t time units after the start of $\mathsf{SR}) = 1$.

Requirement $\mathbf{S}_2$ avoids the two problems we saw with $\mathbf{S}_1$. It says, in a precise sense, that if p and q are up for sufficiently long, then q will $\mathsf{RECEIVE}$ m with high probability (where "sufficiently long" is quantified probabilistically). Moreover, by making only a probabilistic statement, we do not have to worry about unfair runs: They occur with probability 0.

The traditional approach has been to separate specifying the properties that a protocol must satisfy from the problem of finding the best algorithm that meets the specification. But that approach typically assumes that properties are all-or-nothing propositions. That is, it implicitly assumes that a desirable property must be true in every run (or perhaps every fair run) of a protocol. It does not allow a designer to specify that it may be acceptable for a desirable property to sometimes fail to hold, if that results in much better properties holding in general. We believe that, in general, issues of cost should not be separated from the problem of specifying the behavior of an algorithm. A protocol that satisfies a particular traditional specification may do so at the price of having rather undesirable behavior on a significant fraction of runs. For example, to ensure safety, a protocol may block 20% of the time. There may be an alternate protocol that is unsafe only 2% of the time but also blocks only 2% of the time. Whether it is better to violate safety 2% of the time and liveness 2% of the time or to never violate safety but violate liveness 20% of the time obviously depends on the context. The problem with the traditional approach is that this comparison is never even considered (any algorithm that does not satisfy safety is automatically dismissed).

While we believe $\mathbf{S}_2$ is a better specification of what is desired than $\mathbf{S}_1$, it is still not good enough for our purposes, since it does not take costs into account. Without costs, we still cannot decide if it is better to violate liveness 20% of the time or to violate safety 2% of the time and liveness 2% of the time. As a first step to thinking in terms of costs, consider the following specification:

$\mathbf{S}_3$. For each message m, the expected cost of $\mathsf{SR}(m)$ is finite.

As stated, $\mathbf{S}_3$ is not well defined, since we have not specified the cost function. We now consider a particularly simple cost function, much in the spirit of the Alice-Bob example discussed in Section 2. Let SR be a send-receive protocol. Its outcomes are just the possible *runs* or *executions*. We want to associate with each run its utility. There are two types of costs we will take into account: sending messages and waiting. The intuition is that each attempt to send a message consumes some system resources and each time unit spent waiting costs the user. The total cost is a weighted sum of the two.

More precisely, let $\mathbf{c}$-send and $\mathbf{c}$-*wait* be constants representing the cost of sending a message and of waiting one time unit, respectively. Given a run r, let $\#$-$\mathsf{send}(r)$ be the number (possibly

∞) of sends done by the protocol in run r. We now want to define $t\text{-}wait(r)$, which intuitively is the amount of time q spends waiting to RECEIVE m. When should we start counting? In the Alice-Bob example, it was clear, since Alice starts waiting for Bob at 5:00. We do not want to start counting at a fixed time, since we do not assume that the processes will start their protocol at a particular time. What we want is to start at the time when SR is invoked. When do we stop counting, assuming we started? If there are no process crashes, then we stop counting when q finishes RECEIVing m. What if there are process crashes? In traditional specifications (such as $\mathbf{S}_1$), the protocol has no obligations once a process fails. To facilitate comparison between our approach and the traditional approach, we stop counting at the time of a process crash if it happens before q finishes RECEIVing m. (Note that q may never finish RECEIVing if a process crashes.)

Let t_s be the time SR is invoked. (If no such time exists, we let $t\text{-}wait(r) = 0$.) Let t_p be the time p crashes ($t_p = \infty$ if p does not crash); let t_q be the time q crashes ($t_q = \infty$ if q does not crash); let t_f be the time q finishes RECEIVing m ($t_f = \infty$ if q does not finish). Finally let $t\text{-}wait(r) = \max\{\min\{t_p, t_q, t_f\}, t_s\} - t_s$. We take the (total) cost of run r to be

$$\mathbf{c}_0(r) = \#\text{-}\mathsf{send}(r)\mathbf{c}\text{-}\mathsf{send} + t\text{-}wait(r)\mathbf{c}\text{-}wait.$$

Note that $\mathbf{c}_0$ is a random variable on runs. If $\mathbf{c}_0(r)$ captures the cost of run r (as we are assuming here it does), then $\mathbf{S}_3$ says that we want $\mathbf{E}(\mathbf{c}_0) = \mathbf{E}(\#\text{-}\mathsf{send})\mathbf{c}\text{-}\mathsf{send} + \mathbf{E}(t\text{-}wait)\mathbf{c}\text{-}wait$ to be finite.

Note that, if SR is *not* invoked in a run r, then $\mathbf{c}_0(r) = 0$. Since we are interested in the expected cost of SR, we consider only runs in which SR is actually invoked. Also, since we are interested in the expected cost of a *single* invocation in this (and the next) section, we assume for ease of exposition that the protocol is invoked at time 0 (so $t\text{-}wait(r) = \min\{t_p, t_q, t_f\}$) throughout these two sections without further comment.

Proposition 3.1: *$\mathbf{S}_2$ and $\mathbf{S}_3$ are incomparable under cost function $\mathbf{c}_0$.*

Proof: Suppose $\alpha_p = \alpha_q = 1$. Consider a send-receive protocol SR_0 in which p sends m in every round until it receives $\text{ACK}(m)$, and q sends its kth $\text{ACK}(m)$ N^k rounds after receiving m for the kth time, where $N\gamma > 1$. (Recall that γ is the probability of message loss.) It is easy to see that SR_0 satisfies $\mathbf{S}_2$. We show that it does not satisfy $\mathbf{S}_3$ by showing that $\mathbf{E}(\#\text{-}\mathsf{send}) = \infty$.

The basic idea is that q is not acknowledging the receipt of m in a timely fashion, so p will send too many copies of m. Let $A_k = \{r : q\text{'s first } k \text{ ACKs are lost and the } (k+1)\text{st ACK makes it in } r\}$; let $A_\infty = \{r : \text{all of } q\text{'s ACKs are lost}\}$. Note that $\Pr(A_k) = \gamma^k(1-\gamma)$ and $\Pr(A_\infty) = 0$ (so we can ignore runs in A_∞ for the purpose of computing expected cost, since we adopted the convention that $0 \cdot \infty = 0$). Note also that $\mathbf{E}(\#\text{-}\mathsf{send} \mid A_k) \geq N^k$, since p cannot possibly get its first $\text{ACK}(m)$ before time N^k in runs in A_k. Thus

$$\mathbf{E}(\#\text{-}\mathsf{send}) = \sum_{k=0}^{\infty} \mathbf{E}(\#\text{-}\mathsf{send} \mid A_k)\Pr(A_k) \geq \sum_{k=0}^{\infty} N^k\gamma^k(1-\gamma).$$

It is clear that the last sum is not finite, since $N\gamma > 1$; thus the algorithm fails to satisfy $\mathbf{S}_3$.

Suppose $\alpha_p = \alpha_q = 0$. Consider the trivial protocol (i.e., the "do nothing" protocol). In a round in which both p and q are up, one of p or q will crash in the next round with probability $\beta = \beta_p + \beta_q - \beta_p\beta_q$. So the probability that the first crash happens at time k is $(1-\beta)^k\beta$. Thus one of them is expected to crash at time

$$\begin{aligned} \sum_{k=0}^{\infty} k(1-\beta)^k\beta &= \frac{\beta(1-\beta)}{(1-(1-\beta))^2} \\ &= \frac{1-\beta}{\beta}. \end{aligned}$$

(Here and elsewhere in this paper we use the well-known fact that $\sum_{k=0}^{\infty} kx^k = \frac{x}{(1-x)^2}$.) Thus, $\mathbf{E}(\mathbf{c}_0) = \frac{1-\beta}{\beta}\mathbf{c}\text{-}wait$ for the trivial protocol, so the trivial protocol satisfies $\mathbf{S}_3$, although it clearly does not satisfy $\mathbf{S}_2$. ∎

The following theorem characterizes when $\mathbf{S}_3$ is implementable with respect to the cost function $\mathbf{c}_0$. Moreover, it shows that with this cost function, when $\mathbf{S}_3$ is satisfiable, there are in fact protocols that satisfy $\mathbf{S}_3$ and $\mathbf{S}_2$ simultaneously.

Theorem 3.2: *Under cost function $\mathbf{c}_0$, there is a send-receive protocol satisfying $\mathbf{S}_3$ iff $\alpha_p = 0$ or $\alpha_q = 0$ or $\alpha_q = 1$ or $\alpha_p = 1$. Moreover, if $\alpha_p = 0$ or $\alpha_q = 0$ or $\alpha_q = 1$ or $\alpha_p = 1$, then there is a send-receive protocol that satisfies both $\mathbf{S}_2$ and $\mathbf{S}_3$.*

Proof: Suppose $\alpha_q = 1$ or $\alpha_p = 0$. Consider the (sender-driven) protocol SR_1 in which p sends m to q until p receives an ACK(m) from q, and q sends ACK(m) whenever it receives m. SR_1 starts when p first sends m and q finishes RECEIVing m when it first receives m. To see that SR_1 is correct, first consider the case that $\alpha_q = 1$. Let $C_p = \{r : p$ receives ACK(m) at least once from q in $r\}$. Let $N_1(r) = k_1$ if the k_1th copy of m is the first received by q and let $N_2(r) = k_2$ if the k_2th copy of m is the one whose corresponding ACK(m) is the first received by p.

Since the probability that the link may drop a particular message is γ,

$$\mathbf{E}(N_1 \mid C_p) = \sum_{k=1}^{\infty} k\gamma^{k-1}(1-\gamma) = \frac{1-\gamma}{\gamma}\sum_{k=1}^{\infty} k\gamma^k = \frac{1-\gamma}{\gamma}\frac{\gamma}{(1-\gamma)^2} = \frac{1}{1-\gamma}.$$

An analogous argument shows that $\mathbf{E}(N_2 \mid C_p) = \frac{1}{(1-\gamma)^2}$. Note that $t\text{-}wait(r) = N_1(r) + \tau - 1$ for $r \in C_p$, so $\mathbf{E}(t\text{-}wait \mid C_p) = \mathbf{E}(N_1 \mid C_p) + (\tau - 1) = \frac{1}{(1-\gamma)} + \tau - 1$. Moreover, since p stops sending m when it receives ACK(m) from q, it will stop 2τ rounds after the $N_2(r)$th send of m in run r. Thus $\frac{1}{(1-\gamma)^2} + 2\tau - 1$ is the number of times p is expected to send m in runs of C_p. We expect $1-\gamma$ of these to be successful, so the number of times q is expected to send ACK(m) is at most $\frac{1}{(1-\gamma)} + (2\tau-1)(1-\gamma)$. (The actual expected value is slightly less since q may crash shortly after sending the first ACK(m) received by p in runs of C_p). We conclude that $\mathbf{E}(\#\text{-send} \mid C_p) \leq \frac{1}{(1-\gamma)} + \frac{1}{(1-\gamma)^2} + (2\tau - 1)(2-\gamma)$. Thus $\mathbf{E}(\mathbf{c}_0 \mid C_p)$ is finite, since both $\mathbf{E}(\#\text{-send} \mid C_p)$ and $\mathbf{E}(t\text{-}wait \mid C_p)$ are finite.

We now turn to $\mathbf{E}(\mathbf{c}_0 \mid \overline{C_p})$. We first partition $\overline{C_p}$ into two sets:

- $F_1 = \{r : p$ crashes before receiving an ACK(m) from $q\}$ and
- $F_2 = \{r : p$ does not crash and does not receive ACK(m) from $q\}$.

Note that $\Pr(F_2) = 0$ and $\Pr(F_1) = 1 - \Pr(C_p)$. We may ignore runs of F_2 for the purposes of computing the expected cost since we adopted the convention that $0 \cdot \infty = 0$. In runs r of F_1, $t\text{-}wait(r)$ is at most the time it takes for p to crash, which is expected to occur at time $\frac{1-\beta_p}{\beta_p}$. Thus $\mathbf{E}(t\text{-}wait \mid F_1) < \frac{1}{\beta_p}$. Furthermore, if p crashes at time t_c in $r \in F_1$, it sends m exactly t_c times in r (since p does not receive ACK(m) in runs of F_1). In that case, q sends ACK(m) at most t_c times. So $\#\text{-send}(r) \leq 2t_c$ if p crashes at time t_c in $r \in F_1$. Thus $\mathbf{E}(\#\text{-send} \mid F_1) < \frac{2}{\beta_p}$. It follows that $\mathbf{E}(\mathbf{c}_0 \mid \overline{C_p})$ is finite. Since both $\mathbf{E}(\mathbf{c}_0 \mid C_P)$ and $\mathbf{E}(\mathbf{c}_0 \mid \overline{C_p})$ are finite, $\mathbf{E}(\mathbf{c}_0)$ is finite; so SR_1 satisfies $\mathbf{S}_3$. To see that the protocol satisfies $\mathbf{S}_2$, note that for $t \geq \tau$, the probability that q does *not* finish RECEIVing m by time t given that both p and q are still up is $\gamma^{t-\tau}$. Thus $\mathbf{S}_2$ is also satisfied.

Now consider the case that $\alpha_p = 0$. Note that in this case, p is expected to crash at time $\frac{1-\beta_p}{\beta_p}$. Thus, $\mathbf{E}(t\text{-}wait) < \frac{1}{\beta_p}$ and $\mathbf{E}(\#\text{-send}) < \frac{2}{\beta_p}$ (for the same reason as above), regardless of whether q is correct. Thus $\mathbf{E}(\mathbf{c}_0)$ is again finite. The argument that $\mathbf{S}_2$ is satisfied is the same as before.

Now suppose $\alpha_p = 1$ or $\alpha_q = 0$. These cases are somewhat analogous to the ones above, except we need a receiver-driven protocol. Consider a protocol SR_2 in which q queries p in every round until it gets a message from p. More precisely, let REQ denote a *request message*. q sends REQ to p every time unit until it $\mathsf{receive}$s m and p sends m every time it $\mathsf{receive}$s REQ. SR_2 starts when q sends the first REQ and q finishes RECEIVing m when q $\mathsf{receive}$s m for the first time. By reasoning similar to the previous cases, we can show that $\mathbf{E}(\#\text{-}\mathsf{send})$ and $\mathbf{E}(t\text{-}wait)$ are both finite (so $\mathbf{S}_3$ is satisfied) and that $\mathbf{S}_2$ is satisfied.

We now turn to the negative result. It turns out that the negative result is much more general than the positive result. In particular, it holds for any cost function with a certain property. In the following, we use $g \overset{\infty}{\Longrightarrow} f$ to denote that if $g(x) = \infty$ then $f(x) = \infty$.

Lemma 3.3: *Let $\mathbf{c}(r)$ be a cost function such that $t\text{-}wait(r) \overset{\infty}{\Longrightarrow} \mathbf{c}(r)$ and $\#\text{-}\mathsf{send}(r) \overset{\infty}{\Longrightarrow} \mathbf{c}(r)$. If $0 < \alpha_p < 1$ and $0 < \alpha_q < 1$, then for any send-receive protocol SR, $\Pr(\{r : \mathbf{c}(r) = \infty\}) > 0$.*

Proof: Suppose SR is a send-receive protocol for p and q. Let $R_1 = \{r : q$ crashes at time 0 and p is correct in $r\}$. Note that p will do the same thing in all runs in R_1: Either p stops sending after some time t or p never stops sending. If p never stops, then $\#\text{-}\mathsf{send}(r) = \infty$ for all $r \in R_1$. Since, by assumption, $\#\text{-}\mathsf{send}(r) \overset{\infty}{\Longrightarrow} \mathbf{c}(r)$, we have that $\mathbf{c}(r) = \infty$ for each $r \in R_1$. Since $\Pr(R_1) = \alpha_p(1-\alpha_q)\beta_q > 0$, we are done. Now suppose p stops sending after time t. Let $R_2 = \{r : p$ crashes at time 0 and q is correct in $r\}$. Note that q will do the same thing in all runs of R_2: Either q stops sending after some time t' or q never stops sending. If q never stops, then $\mathbf{c}(r) = \infty$ for all $r \in R_2$ and $\Pr(R_2) = \alpha_q(1-\alpha_p)\beta_p > 0$, so again we are done. Finally, suppose that q stops sending at time t' in runs of R_2. Let $t'' = 1 + \max\{t, t'\}$. Consider $R_3 = \{r :$ both processes are correct and all messages up to time t'' are lost in $r\}$. Then $t\text{-}wait(r) = \infty$ for all $r \in R_3$. By assumption, $t\text{-}wait(r) \overset{\infty}{\Longrightarrow} \mathbf{c}(r)$, so $\mathbf{c}(r) = \infty$ for all $r \in R_3$. Let n_p and n_q be the number of invocations of send by p and q, respectively, in runs of R_3 (note that p and q do the same thing in all runs of R_3). Then $\Pr(R_3) = \alpha_p\alpha_q\gamma^{n_p+n_q} > 0$, completing the proof. ▮ (Lemma 3.3)

Clearly $\#\text{-}\mathsf{send}(r) \overset{\infty}{\Longrightarrow} \mathbf{c}_0(r)$ and $t\text{-}wait(r) \overset{\infty}{\Longrightarrow} \mathbf{c}_0(r)$, so Lemma 3.3 applies immediately and we are done. ▮ (Theorem 3.2)

Of course, once we think in terms of utility-based specifications like $\mathbf{S}_3$, we do not want to know just whether a protocol implements $\mathbf{S}_3$; we are in a position to compare the performance of different protocols that implement $\mathbf{S}_3$ (or of variants of one protocol that all implement $\mathbf{S}_3$) by considering their expected utility. Let SR_s^δ and SR_r^δ be generalizations (in the sense that they send messages every δ rounds, where δ need not be 1) of the sender-driven and receiver-driven protocols from Theorem 3.2, respectively. Let SR_{tr} denote the trivial (i.e., "do nothing") protocol. We use $\mathbf{E}^{\mathsf{SR}}$ to denote the expectation operator determined by the probability measure on runs induced by using protocol SR. Thus, for example, $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send})$ is the expected number of messages sent by SR_s^δ. If $\alpha_p = \alpha_q = 0$, then SR_s^δ, SR_r^δ, and SR_{tr} all satisfy $\mathbf{S}_3$ (although SR_{tr} does not satisfy $\mathbf{S}_2$). Which is better?

In practice, process failures and link failures are very unlikely events. We assume in the rest of the paper that β_p, β_q, and γ are all very small, so that we can ignore sums of products of these terms (with coefficients like $2\tau^2$, δ, etc.). One way to formalize this is to say that products involving β_p, β_q, and γ are $O(\varepsilon)$ terms and $2\tau^2$, δ, etc., are $O(1)$ terms. We write $t_1 \approx t_2$ if $|t_1 - t_2|$ is $O(\varepsilon)$. Note that we do not assume expressions like $\frac{\beta_p}{\beta_q}$ and $\frac{\beta_q}{\beta_p}$ are small.

For the following result only, we assume that not only are β_p and β_q $O(\varepsilon)$, they are also $\Theta(\varepsilon)$,[4] so that if $\frac{1}{\beta_p}$ or $\frac{1}{\beta_q}$ is multiplied by an expression that is $O(\varepsilon^2)$, then the result is $O(\varepsilon)$, which can

[4] Recall that x is $\Theta(\varepsilon)$ iff x is $O(\varepsilon)$ and x^{-1} is $O(\varepsilon^{-1})$.

then be ignored.

Proposition 3.4: *If $\alpha_p = \alpha_q = 0$, then*

$$\begin{array}{ll}
\mathbf{E}^{\mathsf{SR}_{tr}}(t\text{-}wait) = \frac{1-(\beta_p+\beta_q-\beta_p\beta_q)}{\beta_p+\beta_q-\beta_p\beta_q}, & \mathbf{E}^{\mathsf{SR}_{tr}}(\#\text{-}\mathsf{send}) = 0, \\
\mathbf{E}^{\mathsf{SR}^\delta_s}(t\text{-}wait) \approx \tau, & \mathbf{E}^{\mathsf{SR}^\delta_s}(\#\text{-}\mathsf{send}) \approx \frac{(\tau+1)\beta_q}{\delta\beta_p} + 2\left\lceil\frac{2\tau}{\delta}\right\rceil, \\
\mathbf{E}^{\mathsf{SR}^\delta_r}(t\text{-}wait) \approx 2\tau, & \mathbf{E}^{\mathsf{SR}^\delta_r}(\#\text{-}\mathsf{send}) \approx \frac{(\tau+1)\beta_p}{\delta\beta_q} + 2\left\lceil\frac{2\tau}{\delta}\right\rceil.
\end{array}$$

Proof: The relatively straightforward (but tedious!) calculations are relegated to the appendix. ∎

Note that the expected cost of messages for SR^δ_s is the same as that for SR^δ_r, except that the roles of β_p and β_q are reversed. The expected time cost of SR^δ_r is roughly τ higher than that of SR^δ_s, because q cannot finish RECEIVing m before time 2τ with a receiver-driven protocol, whereas q may finish RECEIVing m as early as τ with a sender-driven protocol. This says that the choice between the sender-driven and receiver-driven protocol should be based largely on the relative probability of failure of p and q. It also suggests that we should take δ very large to minimize costs. (Intuitively, the larger δ is, the lower the message costs in the case that q crashes before acknowledging p's message.) This conclusion (which may not seem so reasonable) is essentially due to the fact that we are examining a single invocation of SR in isolation. As we shall see in Section 5, this conclusion is no longer justified once we consider repeated invocations of SR. Finally, note that if the cost of messages is high and waiting is cheap, the processes are better off (according to this cost function) using SR_{tr}.

Thus, as far as $\mathbf{S}_3$ is concerned, there are times when SR_{tr} is better than SR^δ_s or SR^δ_r. How much of a problem is it that SR_{tr} does not satisfy $\mathbf{S}_2$? Our claim is that if this desideratum (i.e., $\mathbf{S}_2$) is important, then it should be reflected in the cost function. While the cost function in our example does take into account waiting time, it does not penalize it sufficiently to give us $\mathbf{S}_2$. It is not too hard to find a cost function that captures $\mathbf{S}_2$. For example, suppose we take $\mathbf{c}_1(r) = N^{t\text{-}wait(r)}$, where $N(1-\beta_p-\beta_q+\beta_p\beta_q) > 1$.

Proposition 3.5: *Under cost function $\mathbf{c}_1$, $\mathbf{S}_3$ implies $\mathbf{S}_2$.*

Proof: Suppose SR is a protocol that does not satisfy $\mathbf{S}_2$; we show it does not satisfy $\mathbf{S}_3$ (under cost function $\mathbf{c}_1$). Let $C_p(t)$ and $C_q(t)$ consist of those runs of SR where p and q, respectively, are up for t time units after the start of SR (and perhaps longer). Let $R_q(t)$ consist of the runs of SR where q finishes RECEIVing m no later than time t units after the start of SR. Since SR does not satisfy $\mathbf{S}_2$, there exists $\varepsilon > 0$ and an increasing infinite sequence of times $t_0, t_1, \ldots$, such that $\Pr(\overline{R_q(t_i)} \mid C_p(t_i) \cap C_q(t_i)) > \varepsilon$ for all i. We consider the case $\alpha_p = \alpha_q = 1$ and $\alpha_p\alpha_q < 1$ separately.

Suppose $\alpha_p = \alpha_q = 1$. Then $\Pr(C_p(t) \cap C_q(t)) = 1$ for all t. So

$$\Pr(t\text{-}wait > t_i) = \Pr(\overline{R_q(t_i)}) = \Pr(\overline{R_q(t_i)} \mid C_p(t_i) \cap C_q(t_i)) > \varepsilon$$

for all i. Let $V_i = \{r : t\text{-}wait(r) > t_i\}$ and $V_\infty = \{r : t\text{-}wait(r) = \infty\}$. Note that $V_\infty = \bigcap_{i=0}^\infty V_i$ and that $V_i \supseteq V_{i'}$ for $i' > i$. Thus $\Pr(V_\infty) = \Pr(\bigcap_{i=0}^\infty V_i) > \varepsilon$. So $\mathbf{E}(\mathbf{c}_1) \geq \Pr(V_\infty)N^\infty = \infty$.

Now we turn to the case that $\alpha_p\alpha_q < 1$. Let $W(t) = \{r : t\text{-}wait(r) = t\}$. Note that $t\text{-}wait(r) = t_i + 1$ for all runs $r \in \overline{R_q(t_i)} \cap \overline{C_p(t_i+1)} \cap C_p(t_i) \cap C_q(t_i)$. Thus,

$$\Pr(W(t_i+1) \mid C_p(t_i) \cap C_q(t_i)) \geq \Pr(\overline{C_p(t_i+1)} \cap \overline{R_q(t_i)} \mid C_p(t_i) \cap C_q(t_i)).$$

Given our independence assumptions regarding process failures,

$$\begin{array}{rcl}
\Pr(\overline{C_p(t_i+1)} \cap \overline{R_q(t_i)} \mid C_p(t_i) \cap C_q(t_i)) & = & \Pr(\overline{C_p(t_i+1)} \mid C_p(t_i)) \Pr(\overline{R_q(t_i)} \mid C_p(t_i) \cap C_q(t_i)) \\
& > & (1-\alpha_p)\beta_p\varepsilon.
\end{array}$$

A similar argument (exchanging the roles of C_p and C_q) shows that

$$\Pr(W(t_i+1) \mid C_p(t_i) \cap C_q(t_i)) > (1-\alpha_q)\beta_q\varepsilon.$$

So

$$\begin{aligned}
\mathbf{E}(\mathbf{c}_1) &\geq \sum_{k=0}^{\infty} \Pr(W(k))N^k \\
&\geq \sum_{i=0}^{\infty} \Pr(W(t_i+1))N^{t_i+1} \\
&\geq \sum_{i=0}^{\infty} \Pr(W(t_i+1) \cap C_p(t_i) \cap C_q(t_i))N^{t_i+1} \\
&= \sum_{i=0}^{\infty} \Pr(W(t_i+1) \mid C_p(t_i) \cap C_q(t_i)) \Pr(C_p(t_i) \cap C_q(t_i)))N^{t_i+1} \\
&> \max\{(1-\alpha_p)\beta_p, (1-\alpha_q)\beta_q\}\varepsilon \sum_{i=0}^{\infty}(1-\beta_p-\beta_q+\beta_p\beta_q)^{t_i}N^{t_i+1}.
\end{aligned}$$

Since $(1-\beta_p-\beta_q+\beta_p\beta_q)N > 1$ by assumption, we are done. ∎

The moral here is that $\mathbf{S}_3$ gives us the flexibility to specify what really matters in a protocol, by appropriately describing the cost function. We would like to remind the reader that the cost functions are not ours to choose: They reflect the user's preferences. (Thus we are not saying that $\mathbf{c}_1$ is better than $\mathbf{c}_0$ or vice versa, since each user is entitled to her own preferences.) What we are really saying here is that if $\mathbf{S}_2$ matters to the user, then her cost function would force $\mathbf{S}_3$ to imply $\mathbf{S}_2$—in particular, her cost function could not be $\mathbf{c}_0$.

4 Using Heartbeats

We saw in Section 3 that $\mathbf{S}_3$ is not implementable if we are not certain about the correctness of the processes (i.e., if the probability that they are correct is strictly between 0 and 1) and the cost function $\mathbf{c}(r)$ has the property that $\#\text{-}\mathsf{send}(r) \overset{\infty}{\Longrightarrow} \mathbf{c}(r)$ and $t\text{-}wait(r) \overset{\infty}{\Longrightarrow} \mathbf{c}(r)$. Aguilera, Chen, and Toueg [ACT97] (ACT from now on) suggest an approach that circumvents this problem, using *heartbeat* messages. Informally, a heartbeat from process i is a message sent by i to all other processes to tell them that it is still alive. ACT show that there is a protocol using heartbeats that achieves *quiescent* reliable communication; i.e., in every run of the protocol, only finitely many messages are required to achieve reliable communication (not counting the heartbeats). Moreover, they show that, in a precise sense, quiescent reliable communication is not possible if we are not certain about the correctness of the processes and communication is unreliable, a result much in the spirit of the negative part of Theorem 3.2.[5] In this section, we show that (using the cost function $\mathbf{c}_0$) we can use heartbeats to implement $\mathbf{S}_3$ for all values of α_p and α_q.

For the purposes of this paper, assume that processes send a message we call HBMSG to each other every δ time units. Protocol SR_{hb} in Figure 1 is a protocol for reliable communication based on ACT's protocol. (It is not as general as theirs, but it retains all the features relevant to us.) Briefly, what happens according to this protocol is that the failure detector layer of q sends HBMSG to the corresponding layer of p periodically. If p wants to SEND m, p checks to see if any (new)

[5] ACT actually show that their impossibility result holds even if there is only one process failure, only finitely many messages can be lost, and the processes have access to $\mathcal{S}$ (a *strong* failure detector), which means that eventually every faulty process is permanently suspected and at least one correct process is never suspected. The model used by ACT is somewhat different from the one we are considering, but we can easily modify their results to fit our model.

The sender's protocol (SEND):

1. **while** $\neg\mathsf{receive}(\text{ACK}(m))$ **do**
2. **if** $\mathsf{receive}(\text{HBMSG})$ **then**
3. $\mathsf{send}(m)$
4. **fi**
5. **od**

The receiver's protocol (RECEIVE):

1. **while true do**
2. **if** $\mathsf{receive}(m)$ **then**
3. $\mathsf{send}(\text{ACK}(m))$
4. **fi**
5. **od**

Figure 1: Protocol SR_{hb}

HBMSG has arrived; if so, p sends m to q, provided it has not already received ACK(m) from q; q sends ACK(m) every time it receives m and q finishes RECEIVing m the first time it receives m. Note that q does *not* send any HBMSGs as part of SR_{hb}. That is the job of the failure-detection layer, not the job of the protocol. (We assume that the protocol is built on top of a failure-detection service.) The cost function of the previous section does ot count the costs of HBMSGs. That is, since $\#\text{-}\mathsf{send}(r)$ is the number of messages sent by the protocol, $\mathbf{c}_0(r)$ is not affected by the number of HBMSGs sent in run r. It is also worth noting that this is a sender-driven protocol, quite like that given in the proof of Theorem 3.2.[6] It is straightforward to also design a receiver-driven protocol using heartbeats.

We now want to show that SR_{hb} implements $\mathbf{S}_3$ and get a good estimate of the actual expected cost.

Theorem 4.1: *Under cost function $\mathbf{c}_0$, Protocol SR_{hb} satisfies $\mathbf{S}_3$. Moreover, $\mathbf{E}(t\text{-}wait) \approx 2\tau$ and $\mathbf{E}(\#\text{-}\mathsf{send}) \approx 2\left\lceil \frac{2\tau}{\delta} \right\rceil$, so that $\mathbf{E}(\mathbf{c}_0) \approx 2\tau\mathbf{c}\text{-}wait + 2\left\lceil \frac{2\tau}{\delta} \right\rceil \mathbf{c}\text{-}\mathsf{send}$.*

Proof: Using arguments similar to those of the proof of Proposition 3.4, we can show that $\mathbf{E}(t\text{-}wait) \approx 2\tau$ and $\mathbf{E}(\#\text{-}\mathsf{send}) \approx 2\left\lceil \frac{2\tau}{\delta} \right\rceil$. We leave details to the reader. ∎

The analysis of SR_{hb} is much like that of SR_s^δ in Proposition 3.4. Indeed, in the case that $\alpha_p = \alpha_q = 0$, the two protocols are almost identical. The waiting time is roughly τ more for SR_{hb}, since p does not start sending until it receives the first HBMSG from q. On the other hand, we are better off using SR_{hb} if q crashes before acknowledging p's message. In this case, with SR_s^δ, p continues to send until it crashes, while with SR_{hb}, it stops sending (since it does not get any HBMSGs from q). This leads to an obvious question: Is it really worth sending heartbeats? Of course, if both α_p and α_q are between 0 and 1, we need heartbeats or something like them to get around the impossibility result of Theorem 3.2. But if $\alpha_p = \alpha_q = 0$, then we need to look carefully at the relative size of $\mathbf{c}\text{-}\mathsf{send}$ and $\mathbf{c}\text{-}wait$ to decide which protocol has the lower expected cost.

This suggests that the decision of whether to implement a heartbeat layer must take probabilities and utilities seriously, even if we do not count either the overhead of building such a layer or the cost of heartbeats. What happens if we take the cost of heartbeats into account? This is the subject of the next section.

[6] The reader might notice that the runs induced by this protocol actually resemble those of the *receiver-driven* protocol in the proof of Theorem 3.2 (if we identify HBMSG with REQ). The difference is that in the receiver-driven protocol in the proof of Theorem 3.2, the protocol for the receiver actually sends the REQs whereas here the HBMSGs are sent not by the protocol but by an underlying heartbeat layer, independent of the protocol.

5 The Cost of Heartbeats

In the previous section we showed that $\mathbf{S}_3$ is achievable with the help of heartbeats. When we computed the expected costs, however, we did so with the cost function $\mathbf{c}_0$, which does not count the cost of heartbeats. While someone who takes the heartbeat layer for granted (such as an application programmer or end-user) may have $\mathbf{c}_0$ as their cost function, someone who has to decide whether to implement a heartbeat layer or how frequently heartbeats should be sent (such as a system designer) is likely to have a different cost function—one which takes the cost of heartbeats into account.

As evidence of this, note that it is immediate from Theorem 4.1 that under the cost function $\mathbf{c}_0$, the choice of δ that minimizes the expected cost is clearly at most $2\tau + 1$. Intuitively, if we do not charge for heartbeats, there is no incentive to space them out. On the other hand, if we do charge for heartbeats, then typically we will be charging for heartbeats that are sent long after a given invocation of SR_{hb} has completed.

The whole point of having a heartbeat layer is that heartbeats are meant to be used, not just by one invocation of a single protocol, but by multiple invocations of (possibly) many protocols. We would expect that the optimal frequency of heartbeats should depend in part on how often the protocols that use them are invoked. The picture we have is that the SR_{hb} protocol is invoked from time to time, by different processes in the system. It may well be that various invocations of it are running simultaneously. All these invocations share the heartbeat messages, so their cost can be spread over all of them. If invocations occur often, then there will be few "wasted" heartbeats between invocations, and the analysis of the previous subsection gives a reasonably accurate reading of the costs involved. On the other hand, if δ is small and invocations are infrequent, then there will be many "wasted" heartbeats. We would expect that if there are infrequent invocations, then heartbeats should be spaced further apart.

We now consider a setting that takes this into account. For simplicity, we continue to assume that there are only two processes, p and q, but we now allow both p and q to invoke SR_{hb}. (It is possible to do this with n processes and more than one protocol, but the two-process and single protocol case suffices to illustrate the main point, which is that the optimal δ should depend on how often the protocol is invoked.) We assume that each process, while it is running, invokes SR_{hb} with probability σ at each time unit. Thus, informally, at every round, each running process tosses a coin with probability of σ of landing heads. If it lands heads, the process then invokes SR_{hb} with the other as the recipient. (Note that we no longer assume that the protocol is invoked at time 0 in this section.)

Roughly speaking, in computing the cost of a run, we consider the cost of each invocation of SR_{hb} together with the cost of all the heartbeat messages sent in the run. Our interest will then be in the cost *per invocation of* SR_{hb}. Thus, we apportion the cost of the heartbeat messages among the invocations of SR_{hb}. If there are relatively few invocations of SR_{hb}, then there will be many "wasted" heartbeat messages, whose cost will need to be shared among them.

For simplicity, let us assume that each time SR_{hb} is invoked, a different message is sent. (For example, messages could be numbered and include the name of the sender and recipient.) We say $\mathsf{SR}_{hb}(m)$ *is invoked at time* t_1 *in* r if at time t_1 some process x first executes line 1 of the code of the sender with message m. This invocation of SR_{hb} *completes* at time t_2 if the last message associated with the invocation (either a copy of m or a copy of $\textsc{ack}(m)$) is sent at time t_2. If x received the last heartbeat message from the receiver before invoking $\mathsf{SR}_{hb}(m)$, we take $t_2 = t_1$ (that is, the invocation completes as soon as it starts in this case).

The processes will (eventually) stop $\mathsf{sending}$ m or $\textsc{ack}(m)$ if either process crashes or if the sender $\mathsf{receives}$ $\textsc{ack}(m)$. Thus, with probability 1, all invocations of SR_{hb} will eventually complete.

Let $\#\text{-}\mathsf{SR}(r,t)$ be the number of invocations of SR_{hb} that have completed by time t in r; let $\mathbf{c}\text{-}\mathsf{SR}(r,t)$ be the cost of these invocations. Let $\mathbf{c}\text{-}\textsc{hbmsg}(r,t)$ be the cost of sending $\textsc{hbmsg}$ up to time t in r. This is simply the number of $\textsc{hbmsg}$s sent up to time t (which we denote by $\#\text{-}\textsc{hbmsg}(r,t)$) multiplied by $\mathbf{c}\text{-}\mathsf{send}$. Let $\mathbf{c}^{\text{total}}(r,t) = \mathbf{c}\text{-}\mathsf{SR}(r,t) + \mathbf{c}\text{-}\textsc{hbmsg}(r,t)$. Finally, let

$$\mathbf{c}^{\text{avg}}(r) = \limsup_{t\to\infty} \frac{\mathbf{c}^{\text{total}}(r,t)}{\#\text{-}\mathsf{SR}(r,t)+1},$$

where "lim sup" denotes the limit of the supremum, that is,

$$\mathbf{c}^{\text{avg}}(r) = \lim_{t'\to\infty} \sup_{0\le t\le t'} \frac{\mathbf{c}^{\text{total}}(r,t)}{\#\text{-}\mathsf{SR}(r,t)+1}.[7]$$

Thus $\mathbf{c}^{\text{avg}}(r)$ is essentially the average cost per invocation of SR_{hb}, taking heartbeats into account. We write "lim sup" instead of "lim" since the limit may not exist in general. (However, the proof of the next theorem shows that in fact, with probability 1, the limit does exist.) For the following result only, we assume that $\sqrt{\beta_p}$ and $\sqrt{\beta_q}$ are also $O(\varepsilon)$.

Theorem 5.1: *Under the cost function $\mathbf{c}^{\text{avg}}$, Protocol SR_{hb} satisfies $\mathbf{S}_3$. Furthermore, $\mathbf{E}(\mathbf{c}^{\text{avg}}) \approx ((1-\alpha_p)(1-\alpha_q)\lambda + \alpha_p\alpha_q)\left(2\left\lceil\frac{2\tau}{\delta}\right\rceil \mathbf{c}\text{-}\mathsf{send} + \left(\tau + \frac{\delta-1}{2}\right)\mathbf{c}\text{-}wait\right) + \frac{1}{\delta\sigma}\mathbf{c}\text{-}\mathsf{send}$, where $0 < \lambda < 1$.*

Proof: See the appendix. ∎

Note that with this cost function, we have a real decision to make in terms of how frequently to send heartbeats. As before, there is some benefit to making $\delta > 2\tau$: it minimizes the number of redundant messages sent when SR_{hb} is invoked (that is, messages sent by the sender before receiving the receiver's acknowledgment). Also, by making δ larger we will send fewer heartbeat messages between invocations of SR_{hb}. On the other hand, if we make δ too large, then the sender may have to wait a long time after invoking SR_{hb} before it can send a message to the receiver (since messages are only sent upon receipt of a heartbeat). Intuitively, the greater $\mathbf{c}\text{-}wait$ is relative to $\mathbf{c}\text{-}\mathsf{send}$, the smaller we should make δ. Clearly we can find an optimal choice for δ by standard calculus.

In the model just presented, if $\mathbf{c}\text{-}wait$ is large enough relative to $\mathbf{c}\text{-}\mathsf{send}$, we will take δ to be 1. Taking δ this small is clearly inappropriate once we consider a more refined model, where there are buffers that may overflow. In this case, both the probability of message loss and the time for message delivery will depend on the number of messages in transit. The basic notions of utility still apply, of course, although the calculations become more complicated. This just emphasizes the obvious point is that in deciding what value (or values) δ should have, we need to carefully look at the actual system and the cost function.

6 Discussion

We have tried to argue here for the use of decision theory both in the specification and the design of systems. Our (admittedly rather simple) analysis already shows both how decision theory can help guide the decision made and how much the decision depends on the cost function. None of our results are deep; the cost function just makes precise what could already have been seen from an intuitive calculation. But this is precisely the point: By writing our specification in terms of costs, we can make the intuitive calculations precise. Moreover, the specification forces us to make clear

[7] By adding 1 to the denominator, we guarantee it is never 0; adding 1 also simplifies one of the technical calculations needed in the proof of Theorem 5.1.

exactly what the cost function is and encourages the elicitation of utilities from users. We believe that these are both important features. It is important for the user (and system designer) to spend time thinking about what the important attributes of the system are and to decide on preferences between various tradeoffs.

A possible future direction is to study standard problems in the literature (e.g., Consensus, Byzantine Agreement, Atomic Broadcast, etc.) and recast the specifications in utility-theoretic terms. One way to do this is to replace a liveness requirement by an unbounded increasing cost function (which is essentially the "cost of waiting") and replace a safety requirement by a large penalty. Once we do this, we can analyze the algorithms that have been used to solve these problems, and see to what extent they are optimal given reasonable assumptions about probabilities and utilities.

While we believe that there is a great deal of benefit to be gained from analyzing systems in terms of utility, it is quite often a nontrivial matter. Among the most significant difficulties are the following:

1. Where are the utilities coming from? It is far from clear that a user can or is willing to assign a real-valued utility to all possible outcomes in practice. There may be computational issues (for example, the set of outcomes can be enormous) as well as psychological issues. While the agent may be prepared to assign qualitative utilities like "good", "fair", or "bad", he may not be prepared to assign 20.7. While to some extent the system can convert qualitative utilities to a numerical representation, this conversion may not precisely captures the user's intent. There are also nontrivial user-interface issues involved in eliciting utilities from users. In light of this, we need to be very careful if results depend in sensitive ways on the details of the utilities.

2. Where are the probabilities coming from? We do not expect users to be experts at probability. Rather, we expect the system to be gathering statistics and using them to estimate the probabilities. Of course, someone still has to tell the system what statistics to gather. Moreover, our statistics may be so sparse that we cannot easily obtain a reliable estimate of the probability.

3. Why is it even appropriate to maximize expected utility? There are times when it is far from clear that this is the best thing to do, especially if our estimates of the probability and utility are suspect. For example, suppose one action has a guaranteed utility of 100 (on some appropriate scale), while another has an expected utility of 101, but has a nontrivial probability of having utility 0. If the probabilities and utilities that were used to calculate the expectation are reliable, and we anticipate performing these actions frequently, then there is a good case to be made for taking the action with the higher expected utility. On the other hand, if the underlying numbers are suspect, then the action with the guaranteed utility might well be preferable.

We see these difficulties not as ones that should prevent us from using decision theory, but rather as directions for further research. It may be possible in many cases to learn a user's utility. Moreover, we expect that in many applications, except for a small region of doubt, the choice of which decision to make will be quite robust, in that perturbations to the probability and utility will not change the decision. Even in cases where perturbations do change the decision, both decisions will have roughly equal expected utility. Thus, as long as we can get somewhat reasonable estimates of the probability and utility, decision theory may have something to offer.

Another important direction for research is to consider *qualitative decision theory*, where both utility and likelihood are more qualitative, and not necessarily real numbers. This is, in fact, an active area of current research, as `http://www.medg.lcs.mit.edu/qdt/bib/unsorted.bib` (a bibliography of over 290 papers) attests. Note that once we use more qualitative notions, then we may not be able to compute expected utilities at all (since utilities may not be numeric) let alone take the action with maximum expected utility, so we will have to consider other decision rules.

Finally, we might consider what would be an appropriate language to specify and reason about utilities, both for the user and the system designer.

While it is clear that there is still a great deal of work to be done in order to use decision-theoretic techniques in systems design and specification, we hope that this discussion has convinced the reader of the utility of the approach.

Acknowledgments

We thank Sam Toueg for numerous discussions regarding heartbeats and Jim Gray for giving us some insight on costs in database computations and for pointing out the use of costs in deadlock detection. We also thank the anonymous referees for their helpful comments.

Appendix: Proofs

We present the proofs of Proposition 3.4 and Theorem 5.1. We repeat the statements of the results for the convenience of the reader. Recall that for Proposition 3.4, we are assuming that β_p and β_q are both $\Theta(\varepsilon)$, and that for Theorem 5.1, we are assuming that $\sqrt{\beta_p}$ and $\sqrt{\beta_q}$ are both $O(\varepsilon)$.

Proposition 3.4: *If $\alpha_p = \alpha_q = 0$, then*

$$\begin{array}{ll} \mathbf{E}^{\mathsf{SR}_{tr}}(t\text{-}wait) = \frac{1-(\beta_p+\beta_q-\beta_p\beta_q)}{\beta_p+\beta_q-\beta_p\beta_q}, & \mathbf{E}^{\mathsf{SR}_{tr}}(\#\text{-}\mathsf{send}) = 0, \\ \mathbf{E}^{\mathsf{SR}_s^\delta}(t\text{-}wait) \approx \tau, & \mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send}) \approx \frac{(\tau+1)\beta_q}{\delta\beta_p} + 2\left\lceil\frac{2\tau}{\delta}\right\rceil, \\ \mathbf{E}^{\mathsf{SR}_r^\delta}(t\text{-}wait) \approx 2\tau, & \mathbf{E}^{\mathsf{SR}_r^\delta}(\#\text{-}\mathsf{send}) \approx \frac{(\tau+1)\beta_p}{\delta\beta_q} + 2\left\lceil\frac{2\tau}{\delta}\right\rceil. \end{array}$$

Proof: For SR_{tr}, note that $\#\text{-}\mathsf{send}(r) = 0$ for all r, so $\mathbf{E}^{\mathsf{SR}_{tr}}(\#\text{-}\mathsf{send}) = 0$. We also have that $t\text{-}wait(r)$ is the time of the first crash in r. Since the probability of a crash during a time unit is $\beta = \beta_p + \beta_q - \beta_p\beta_q$, we have that the expected time of the first crash, and hence $\mathbf{E}^{\mathsf{SR}_{tr}}(t\text{-}wait)$, is

$$\sum_{k=0}^{\infty} k(1-\beta)^k\beta = \frac{\beta(1-\beta)}{(1-(1-\beta))^2} = \frac{1-\beta}{\beta} = \frac{1-(\beta_p+\beta_q-\beta_p\beta_q)}{\beta_p+\beta_q-\beta_p\beta_q}.$$

For SR_s^δ, we first show that $\mathbf{E}^{\mathsf{SR}_s^\delta}(t\text{-}wait) \approx \tau$. Since $\alpha_p = \alpha_q = 0$, $\Pr(t\text{-}wait(r) = \infty) = 0$, thus $\mathbf{E}^{\mathsf{SR}_s^\delta}(t\text{-}wait) = \sum_{k=1}^{\infty} k \Pr(t\text{-}wait = k)$. We break the sum into three pieces,

- $\sum_{k=1}^{\tau-1} k \Pr(t\text{-}wait = k)$,
- $\Pr(t\text{-}wait = \tau)$, and
- $\sum_{k=\tau+1}^{\infty} k \Pr(t\text{-}wait = k)$,

and analyze each one separately.

For the first part, note that the only way that $t\text{-}wait = k$ for $1 \leq k < \tau$ is for there to be a crash before τ. Thus

$$\Pr(t\text{-}wait = k) = ((1-\beta_p)(1-\beta_q))^k(\beta_p + \beta_q - \beta_p\beta_q) < \beta_p + \beta_q.$$

It follows that

$$\sum_{k=1}^{\tau-1} k \Pr(t\text{-}wait = k) < (\beta_p + \beta_q)\sum_{k=1}^{\tau-1} k = (\beta_p + \beta_q)\frac{\tau(\tau-1)}{2} \approx 0.$$

Thus we may drop the first part.

For the second part, note that $t\text{-}wait = \tau$ if p and q are up until τ and q received the first copy of m p sent. (We may also have $t\text{-}wait = \tau$ if one of p or q crashes at time τ.) Thus,

$$\Pr(t\text{-}wait = \tau) \geq ((1-\beta_p)(1-\beta_q))^\tau(1-\gamma) \approx 1,$$

so the second part is $\approx \tau$.

Finally, for the third part, if $k > \tau$, then k has the form $\tau + a\delta + b$, where $a \geq 0$ and $0 \leq b < \delta$ (and $a + b > 0$). If $t\text{-}wait = k = \tau + a\delta + b$, then $a+1$ messages are lost by the link, so $\Pr(t\text{-}wait = k) \leq \gamma^{a+1}$. A straightforward calculation shows that

$$\begin{aligned}
\sum_{k=\tau+1}^{\infty} k \Pr(t\text{-}wait = k) &= \sum_{b=1}^{\delta-1}(\tau+b)\Pr(t\text{-}wait = \tau+b) \\
&\quad + \sum_{a=1}^{\infty}\sum_{b=0}^{\delta-1}(\tau + a\delta + b)\Pr(t\text{-}wait = \tau + a\delta + b) \\
&\leq \sum_{a=0}^{\infty} \delta(\tau + (a+1)\delta)\gamma^{a+1} \\
&\leq \sum_{a=0}^{\infty}((a+1)\delta^2 + \delta\tau)\gamma^{a+1} \\
&= \delta^2\sum_{a=1}^{\infty} a\gamma^a + \delta\tau\sum_{a=1}^{\infty}\gamma^a \\
&\approx 0.
\end{aligned}$$

Thus, we can also ignore the third part. This gives us $\mathbf{E}^{\mathsf{SR}_s^\delta}(t\text{-}wait)) \approx \tau$, as desired.

Now let us turn to $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send})$. Let us say that a send is *successful* iff the link does not drop the message (which could be an ACK). Consider the set of runs $A = \{r : q$ successfully sends ACK(m) before crashing in $r\}$. Roughly speaking, what happens is that in runs of A, p is receives ACK(m) at time 2τ with probability ≈ 1. In the meantime, p has sent m exactly $\left\lceil \frac{2\tau}{\delta} \right\rceil$ times with probability ≈ 1. With probability ≈ 1, all of these are received by q; q in turn acknowledges all copies and thus $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid A) \approx 2\left\lceil \frac{2\tau}{\delta} \right\rceil$; that is why this term appears in $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send})$. In $\overline{A}$, the expected value of #-send is very large, since p will send m until it crashes, so despite the low probability of $\overline{A}$, it contributes the term $\frac{(\tau+1)\beta_q}{\delta\beta_p}$. We now turn to the details.

We first compute $\Pr(A)$. Note that q can send ACK(m) only at times of the form $\tau + k\delta$. Let $B_k = \{r : q$ sends the first successful ACK(m) at time $\tau + k\delta\}$. Note that $A = \bigcup_{k=0}^\infty B_k$ and that $B_i \cap B_j = \emptyset$ if $i \neq j$. Thus $\Pr(A) = \sum_{k=0}^\infty \Pr(B_k)$. Since q sends the first successful ACK(m) at time $\tau + k\delta$ in runs of B_k, p must (successfully) send m at time $k\delta$ in runs of B_k. Thus

$$\Pr(B_k) = (1-\beta_p)^{k\delta+1}(1-\beta_q)^{\tau+k\delta+1}(2\gamma-\gamma^2)^k(1-\gamma)^2.$$

The first factor reflects the fact that p must have been up at time $k\delta$ (to $\mathsf{send}\ m$) while the second factor reflects the fact that q must have been up at time $\tau + k\delta$ (to $\mathsf{receive}\ m$ and send $\mathrm{ACK}(m)$). The third factor reflects the fact that the previous k attempts have failed: either m was lost or the corresponding $\mathrm{ACK}(m)$ was lost, which occurs with probability $(\gamma + (1-\gamma)\gamma) = 2\gamma - \gamma^2$. The final factor reflects the fact that the $(k+1)$st attempt succeeded: both messages got through. So

$$
\begin{aligned}
\Pr(A) &= \sum_{k=0}^{\infty} \Pr(B_k) \\
&= \sum_{k=0}^{\infty} (1-\beta_p)^{k\delta+1}(1-\beta_q)^{\tau+k\delta+1}(2\gamma-\gamma^2)^k(1-\gamma)^2 \\
&= (1-\beta_p)(1-\beta_q)^{\tau+1}(1-\gamma)^2 \sum_{k=0}^{\infty}(1-\beta_p)^{k\delta}(1-\beta_q)^{k\delta}(2\gamma-\gamma^2)^k \\
&= (1-\beta_p)(1-\beta_q)^{\tau+1}(1-\gamma)^2 \frac{1}{1-(1-\beta_p)^{\delta}(1-\beta_q)^{\delta}(2\gamma-\gamma^2)} \\
&= (1-\beta_p)(1-\beta_q)^{\tau+1}(1-\gamma)^2((1+2\gamma)+O(\varepsilon^2)) \\
&= 1-\beta_p-(\tau+1)\beta_q+O(\varepsilon^2) \\
&\approx 1.
\end{aligned}
$$

We now want to compute $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid A)$. Again, we break $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid A)$ into three pieces,

- $\displaystyle\sum_{k=0}^{2\lceil\frac{2\tau}{\delta}\rceil-1} k\Pr(\#\text{-}\mathsf{send} = k \mid A)$,
- $2\left\lceil\frac{2\tau}{\delta}\right\rceil \Pr(\#\text{-}\mathsf{send} = 2\left\lceil\frac{2\tau}{\delta}\right\rceil \mid A)$, and
- $\displaystyle\sum_{k=2\lceil\frac{2\tau}{\delta}\rceil+1}^{\infty} k\Pr(\#\text{-}\mathsf{send} = k \mid A)$,

and compute each part separately.

Note that $\Pr(\#\text{-}\mathsf{send} = k \mid A) \le \beta_p + \beta_q + \gamma$ for $k < 2\left\lceil\frac{2\tau}{\delta}\right\rceil$, since either a process crashed or a message is lost. Thus the first part is no more than $2\left\lceil\frac{2\tau}{\delta}\right\rceil(\beta_p+\beta_q+\gamma) \approx 0$, so we may ignore it. For the second part, we have

$$\Pr(\#\text{-}\mathsf{send} = 2\left\lceil\frac{2\tau}{\delta}\right\rceil \mid A) \ge (1-\beta_p)^{2\tau+1}(1-\beta_q)^{\lceil\frac{2\tau}{\delta}\rceil\delta+\tau+1}(1-\gamma)^{\lceil\frac{2\tau}{\delta}\rceil+1} \approx 1,$$

since if p is up at time 2τ, q is up at time $\left\lceil\frac{2\tau}{\delta}\right\rceil\delta+\tau$, all of p's sends got through, and q's first $\mathrm{ACK}(m)$ got through, then $\#\text{-}\mathsf{send} = 2\left\lceil\frac{2\tau}{\delta}\right\rceil$; thus the second part is $\approx 2\left\lceil\frac{2\tau}{\delta}\right\rceil$. We now turn our attention to the last part.

Note that p sends at least half the messages in every run r (whether $r \in A$ or $r \in \overline{A}$). Note also that, after the first successful attempt (that is, after the first message sent by p which is $\mathsf{received}$ by q whose corresponding acknowledgment is not lost by the link), p will send at most $\left\lceil\frac{2\tau}{\delta}\right\rceil$ messages, since p would stop $\mathsf{sending}$ 2τ time units after the first successful attempt (either because p $\mathsf{received}$ $\mathrm{ACK}(m)$ or p crashed). Combining the above two observations, we see that if $\#\text{-}\mathsf{send}(r) = 2\left\lceil\frac{2\tau}{\delta}\right\rceil + k$ for $k > 0$, then p must have sent at least $\left\lceil\frac{2\tau}{\delta}\right\rceil + \left\lceil\frac{k}{2}\right\rceil$ messages and there are at least $\left\lceil\frac{k}{2}\right\rceil$ unsuccessful attempts in r. Thus, $\Pr(\#\text{-}\mathsf{send} = 2\left\lceil\frac{2\tau}{\delta}\right\rceil + k \mid A) \le (2\gamma-\gamma^2)^{\lceil\frac{k}{2}\rceil}$. So we have

$$\begin{aligned}
\sum_{k=2\lceil\frac{2\tau}{\delta}\rceil+1}^{\infty} k\Pr(\#\text{-}\mathsf{send} = k \mid A) &\le \sum_{k=2\lceil\frac{2\tau}{\delta}\rceil+1}^{\infty} k(2\gamma-\gamma^2)^{\lceil\frac{k}{2}\rceil} \\
&= \sum_{k=\lceil\frac{2\tau}{\delta}\rceil}^{\infty} ((2k+1)+(2k+2))(2\gamma-\gamma^2)^{k+1} \\
&= \sum_{k=\lceil\frac{2\tau}{\delta}\rceil}^{\infty} (4k+3)(2\gamma-\gamma^2)^{k+1} \\
&\approx 0.
\end{aligned}$$

So we may ignore the last part as well. Thus $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid A) \approx 2\left\lceil\frac{2\tau}{\delta}\right\rceil$. Since $\Pr(A) \approx 1$, we have $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid A)\Pr(A) \approx 2\left\lceil\frac{2\tau}{\delta}\right\rceil$.

We now focus on $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid \overline{A})\Pr(\overline{A})$. Recall that for $r \in \overline{A}$, q fails to successfully send $\text{ACK}(m)$ in r. Consider the following three sets (which is a partition of the set of all runs):

- $C_1 = \{r : p \text{ crashes at time } 0 \text{ in } r\}$,
- $C_2 = \{r : p \text{ does not crash at time } 0 \text{ and } q \text{ crashes at or before time } \tau \text{ in } r\}$, and
- $C_3 = \{r : p \text{ does not crash at time } 0 \text{ and } q \text{ does not crash at or before time } \tau \text{ in } r\}$.

We now show that these are their probabilities:

- $\Pr(C_1 \cap \overline{A}) = \beta_p$,
- $\Pr(C_2 \cap \overline{A}) = (1-\beta_p)(1-(1-\beta_q)^{\tau+1}) = (\tau+1)\beta_q + O(\varepsilon^2)$, and
- $\Pr(C_3 \cap \overline{A}) = O(\varepsilon^2)$.

First note that $\Pr(C_1) = \beta_p$ and $\Pr(C_2) = (1-\beta_p)(1-(1-\beta_q)^{\tau+1}) = (\tau+1)\beta_q + O(\varepsilon^2)$. Furthermore, $C_1 \cup C_2 \subseteq \overline{A}$, since if $r \in C_1 \cup C_2$, q does not send $\text{ACK}(m)$ successfully before crashing. Thus $\Pr(C_1 \cap \overline{A}) = \beta_p$ and $\Pr(C_2 \cap \overline{A}) = (\tau+1)\beta_q + O(\varepsilon^2)$. Since, as we showed earlier, $\Pr(A) = 1 - \beta_p - (\tau+1)\beta_q + O(\varepsilon^2)$, it also follows that $\Pr(C_3 \cap \overline{A}) = O(\varepsilon^2)$.

Now that we have $\Pr(C_i \cap \overline{A})$, let us turn to $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid C_i \cap \overline{A})$. Note that for $r \in \overline{A}$, p will send messages until it crashes. For $r \in C_1$, p crashes immediately, so $\#\text{-}\mathsf{send}(r) = 0$ for $r \in C_1$. For $r \in C_2$, q crashes before it can possibly send any messages, so all the messages are sent by p. Thus

$$\Pr(\#\text{-}\mathsf{send} = k \mid C_2) = (1-\beta_p)^{(k-1)\delta+1}(1-(1-\beta_p)^\delta),$$

since p must be up at time $(k-1)\delta$ and crash before time $k\delta$ to send m exactly k times. So

$$\begin{aligned}
\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid C_2 \cap \overline{A}) &= \sum_{k=1}^{\infty} k(1-\beta_p)^{(k-1)\delta+1}(1-(1-\beta_p)^\delta) \\
&= \frac{(1-\beta_p)(1-(1-\beta_p)^\delta)}{(1-\beta_p)^\delta}\sum_{k=1}^{\infty} k((1-\beta_p)^\delta)^k \\
&= \frac{(1-\beta_p)(1-(1-\beta_p)^\delta)}{(1-\beta_p)^\delta}\frac{(1-\beta_p)^\delta}{(1-(1-\beta_p)^\delta)^2} \\
&= \frac{1-\beta_p}{1-(1-\beta_p)^\delta} \\
&= \frac{1}{\delta\beta_p} + O(1).
\end{aligned}$$

The $O(1)$ term is there because $\left|\frac{1}{\delta\beta_p} - \frac{1-\beta_p}{1-(1-\beta_p)^\delta}\right| = \left|\frac{1}{\delta\beta_p} - \frac{1-\beta_p}{\delta\beta_p+O(\varepsilon^2)}\right| = \left|\frac{O(\varepsilon^2)}{(\delta\beta_p)^2+O(\varepsilon^3)}\right|$, which is $O(1)$, since we assumed that β_p is $\Theta(\varepsilon)$ for this proposition.

For $r \in C_3 \cap \overline{A}$, q might send messages (none of which, however, will get through). Let $E_k = \{r \in C_3 \cap \overline{A} : p$ crashes at time $k\}$. We have $\Pr(E_k) \le (1-\beta_p)^k\beta_p$. Furthermore, $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid E_k) \le 2\left\lceil\frac{k}{\delta}\right\rceil$, since p sends $\left\lceil\frac{k}{\delta}\right\rceil$ messages in E_k and q sends at most that many messages. So we have

$$\begin{aligned}
\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid C_3 \cap \overline{A}) &= \sum_{k=1}^{\infty} \mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid E_k)\Pr(E_k) \\
&\le \sum_{k=1}^{\infty} 2\left\lceil\frac{k}{\delta}\right\rceil(1-\beta_p)^k\beta_p \\
&\le \sum_{k=1}^{\infty} 2\left(\frac{k}{\delta}+1\right)(1-\beta_p)^k\beta_p \\
&= \frac{2\beta_p}{\delta}\sum_{k=1}^{\infty} k(1-\beta_p)^k + 2\beta_p\sum_{k=1}^{\infty}(1-\beta_p)^k \\
&= \frac{2\beta_p}{\delta}\frac{1-\beta_p}{\beta_p^2} + 2(1-\beta_p) \\
&= \frac{2}{\delta\beta_p} + O(1).
\end{aligned}$$

Since we assumed that β_p is $\Theta(\varepsilon)$, $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid C_3 \cap \overline{A})\Pr(C_3 \cap \overline{A}) = O(\varepsilon)$. Recall that $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid C_1 \cap \overline{A}) = 0$, so

$$\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid \overline{A})\Pr(\overline{A}) \approx \mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send} \mid C_2 \cap \overline{A})\Pr(C_2 \cap \overline{A}) \approx \frac{(\tau+1)\beta_q}{\delta\beta_p}.$$

This gives us $\mathbf{E}^{\mathsf{SR}_s^\delta}(\#\text{-}\mathsf{send}) \approx \frac{(\tau+1)\beta_q}{\delta\beta_p} + 2\left\lceil\frac{2\tau}{\delta}\right\rceil$ as desired.

The reasoning for the SR_r^δ case is similar to the SR_s^δ case. The only major difference is that q cannot possibly finish RECEIVing m before time 2τ. We leave details to the reader. ∎

Theorem 5.1: *Under the cost function $\mathbf{c}^{\text{avg}}$, Protocol SR_{hb} satisfies $\mathbf{S}_3$. Furthermore,* $\mathbf{E}(\mathbf{c}^{\text{avg}}) \approx ((1-\alpha_p)(1-\alpha_q)\lambda + \alpha_p\alpha_q)\left(2\left\lceil\frac{2\tau}{\delta}\right\rceil \mathbf{c}\text{-}\mathsf{send} + \left(\tau + \frac{\delta-1}{2}\right)\mathbf{c}\text{-}wait\right) + \frac{1}{\delta\sigma}\mathbf{c}\text{-}\mathsf{send}$, *where* $0 < \lambda < 1$.

Proof: Roughly speaking, the first summand corresponds to the expected per-invocation cost of the protocol and the second corresponds to the expected per-invocation cost of the heartbeats. To do the analysis carefully, we divide the set of runs into three subsets:

- $F_1 = \{r :$ one process is correct and the other eventually crashes in $r\}$,
- $F_2 = \{r :$ both processes are correct in $r\}$, and
- $F_3 = \{r :$ both processes eventually crash in $r\}$.

These are their probabilities:

- $\Pr(F_1) = \alpha_p(1-\alpha_q) + \alpha_q(1-\alpha_p)$,
- $\Pr(F_2) = \alpha_p\alpha_q$, and
- $\Pr(F_3) = (1-\alpha_p)(1-\alpha_q)$.

For $r \in F_1$, we expect the lone correct process to invoke SR_{hb} infinitely often. All but finitely many of these invocations will take place after the other process crashed. Thus the average cost of an invocation in r will be 0. For $r \in F_2$, on the other hand, both processes are expected to invoke SR_{hb} infinitely often and the average cost of the invocation in r is expected to be close to the expected cost of a single invocation of SR_{hb}. The computation of the expected cost of an invocation in a run in F_3 is more delicate. We now examine the details.

Let G_1 be the subset of F_1 consisting of runs r in which the correct process tries to invoke the protocol infinitely often. Clearly $\Pr(G_1 \mid F_1) = 1$, since the protocol is invoked with probability σ at each time unit. Moreover, for each run $r \in G_1$, we have

$$\lim_{t\to\infty} \frac{\mathbf{c}\text{-}\mathsf{SR}(r,t)}{\#\text{-}\mathsf{SR}(r,t)+1} = 0,$$

since there are only finitely many complete invocations with non-zero cost and there are infinitely many complete invocations. Thus, $\mathbf{E}(\mathbf{c}^{\text{avg}} \mid F_1) = 0$.

Let G_2 be the subset of F_2 where there are infinitely many invocations of SR_{hb}. Clearly $\Pr(G_2 \mid F_2) = 1$. Let $Z = 2\left\lceil \frac{2\tau}{\delta} \right\rceil \mathbf{c}\text{-}\mathsf{send} + \left(\tau + \frac{\delta-1}{2}\right) \mathbf{c}\text{-}\mathit{wait}$. By the Law of Large Numbers, for almost all runs r of G_2, the analysis of Proposition 3.4 shows that

$$\lim_{t\to\infty} \frac{\mathbf{c}\text{-}\mathsf{SR}(r,t)}{\#\text{-}\mathsf{SR}(r,t)+1} \approx Z.$$

(Note that we have $\tau + \frac{\delta-1}{2}$ instead of 2τ as in Theorem 4.1. This is because in the current setting, the expected amount of time elapsed between the start of an invocation and the arrival of the first HBMSG is $\frac{\delta-1}{2}$. In the setting of Theorem 4.1, however, the first HBMSG cannot arrive until time τ, since the invocation starts at time 0 and the first HBMSG is sent at time 0. Note that in both cases, the expected time of waiting is τ plus the expected time elapsed between the start of the invocation and the arrival of the next HBMSG.) Thus $\Pr(\mathbf{c}^{\text{avg}}(r) \approx Z \mid F_2) = 1$.

We now turn our attention to F_3. Let $F_3(t_1, t_2, i_1, i_2, i_3)$ be a subset of F_3 with the following properties:

- the first crash in r happens at time t_1,
- the second crash in r happens at time t_2,
- the number of invocations starting before time $t_1 - 3\tau - \delta$ is i_1,
- the number of invocations starting between times $t_1 - 3\tau - \delta$ and $t_1 + \tau$ is i_2, and
- the number of invocations starting after time $t_1 + \tau$ is i_3.

It is clear that each of these sets are measurable. (Some of them are empty, so they will have probability 0; we could introduce restrictions to rule out the empty ones, but leaving them in is not a problem.)

Suppose $F_3(t_1, t_2, i_1, i_2, i_3)$ is not empty. Then

$$\mathbf{E}(\mathbf{c}^{\text{avg}} \mid F_3(t_1, t_2, i_1, i_2, i_3)) \approx \frac{i_1 + \kappa(t_1, t_2, i_1, i_2, i_3) i_2}{i_1 + i_2 + i_3 + 1} Z,$$

where $0 < \kappa(t_1, t_2, i_1, i_2, i_3) < 1$. Roughly speaking, the expected cost of an invocation in the first group is Z, since if no messages are lost (which happens with probability ≈ 1), the number of messages sent is exactly $2\left\lceil\frac{2\tau}{\delta}\right\rceil$ and the time of waiting is between τ and $\tau+\delta-1$, depending on when the first HBMSG arrives after the invocation starts. If no messages are lost, a HBMSG is $\mathsf{received}$ every δ time units, so the wait for a HBMSG is $\frac{\delta-1}{2}$ on average. Thus the first group of invocations contribute $i_1 Z$ to $\mathbf{c}\text{-}\mathsf{SR}(r)$, on average. As for the second group, they contribute something less than $i_2 Z$ to $\mathbf{c}\text{-}\mathsf{SR}(r)$ on average; in many of these invocation, the first process crash (which happens at most $3\tau+\delta$ after the beginning of an invocation in the second group) may reduce the time of waiting or the number of messages sent. That is why we have a multiplicative constant $\kappa(t_1,t_2,i_1,i_2,i_3)$ in front of i_2. The last group of invocations all have zero cost, since by the time they started, the surviving process (which must be the invoker) will never $\mathsf{receive}$ any new HBMSGs from the crashed process; so the time of waiting and the number of messages sent are both zero.

Thus we have

$$\begin{aligned}\mathbf{E}(\mathbf{c}^{\text{avg}} \mid F_3) &= \sum_{t_1,t_2,i_1,i_2,i_3} \mathbf{E}(\mathbf{c}^{\text{avg}} \mid F_3(t_1,t_2,i_1,i_2,i_3))\Pr(F_3(t_1,t_2,i_1,i_2,i_3)) \\ &\approx Z \sum_{t_1,t_2,i_1,i_2,i_3} \frac{i_1+\kappa(t_1,t_2,i_1,i_2,i_3)i_2}{i_1+i_2+i_3+1}\Pr(F_3(t_1,t_2,i_1,i_2,i_3)).\end{aligned}$$

Let

$$\lambda = \sum_{t_1,t_2,i_1,i_2,i_3} \frac{i_1+\kappa(t_1,t_2,i_1,i_2,i_3)i_2}{i_1+i_2+i_3+1}\Pr(F_3(t_1,t_2,i_1,i_2,i_3)).$$

Clearly $\lambda<1$ and $\mathbf{E}(\mathbf{c}^{\text{avg}} \mid F_3) \approx \lambda Z$, as desired.

Now we turn to the expected heartbeat costs per invocation. Each process will send a HBMSG every δ time units for as long as it is up. So if in r a process is up at time t, then it sent $\left\lceil\frac{t}{\delta}\right\rceil$ HBMSGs in r up to time t. Suppose $r\in F_2$. Then, $\#\text{-HBMSG}(r,t) = 2\left\lceil\frac{t}{\delta}\right\rceil$, and by the Law of Large Numbers, for all $\eta>0$,

$$\Pr\left(\lim_{t\to\infty} |\#\text{-}\mathsf{SR}(r,t) - 2t\sigma| \le \eta t \;\middle|\; F_2\right) = 1.$$

Thus,

$$\Pr\left(\lim_{t\to\infty} \frac{\#\text{-HBMSG}(r,t)}{\#\text{-}\mathsf{SR}(r,t)+1} = \frac{1}{\delta\sigma} \;\middle|\; F_2\right) = 1.$$

Next, suppose $r\in F_1$. Then one of the processes will send only finitely many HBMSGs and invoke SR_{hb} finitely often. Thus after the crash, we have

$$\frac{\#\text{-HBMSG}(r,t)}{\#\text{-}\mathsf{SR}(r,t)+1} = \frac{\left\lceil\frac{t}{\delta}\right\rceil + H}{I_2+I_1+1},$$

where H is the number of times the crashed process sends HBMSG in r, I_1 is the number of times the crashed process invoked SR_{hb} in r, and I_2 is the number of times the live process invoked SR_{hb} in r. For all $\eta>0$, we have that

$$\Pr\left(\lim_{t\to\infty} |I_2 - t\sigma| \le \eta t \;\middle|\; F_1\right) = 1.$$

Thus,

$$\Pr\left(\lim_{t\to\infty} \frac{\#\text{-HBMSG}(r,t)}{\#\text{-}\mathsf{SR}(r,t)+1} = \frac{1}{\delta\sigma} \;\middle|\; F_1\right) = 1.$$

Finally, consider the set F_3, where both processes crash. Again, the situation here is more complicated, since there are only finitely many complete invocations and HBMSGs in each run, so we cannot resort to the Law of Large Numbers. Let $F_3(j,k)$ be the set of runs where p crashes at time j and q crashes at time k. Clearly $\Pr(F_3(j,k) \mid F_3) = (1-\beta_p)^j(1-\beta_q)^k\beta_p\beta_q$ and the number of heartbeats sent in runs of $F_3(j,k)$ is $\left\lceil \frac{j}{\delta} \right\rceil + \left\lceil \frac{k}{\delta} \right\rceil$. Let $\#\text{-}\textsc{hbmsg}^{\text{avg}}(r) = \lim_{t\to\infty} \frac{\#\text{-HBMSG}(r,t)}{\#\text{-}\mathsf{SR}(r,t)+1}$. Observe that

$$\begin{aligned}
\mathbf{E}(\#\text{-}\textsc{hbmsg}^{\text{avg}} \mid F_3(j,k)) &= \sum_{i=0}^{j+k} \frac{\left\lceil \frac{j}{\delta} \right\rceil + \left\lceil \frac{k}{\delta} \right\rceil}{i+1} \sigma^i (1-\sigma)^{j+k-i} \binom{j+k}{i} \\
&= \frac{\left\lceil \frac{j}{\delta} \right\rceil + \left\lceil \frac{k}{\delta} \right\rceil}{\sigma(j+k+1)} \sum_{i=0}^{j+k} \sigma^{i+1}(1-\sigma)^{j+k-i} \binom{j+k+1}{i+1} \\
&= \frac{\left\lceil \frac{j}{\delta} \right\rceil + \left\lceil \frac{k}{\delta} \right\rceil}{\sigma(j+k+1)} \sum_{i=1}^{j+k+1} \sigma^{i}(1-\sigma)^{j+k+1-i} \binom{j+k+1}{i} \\
&= \frac{\left\lceil \frac{j}{\delta} \right\rceil + \left\lceil \frac{k}{\delta} \right\rceil}{\sigma(j+k+1)} (1-(1-\sigma)^{j+k+1}).
\end{aligned}$$

Thus,

$$\begin{aligned}
\mathbf{E}(\#\text{-}\textsc{hbmsg}^{\text{avg}} \mid F_3) &= \sum_{j,k} \mathbf{E}(\#\text{-}\textsc{hbmsg}^{\text{avg}} \mid F_3(j,k)) \Pr(F_3(j,k)) \\
&= \sum_{j,k} \frac{\left\lceil \frac{j}{\delta} \right\rceil + \left\lceil \frac{k}{\delta} \right\rceil}{\sigma(j+k+1)} (1-(1-\sigma)^{j+k+1})(1-\beta_p)^j(1-\beta_q)^k\beta_p\beta_q \\
&= \sum_{j,k} \frac{\left\lceil \frac{j}{\delta} \right\rceil + \left\lceil \frac{k}{\delta} \right\rceil}{\sigma(j+k+1)} (1-\beta_p)^j(1-\beta_q)^k\beta_p\beta_q \\
&\quad - \sum_{j,k} \frac{\left\lceil \frac{j}{\delta} \right\rceil + \left\lceil \frac{k}{\delta} \right\rceil}{\sigma(j+k+1)} (1-\sigma)^{j+k+1}(1-\beta_p)^j(1-\beta_q)^k\beta_p\beta_q.
\end{aligned}$$

Note that

$$\frac{\left\lceil \frac{j}{\delta} \right\rceil + \left\lceil \frac{k}{\delta} \right\rceil}{\sigma(j+k+1)} < L_1$$

for some constant L_1 (roughly $\frac{1}{\delta\sigma}$). Thus the second summand above is bounded above by

$$\begin{aligned}
L_1\beta_p\beta_q(1-\sigma)\sum_{j,k}((1-\sigma)(1-\beta_p))^j((1-\sigma)(1-\beta_q))^k &= \frac{L_1\beta_p\beta_q(1-\sigma)}{(\sigma+\beta_p-\sigma\beta_p)(\sigma+\beta_q-\sigma\beta_q)} \\
&\approx \frac{L_1\beta_p\beta_q(1-\sigma)}{\sigma^2},
\end{aligned}$$

which is $O(\varepsilon^2)$. Thus we can ignore the second summand. Taking $L(j,k) = \left\lceil \frac{j}{\delta} \right\rceil + \left\lceil \frac{k}{\delta} \right\rceil - \frac{j+k+1}{\delta}$, we get that

$$\begin{aligned}
\mathbf{E}(\#\text{-}\textsc{hbmsg}^{\text{avg}} \mid F_3) &\approx \sum_{j,k} \frac{\left\lceil \frac{j}{\delta} \right\rceil + \left\lceil \frac{k}{\delta} \right\rceil}{\sigma(j+k+1)} (1-\beta_p)^j (1-\beta_q)^k \beta_p \beta_q \\
&= \sum_{j,k} \frac{1}{\sigma\delta} (1-\beta_p)^j (1-\beta_q)^k \beta_p \beta_q \\
&\quad + \sum_{j,k} \frac{L(j,k)}{\sigma(j+k+1)} (1-\beta_p)^j (1-\beta_q)^k \beta_p \beta_q \\
&= \frac{1}{\sigma\delta} + \frac{1}{\sigma} \sum_{j,k} \frac{L(j,k)}{j+k+1} (1-\beta_p)^j (1-\beta_q)^k \beta_p \beta_q .
\end{aligned}$$

It clearly suffices to show that the second summand above is $O(\varepsilon)$. Note that $\frac{1}{j+k+1} < \sqrt{\beta_p}$ if $j > \frac{1}{\sqrt{\beta_p}}$; similarly, $\frac{1}{j+k+1} < \sqrt{\beta_q}$ if $k > \frac{1}{\sqrt{\beta_q}}$. Finally, it is clear that $\frac{1}{j+k+1} \leq 1$ for all $j,k \geq 0$. Call the second summand above S. Since $L(j,k) < 2$, we have that

$$\begin{aligned}
\sigma S &\leq 2\sqrt{\beta_p} \sum_{j > \frac{1}{\sqrt{\beta_p}}} \sum_k (1-\beta_p)^j (1-\beta_q)^k \beta_p \beta_q \\
&\quad + 2\sqrt{\beta_q} \sum_j \sum_{k > \frac{1}{\sqrt{\beta_q}}} (1-\beta_p)^j (1-\beta_q)^k \beta_p \beta_q \\
&\quad + \sum_{j \leq \frac{1}{\sqrt{\beta_p}}} \sum_{k \leq \frac{1}{\sqrt{\beta_q}}} 2\beta_p \beta_q \\
&\leq 2(\sqrt{\beta_p} + \sqrt{\beta_q}) + 2\sqrt{\beta_p \beta_q}.
\end{aligned}$$

Since we assumed that $\sqrt{\beta_p}$ and $\sqrt{\beta_q}$ are both $O(\varepsilon)$ for this theorem, the second summand above is $O(\varepsilon)$. Thus, $\mathbf{E}(\#\text{-}\textsc{hbmsg}^{\text{avg}} \mid F_3) \approx \frac{1}{\sigma\delta}$. It follows that $\mathbf{E}(\#\text{-}\textsc{hbmsg}^{\text{avg}}) \approx \frac{1}{\sigma\delta}$, as desired. ∎

References

[ACT97] M. K. Aguilera, W. Chen, and S. Toueg. Heartbeat: a timeout-free failure detector for quiescent reliable communication. In *Proceedings of the 11th International Workshop on Distributed Algorithms*, pages 126–140. Springer-Verlag, Berlin/Heidelberg/New York, 1997. A full version is also available as Technical Report 97-1631, Department of Computer Science, Cornell University, 1997.

[BBS98] S. Bajaj, L. Breslau, and S. Shenker. Uniform versus priority dropping for layered video. Submitted for publication, 1998.

[BMPP98] J.-C. Bermond, N. Marlin, D. Peleg, and S. Perennes. Directed virtual path layouts in ATM networks. In *Proceedings of the 12th International Symposium on Distributed Computing*, Lecture Notes in Computer Science, Vol. 1499, pages 75–88, Berlin/Heidelberg/New York, 1998. Springer-Verlag.

[BS98] L. Breslau and S. Shenker. Best-effort versus reservations: A simple comparative analysis. Submitted for publication, 1998.

[CM98] I. Cidon and O. Mokryn. Propagation and leader election in a multihop broadcast environment. In *Proceedings of the 12th International Symposium on Distributed Computing*, Lecture Notes in Computer Science, Vol. 1499, pages 104–118. Springer-Verlag, Berlin/Heidelberg/New York, 1998.

[EHWG98] W. Eberly, L. Higham, and J. Warpechowska-Gruca. Long-lived, fast, waitfree renaming with optimal name space and high throughput (extended abstract). In *Proceedings of the 12th International Symposium on Distributed Computing*, Lecture Notes in Computer Science, Vol. 1499, pages 149–160. Springer-Verlag, Berlin/Heidelberg/New York, 1998.

[FMS98] P. Flocchini, B. Mans, and N. Santoro. Sense of direction in distributed computing. In *Proceedings of the 12th International Symposium on Distributed Computing*, Lecture Notes in Computer Science, Vol. 1499, pages 1–15. Springer-Verlag, Berlin/Heidelberg/New York, 1998.

[Fra86] N. Francez. *Fairness.* Springer-Verlag, Berlin/New York, 1986.

[Kes97] S. Keshav. *An Engineering Approach to Computer Networking.* Professional Computing Series. Addison-Wesley, Reading, Massachusetts, 1997.

[KL95] Y. A. Korilis and A. A. Lazar. On the existence of equilibria for noncooperative flow control. *Journal of the Association for Computing Machinery*, 42(3):584–613, May 1995.

[LS98] A. A. Lazar and N. Semret. The PSP auction mechanism for network resource sharing. In *8th International Symposium on Dynamic Games and Applications*, pages 359–365, 1998.

[MHW96] A. R. Mikler, V. Honavar, and J. S.K. Wong. Analysis of utility-theoretic heuristics for intelligent adaptive network routing. In *Proceedings, Thirteenth National Conference on Artificial Intelligence (AAAI '96)*, volume 1, pages 96–101, 1996.

[MIB98] J.-M. Menaud, V. Issarny, and M Banâtre. A new protocol for efficient cooperative transversal web caching. In *Proceedings of the 12th International Symposium on Distributed Computing*, Lecture Notes in Computer Science, Vol. 1499, pages 288–302. Springer-Verlag, Berlin/Heidelberg/New York, 1998.

[Res87] M. D. Resnik. *Choices: An Introduction to Decision Theory.* University of Minnesota Press, Minneapolis, 1987.

[RS89] J. E. Russo and P. J. H. Schoemaker. *Decision Traps: Ten Barriers to Brilliant Decision-Making and How to Overcome Them.* Doubleday, New York, 1989.

[SKS97] A. Silberschatz, H. Korth, and S. Sudarshan. *Database System Concepts.* McGraw-Hill Companies, Inc., New York, third edition, 1997.

[TRAR98] F. J. Torres-Rojas, M. Ahamad, and M. Raynal. Lifetime based consistency protocols for distributed objects. In *Proceedings of the 12th International Symposium on Distributed Computing*, Lecture Notes in Computer Science, Vol. 1499, pages 378–392. Springer-Verlag, Berlin/Heidelberg/New York, 1998.

[YAGW98] C. H. Young, N. B. Abu-Ghazaleh, and P. A. Wilsey. OFC: a distributed fossil-collection algorithm for time-warp. In *Proceedings of the 12th International Symposium on Distributed Computing*, Lecture Notes in Computer Science, Vol. 1499, pages 408–418. Springer-Verlag, Berlin/Heidelberg/New York, 1998.