

Determining a Maximum Tolerated Schedule of a Cytotoxic Agent

Thomas M. Braun,¹ Peter F. Thall² and Zheng Yuan¹

¹Department of Biostatistics, School of Public Health, University of Michigan,
Ann Arbor, MI 48109, U.S.A.

²Department of Biostatistics and Applied Mathematics, Box 447,
The University of Texas, M.D. Anderson Cancer Center, Houston, TX 77030,
U.S.A.

April 2, 2004

SUMMARY. Most phase I clinical trials are designed to determine a maximum tolerated dose (MTD) for a single initial administration or treatment course of an experimental cytotoxic agent. Toxicity usually is defined as the indicator of whether one or more particular adverse events occur within a short time period from the start of therapy. In actual clinical practice, however, physicians often administer the agent to the patient repeatedly and monitor long-term toxicity due to cumulative effects. We propose a new dose-finding method for such settings, based on the time to toxicity, rather than on a binary outcome. The model and method account for the patient's entire sequence of administrations, with the overall hazard of toxicity modelled as the sum of a sequence of hazards, each associated with one administration. The goal is to determine a maximum tolerated schedule, rather than a conventional MTD for only the first course. Data monitoring and decision making are done continuously throughout the trial. We illustrate the method by

the motivating application, an allogeneic bone marrow transplantation trial to determine how long a recombinant human growth factor can be administered as prophylaxis for acute graft-versus-host disease. We also present a simulation study in the context of this trial.

KEY WORDS: Phase I trial; bone marrow transplantation; dose escalation; CRM; maximum tolerated dose; KGF

1. Introduction

Pre-clinical studies have shown that recombinant human keratinocyte growth factor (KGF) markedly reduces chemotherapy- or radiation-induced injury to the mucosal lining of the lower gastro-intestinal tract (Farrell et al., 1998). These pre-clinical results motivate the hypothesis that KGF can shield the gastro-intestinal tract of allogeneic bone marrow transplant (BMT) recipients from the detrimental effects of acute graft-versus-host disease (aGVHD) while preserving the beneficial graft-versus-leukemia effect (Panoskaltsis-Mortari et al., 1998; Krijanovski et al., 1999). However, administration of KGF is not without risk; the most common toxicities related to KGF administration include skin-related events such as rash, flushing, and edema, as well as increases in amylase and lipase, both of which are indicators of pancreas dysfunction.

A Phase I study in colorectal cancer patients found that 40 mg/kg of KGF could be safely administered on each of 3 consecutive days (Meropol et al., 2003). In another recently completed Phase I study of BMT patients, investigators observed minimal toxicity after increasing the daily dose of KGF to 60mg/kg as well as doubling the total days of administration. Specifically, each patient was given 60 mg/kg/day of KGF on each of the two days prior to BMT, as well as

on the day of BMT. The patient then received four days of rest with no KGF, then received KGF for an additional three consecutive days. Thus, KGF was administered using the 10-day schedule (3-days-on/4-days-off/3-days-on), which we denote by (3+, 4-, 3+). Toxicity was monitored for 28 days, motivated by the assumption that any adverse effect due to a single administration of KGF is certain to occur within 18 days. Although one course of KGF using the (3+, 4-, 3+) schedule is considered safe, investigators believed that one course may not be sufficient prophylaxis for aGVHD, since aGVHD may take up to roughly 100 days after BMT to develop. Consequently, it was desired to investigate the safety of multiple courses of KGF, with 4 days of rest between consecutive courses, and focus on toxicity associated with the entire period of therapy. Thus, two courses would consist of the 24-day schedule (3+, 4-, 3+, 4-, 3+, 4-, 3+), and so on.

Conventional phase I studies determine the maximum tolerated dose (MTD) for a single administration or course by characterizing a patient's outcome as a binary variable indicating whether toxicity occurs within a relatively short, pre-specified time period from the start of therapy. The MTD is generally considered the highest dose that does not present a practical limitation to therapy (Storer, 1989; Goodman et al., 1995; Babb et al., 1998). This general approach has seen widespread use in large part because it facilitates adaptive dose-finding methods that use the doses and outcomes of previous patients to define the doses for successive patients. However, conventional dose-finding methods are generally inadequate for trials where the agent is administered repeatedly over time and evaluation of long-term, cumulative effects is important. This inadequacy occurs because conventional approaches (1) base dose-finding on one, initial administration or course of therapy, and (2) require that toxicity be evaluated quickly enough

to allow for adaptive dose assignment. One exception is the TiTE CRM method proposed by Cheung and Chappell (2000), which allows evaluation of long-term toxicity without delaying accrual. The TiTe CRM still requires assumption (1), however, since it does not accommodate settings where treatment continues past the first course.

These considerations motivated us to construct a new type of phase I design. The goal is to evaluate and compare the cumulative toxicities of several treatment schedules, each consisting of a sequence of administration times. Our method uses the patient's time to toxicity as the outcome, with the hazard of toxicity modelled as the sum of a sequence of hazards, each associated with one administration. The goal is to determine the maximum tolerated schedule (MTS) that the patient may receive, based on the risk of toxicity occurring within a specified period of follow-up that includes the maximum schedule being considered. Patient accrual, data monitoring, and outcome-adaptive decision-making are done continuously throughout the trial under a Bayesian formulation. Each time a new patient is accrued, the most recent data are used to evaluate pre-specified criteria that define the optimal schedule, which is assigned to the newly-accrued patient.

Section 2 establishes notation and presents the probability model. Section 3 describes a method for prior elicitation, and Section 4 describes the trial design. Section 5 describes application of the method to the KGF trial, and presents a computer simulation study to assess the design's performance. Section 6 contains concluding remarks.

2. Probability Model for Toxicity

2.1 Preliminary Notation

Let t^* denote any given time, from the start of the trial, at which it is desired to evaluate the data and make a decision, and denote by n^* the number of patients enrolled up to t^* . For each $i = 1, \dots, n^*$, let e_i denote the *study* time when patient i enters the trial, and let $\mathbf{s}_i = \{s_{i,1}, \dots, s_{i,m_i}\}$ denote the successive *patient* times at which the i^{th} patient receives the agent, where $s_{i,1}$ coincides with study entry. Then, in *study* time, patient i receives the agent at times

$$0 < e_i + s_{i,1} < e_i + s_{i,2} < \dots < e_i + s_{i,m_i}.$$

This notation accounts for the possibility that the agent may be administered whenever and as frequently as desired to each patient, and it allows for an arbitrary number of different treatment sequences to be studied in the trial. Moreover, while investigators typically have a fixed number of schedules they wish to study, our formulation accommodates the common circumstance where some of a patient's actual administration times deviate from the scheduled times, for any of a variety of practical reasons.

Our trial design will restrict its focus to k treatment sequences, denoted by $\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(k)}$, such that $\mathbf{s}^{(j)} = (s_1, s_2, \dots, s_{m(j)})$, so that the j^{th} sequence has a total of $m(j)$ administrations at *patient* times $s_{i,1}, \dots, s_{i,m(j)}$ for the i^{th} patient. Furthermore, $\mathbf{s}^{(j)}$ is a subsequence of $\mathbf{s}^{(j+1)}$ for each $j = 1, \dots, k-1$, so that the duration of therapy increases with j and $m(1) < m(2) < \dots < m(k)$. In the KGF trial, one course of the (3+, 4-, 3+) schedule corresponds to $\mathbf{s}^{(1)} = (1, 2, 3, 8, 9, 10)$, two courses correspond to $\mathbf{s}^{(2)} = (1, 2, 3, 8, 9, 10, 15, 16, 17, 22, 23, 24) = (\mathbf{s}^{(1)}, \mathbf{s}^{(1)} + 14)$, and so on, with BMT at day 3 in any case.

The statistical design problem is to construct a method for choosing a sequence of administration times, from the set of sequences $\{\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(k)}\}$, for each successive patient. Our method will do this adaptively, with the decision for a newly accrued patient utilizing the current data, \mathcal{D}^* , from all patients entering the study previously. Specifically, we will assume a Bayesian model, summarize the information in terms of the posterior $p(\boldsymbol{\theta} \mid \mathcal{D}^*)$, and use this as a basis for deciding which sequence to assign to the current patient. Although $m(j)$ denotes the scheduled number of administrations for any patient assigned to schedule j , we define m_i to be the index of the last administration received by patient i at interim study time t^* . This notation emphasizes the fact that at t^* , $m_i \leq m(j)$ either due to administrative censoring at t^* or because patient i had toxicity at study time $e_i + s_{i,m_i}$ and thus received no further administrations.

Let Y_i be the actual, although possibly not yet observed, amount of time after study entry at which patient i experiences toxicity. At interim study time t^* , we denote the amount of time that patient i has been observed by

$$Y_i^o = \begin{cases} Y_i & \text{if } e_i + Y_i \leq t^* \\ t^* - e_i & \text{if } e_i + Y_i > t^*, \end{cases}$$

and we define $\delta_i = 1$ if $Y_i^o = Y_i$ and $\delta_i = 0$ if $Y_i^o \neq Y_i$, that is, δ_i is the indicator that patient i has had toxicity by t^* .

We denote by τ the fixed maximum length of follow-up for each patient, which is chosen for clinical reasons and must be large enough to accommodate the longest sequence, $\mathbf{s}^{(k)}$. In the KGF trial, $\tau = 100$ days. A fixed target probability, p_τ , is elicited from the physician and is defined as the targeted threshold for the schedule-specific probabilities $F(\tau \mid \boldsymbol{\theta}, \mathbf{s}, 1)$ of toxicity at any time from enrollment to τ .

2.2 Toxicity Distribution

Consider a single patient, and for simplicity temporarily suppress the index i . We assume that the same dose of the agent is given at each administration. Let $h(u \mid \boldsymbol{\theta})$, $u \geq 0$, be the hazard function attributed to a single administration of the agent, where $\boldsymbol{\theta}$ is a parameter vector. Below, we will consider several specific forms for h . We define the total hazard of toxicity at time t^* for a patient treated with schedule \mathbf{s} to be

$$\lambda(t^* \mid \boldsymbol{\theta}, \mathbf{s}, Y^o) = \sum_{\ell=1}^m h(Y^o - s_\ell \mid \boldsymbol{\theta}), \quad (1)$$

where $h(u) = 0$ if $u < 0$. Under this definition, the patient's risk of toxicity at study time t^* depends upon three quantities: 1) Y^o , the length of time the patient has been on study, 2) m , the number of administrations received up to t^* , and 3) \mathbf{s} , the times at which the treatment was administered. Note that the time of study entry, e , is subsumed by Y^o , and the function h does not change with successive administrations. In the sequel, we will discuss how we can generalize this assumption to obtain a more general model.

The patient's cumulative hazard function (chf) is

$$\begin{aligned} \Lambda(t^* \mid \boldsymbol{\theta}, \mathbf{s}, Y^o) &= \int_0^{Y^o} \sum_{\ell=1}^m h(u - s_\ell \mid \boldsymbol{\theta}) du \\ &= \sum_{\ell=1}^m H(Y^o - s_\ell \mid \boldsymbol{\theta}), \end{aligned} \quad (2)$$

where

$$H(Y^o - s_\ell \mid \boldsymbol{\theta}) = \int_0^{Y^o} h(u - s_\ell \mid \boldsymbol{\theta}) du.$$

From Equation (2), the cumulative distribution function (cdf) of Y is

$$F(t^* \mid \boldsymbol{\theta}, \mathbf{s}, Y^o) = \Pr(Y \leq t^* \mid \boldsymbol{\theta}, \mathbf{s}, Y^o) = 1 - \exp\{-\Lambda(t^* \mid \boldsymbol{\theta}, \mathbf{s}, Y^o)\}, \quad (3)$$

and the probability density function (pdf) of Y is

$$f(t^* | \boldsymbol{\theta}, \mathbf{s}, Y^o) = \lambda(t^* | \boldsymbol{\theta}, \mathbf{s}, Y^o) \exp\{-\Lambda(t^* | \boldsymbol{\theta}, \mathbf{s}, Y^o)\}.$$

2.3 Specifying the Single-Administration Hazard Function

The single-administration hazard $h(u | \boldsymbol{\theta})$ can be quite general, provided that it reasonably reflects the risk of toxicity for the agent under study and is sufficiently tractable to facilitate the computations necessary to implement the schedule-finding method we describe below. Our first assumption is that the hazard of toxicity from a single administration has a finite duration and vanishes to zero within θ_3 days. In the KGF trial, based on the physicians' experience, the hazard vanishes after $\theta_3 = 18$ days. Because we assume that $h(\cdot)$ has finite duration, we cannot model $h(\cdot)$ as the hazard of a typical parametric lifetime distribution, such as the gamma or Weibull, unless $h(\cdot)$ is truncated appropriately. As a simple, practical alternative, we assume that h increases linearly to a maximum and decreases linearly thereafter. Specifically, we define

$$h(u | \boldsymbol{\theta}) = \begin{cases} \theta_2 \frac{u}{\theta_1} & 0 \leq u \leq \theta_1 \\ \theta_2 \frac{\theta_3 - u}{\theta_3 - \theta_1} & \theta_1 < u \leq \theta_3 \\ 0 & u > \theta_3. \end{cases} \quad (4)$$

Thus, $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$, with θ_1 the time at which $h(u | \boldsymbol{\theta})$ reaches its maximum, θ_2 the maximum hazard, and θ_3 the time when the hazard vanishes to zero. Figure 1 illustrates this function. Initially, we assumed the triangular hazard (4) only had two parameters θ_1 and θ_2 , with θ_3 fixed and assumed known. However, we found that our ability to locate the optimal schedule was severely hindered when the actual duration of $h(\cdot)$ was much longer than the assumed value of θ_3 .

[Figure 1 about here.]

Other forms for $h(u \mid \boldsymbol{\theta})$ certainly are possible, depending on the particular application. For example, the Weibull hazard $h_1(u \mid \boldsymbol{\theta}) = \theta_1 \theta_2 (\theta_1 u)^{\theta_2 - 1}$ allows the risk of toxicity to continue indefinitely, with the shape parameter θ_2 determining whether the risk of toxicity increases, decreases, or remains constant over time. We could also generalize our model so that the hazard of toxicity varies for each administration. For example, in the context of our motivating example, one may believe that each administration has a hazard that is proportional to its proximity to BMT. Since each patient has their BMT at time s_3 and the last possible administration would occur 77 days after BMT, an extended version of (1) is $\lambda(t^* \mid \boldsymbol{\theta}, \mathbf{s}, Y^o) = \sum_{\ell} h_{\ell}(Y^o - s_{\ell} \mid \boldsymbol{\theta})$, where $h_{\ell}(Y^o - s_{\ell} \mid \boldsymbol{\theta}) = \alpha(s_{\ell}) h(Y^o - s_{\ell} \mid \boldsymbol{\theta})$, and

$$\alpha(s_{\ell}) = 2 - \left| \frac{s_{\ell} - s_3}{77 - s_3} \right|$$

is an inflation factor that is largest at s_3 and decreases linearly as the time of administration moves away from s_3 .

2.4 Likelihood and Posterior

The design requires adaptive decisions to be made repeatedly throughout the trial. The most recent data at study time t^* collected on patient i , for $i = 1, 2, \dots, n^*$, are $\mathcal{D}_i = (s_i, Y_i^o, \delta_i)$. The optimal treatment sequence assigned to patient $n^* + 1$ who enters the trial at t^* is based on the posterior of $\boldsymbol{\theta}$ given the data available at t^* , which we denote by $\mathcal{D}^* = (t^*, \mathcal{D}_1, \mathcal{D}_2, \dots, \mathcal{D}_{n^*})$. The likelihood at t^* is

$$\mathcal{L}(\mathcal{D}^* \mid \boldsymbol{\theta}) = \prod_{i=1}^{n^*} \{f(t^* \mid \boldsymbol{\theta}, s_i, Y_i^o)\}^{\delta_i} \{1 - F(t^* \mid \boldsymbol{\theta}, s_i, Y_i^o)\}^{1 - \delta_i} \quad (5)$$

Denoting the prior by $p(\boldsymbol{\theta})$, the posterior of $\boldsymbol{\theta}$ is

$$g(\boldsymbol{\theta} \mid \mathcal{D}^*) = \frac{\mathcal{L}(\mathcal{D}^* \mid \boldsymbol{\theta})p(\boldsymbol{\theta})}{\int \mathcal{L}(\mathcal{D}^* \mid \boldsymbol{\theta})p(\boldsymbol{\theta})d\boldsymbol{\theta}}.$$

Because the above integral cannot be obtained analytically under our assumed model, we compute posterior quantities via Markov chain Monte Carlo (MCMC) methods (Robert and Casella, 1999).

3. Establishing Prior Distribution for $\boldsymbol{\theta}$

3.1 Eliciting Prior Parameters from Investigators

In practice, $p(\boldsymbol{\theta})$ must be sufficiently uninformative so that it is dominated by the data, which in turn will allow the schedule-finding algorithm to provide a safe and reliable design. We specify the prior distribution of θ_1 to be conditional upon θ_3 because the range of θ_1 is bounded above by θ_3 . Since no such restriction is necessary for θ_2 , we assume a priori that θ_2 is independent of both θ_1 and θ_3 . Although a number of approaches are possible for eliciting $p(\boldsymbol{\theta})$ from investigators, we factor $p(\boldsymbol{\theta})$ into $p_3(\theta_3)p_1(\theta_1 \mid \theta_3)p_2(\theta_2)$ and derive plausible choices of each of those three factors sequentially as follows.

We first determine the prior distribution for θ_3 by asking the investigator to specify a range $[T_\ell, T_u]$ of plausible values for the duration of the hazard for a single administration, as well as μ_{θ_3} , the average anticipated duration. From this information, we assume θ_3 to have a generalized beta distribution

$$p_3(u) = \frac{(u - T_\ell)^{a_3-1}(T_u - u)^{b_3-1}}{(T_u - T_\ell)^{a_3+b_3-1}B(a_3, b_3)}, \quad T_\ell \leq u \leq T_u,$$

where $B(a, b) = \int_0^1 x^{a-1}(1-x)^{b-1}dx$, $a_3 = k_3(\mu_{\theta_3} - T_\ell)$, and $b_3 = k_3(T_u - \mu_{\theta_3})$. The tuning constant k_3 will scale the values of a_3 and b_3 to modulate the

variability of $p_3(u)$, which will increase as k_3 decreases. For example, in the KGF trial, if the investigators believed that the duration of the hazard could range from 4 to 100 days, but was 18 days on average, $a_3 \propto 14$ and $b_3 \propto 82$. Through preliminary simulations, an appropriate value for k_3 may be determined so that $p_3(u)$ is sufficiently sensitive to the data collected on only the first few patients.

Next, we assume that $\theta_1 | \theta_3$ also follows a generalized beta distribution with domain $[0, \theta_3]$, formally

$$p_1(u|\theta_3) = \frac{u^{a_1-1}(\theta_3 - u)^{b_1-1}}{\theta_3^{a_1+b_1-1} B(a_1, b_1)}, \quad 0 \leq u \leq \theta_3.$$

The parameters a_1 and b_1 could be determined similarly to a_3 and b_3 above. However, a different approach for determining a_1 and b_1 is appropriate when the investigator summarizes the range of plausible values as the interval $m \pm d$. By viewing this interval as a 95% credible interval and assuming approximate symmetry about the conditional mean $E(\theta_1 | \theta_3) = m$, we obtain the equations $\text{var}(\theta_1 | \theta_3) = \theta_3^2 a_1 b_1 / \{(a_1 + b_1)^2 (a_1 + b_1 + 1)\} = d^2/4$ and $m = \theta_3 a_1 / (a_1 + b_1)$, which together yield the prior parameters

$$\begin{aligned} a_1 &= \frac{m}{\theta_3} \left[\frac{4m(\theta_3 - m)}{d^2} - 1 \right] \\ b_1 &= \frac{\theta_3 - m}{\theta_3} \left[\frac{4m(\theta_3 - m)}{d^2} - 1 \right]. \end{aligned}$$

Conventional alternatives to the two approaches given above include eliciting percentiles of the prior distributions and estimating the hyperparameters from available historical data (cf. Robert, 2001, Chapter 3).

In order to determine $p_2(u)$, the investigator must select the schedule, \mathbf{s}^* , from the set $\{\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(k)}\}$ that (s)he believes *a priori* to be the MTS. Recall that each

patient is followed for up to τ days, and p_τ is the targeted probability of toxicity by τ . From the *a priori* hypothesis of the optimal schedule, we determine μ_{θ_2} , the prior mean of θ_2 , as the value for which \mathbf{s}^* satisfies $F(\tau \mid \theta_2 = \mu_{\theta_2}, \theta_1, \theta_3 = \theta_3^*, \mathbf{s}^*, 1) = p_\tau$, where θ_3^* is a fixed value for θ_3 . A fixed value for θ_1 is unnecessary because the triangle in Figure 1 has area $\theta_2\theta_3/2$, making the cumulative hazard (and cumulative probability) of toxicity independent of θ_1 . For example, in the KGF trial, suppose that investigators believed that schedule $\mathbf{s}^* = \mathbf{s}^{(2)}$ was optimal when $\theta_3^* = 18$. Then each administration has a cumulative hazard of $9\mu_{\theta_2}$, and because there are 12 administrations, the entire schedule has a cumulative hazard of $108\mu_{\theta_2}$. As a result, $p_\tau \approx 1 - \exp(-108\mu_{\theta_2})$, yielding $\mu_{\theta_2} = -\log(1 - p_\tau)/108$.

Given that θ_2 is the height of the single-administration hazard, it is qualitatively different from θ_1 and θ_3 , which are durations of time. Therefore, we do not wish to restrict the upper bound of $p_2(u)$, making a generalized beta distribution an inappropriate choice for $p_2(u)$. Instead, we assume θ_2 has a Gamma prior distribution with parameters $a_2 = k_2$ and $b_2 = k_2/\mu_{\theta_2}$, where k_2 is a tuning constant used to modulate $\text{Var}(\theta_2) = \mu_{\theta_2}^2/k_2$.

3.2 Calibrating the Prior Distribution of θ

The ability of the data to dominate the prior distribution of θ is heavily influenced by the variances of $p_3(\theta_3)$, $p_1(\theta_1 \mid \theta_3)$, and $p_2(\theta_2)$, as quantified by the tuning constants k_2 and k_3 , and d , the width of the credibility interval for θ_1 . Thus, we strongly recommend that investigators carefully evaluate the design's sensitivity to these parameters. Specifically, by simulating the toxicity times of a small number of patients using various values of k_2 , k_3 and d , investigators can examine the design's sensitivity to these parameters thus determine values yielding a design with desirable behavior. By comparing the prior means for θ to their respective

posterior values, we can measure how a small amount of data can dominate the prior distribution. For example, imagine that the prior distribution for θ_3 reflected the belief that toxicities were unlikely after 25 days. By simulating a few patients to have toxicities that occur far beyond 25 days, one can determine whether the prior allows the posterior mean of θ_3 to shift beyond the prior mean and reflect the data appropriately. If not, k_2, k_3 , and d may be calibrated and the exercise repeated until the desired effect is achieved.

The prior variances cannot be made arbitrarily large, as is usually done with Bayesian analyses of large datasets. In any small scale clinical trial using adaptive methods, very little data are available, especially early in the trial. With overly variable (uninformative) prior distributions, the prior probability mass on large variances often cannot be overcome by the small amount of data, depending on the particular model, data structure, and decision making algorithm. In the present setting, unduly large prior variances would severely hinder the algorithm's ability to assign optimal schedules during the trial and select an optimal MTS at the end.

A given prior $p(\boldsymbol{\theta})$ in turn determines prior distributions on the cumulative probabilities of toxicity at each of the schedules, $F(\tau | \boldsymbol{\theta}, \mathbf{s}^{(j)}, 1)$, $j = 1, \dots, k$. It is very useful to examine the prior on $F(\tau | \boldsymbol{\theta}, \mathbf{s}^{(j)}, 1)$ for each $s^{(j)}$. This may be done by simply sampling from $p(\boldsymbol{\theta})$ and examining the resulting distribution of $F(\tau | \boldsymbol{\theta}, \mathbf{s}^{(j)}, 1)$ for each $s^{(j)}$. This may be used as a diagnostic tool to determine whether a given prior is producing pathological behavior by placing too much probability mass near 0 or 1 for a given $F(\tau | \boldsymbol{\theta}, \mathbf{s}^{(j)}, 1)$. The point is that, in terms of decision making, it is not $p(\boldsymbol{\theta})$ or $p(\boldsymbol{\theta} | \text{data})$ *per se* that matter, but rather the consequent distributions of the decision criteria.

For example, based on the examples discussed previously, suppose that one

assumes *a priori* that θ_3 varies uniformly over $[4, 100]$ days, with a mean of 18 days, while θ_1 is within 2 ± 2 days after administration. Suppose that the targeted threshold $p_\tau = 0.20$ for $\tau = 100$ days and Schedule 2 is optimal, which motivates our assumption that θ_2 has a Gamma distribution with mean 0.0021. Based upon a preliminary calibration as discussed above, we defined $k_2 = k_3 = 1$. Figure 2 displays the resulting priors of $F(100 | \boldsymbol{\theta}, \mathbf{s}^{(j)}, 1)$ for $j = 1, \dots, 6$. As desired, $E\{F(100 | \boldsymbol{\theta}, \mathbf{s}^{(j)}, 1)\}$ is closest to 0.20 for $j = 2$, and the dispersion of $F(100 | \boldsymbol{\theta}, \mathbf{s}^{(j)}, 1)$ increases with j due the cumulative nature of the schedules.

[Figure 2 about here.]

4. Trial Conduct

A maximum of N patients are enrolled in the trial, with each assigned a treatment administration sequence upon arrival in the clinic. The first patient is assigned the shortest sequence under consideration, $\mathbf{s}^{(1)}$. Each patient is followed for up to τ days, with treatment terminated if toxicity is observed prior to τ . Given a desired threshold p_τ for $F(\tau | \boldsymbol{\theta}, \mathbf{s}^{(j)}, 1)$, we will consider two alternative criteria for choosing each patient's sequence.

Criterion 1: At time t^* , for each $j = 1, \dots, k$, compute

$$F_j^*(\tau) = E\{F(\tau | \boldsymbol{\theta}, \mathbf{s}^{(j)}, 1) | \mathcal{D}^*\}. \quad (6)$$

The best sequence is defined as that having $F_j^*(\tau)$ closest to p_τ , that is, minimizing $|F_j^*(\tau) - p_\tau|$. This criterion, as a function of treatment sequence, is analogous to the CRM criterion (O'Quigley et al., 1990) based upon the posterior mean probability of the more usual binary toxicity, as a function of dose.

Criterion 2: At time t^* , for each $j = 1, \dots, k$, compute

$$\phi_j(\tau) = \Pr\{F(\tau \mid \boldsymbol{\theta}, \mathbf{s}^{(j)}, 1) > p_\tau \mid \mathcal{D}^*\}. \quad (7)$$

Since $F(\tau \mid \boldsymbol{\theta}, \mathbf{s}^{(j)}, 1)$ is monotone increasing in j , it follows that $\phi_1(\tau) \leq \phi_2(\tau) \leq \dots \leq \phi_k(\tau)$. Given a fixed upper limit, \bar{p} , the best sequence is defined as the longest sequence for which $\phi_j(\tau) < \bar{p}$, that is, for which the risk of toxicity is acceptable. The second criterion is similar to the acceptability criteria used by Thall and Russell (1998) in the setting where both efficacy and toxicity are used for dose-finding, and it is also similar to the criterion for overdose control proposed by Babb et al. (1998) when doing dose-finding based on toxicity. Under either Criterion 1 or 2, the best sequence is assigned to patient $n^* + 1$, and the MTS is defined as the best sequence based on the final data at the end of the trial.

To protect patient safety, we impose the additional constraint that only incremental schedule escalation is permitted. Formally, if $\mathbf{s}^{(j)}$ is determined to be the best schedule, the next patient is assigned $\mathbf{s}^{(j)}$ only if each of the schedules $\mathbf{s}^{(1)}, \dots, \mathbf{s}^{(j-1)}$ previously has been assigned to at least M patients, for a predetermined $M \geq 1$. Otherwise, the next patient enrolled is assigned to the longest schedule that meets this criterion. In our application, we have set $M = 1$, although larger values of M are possible. However, an important practical consideration is that larger values of M will slow the speed of schedule escalation, so that the trial will take longer to accrue patients on longer schedules. Schedule de-escalation is permitted per the criterion, without any constraint.

5. Application

5.1 Simulation Study Design

We examined the performance of our study design by simulation in the context of the KGF trial. The investigators decided that $k = 6$ schedules should be investigated, corresponding to 2, 4, 6, 8, 10 or 12 weeks of therapy with KGF. Because aGVHD is defined to occur during the first 100 days after transplant, the maximum period to monitor toxicity was specified to be $\tau = 100$ days. Per the adaptive design, a schedule is specified for each patient, and KGF is discontinued if the patient experiences toxicity before 100 days. The goal is to determine how long a BMT patient can receive KGF as prophylaxis for aGVHD while controlling the risk of toxicity within 100 days to be 20%, on average.

The investigators believed, *a priori*, that the hazard of toxicity for a single administration vanishes after an average of 18 days, with a range of 4 to 100 days. They also believed that θ_1 , the time when the hazard of toxicity from one administration of KGF is largest, should be at most 4 days. Using the methods described in Section 3, we derived parameters a_3 and b_3 for $p(\theta_3)$, and a_1 and b_1 for $p(\theta_1 | \theta_3)$. Based upon results of an earlier trial, the rate of toxicity on the lowest schedule of KGF was very low, and the investigators believed that even 12 weeks of KGF would not cause appreciably more toxicity, leading us to assume that $p(\theta_2)$ was a Gamma distribution with mean 0.0007. Through a detailed sensitivity analysis, we determined that $k_3 = 0.1$ and $k_2 = 0.2$ allowed the data to have adequate influence on the posterior of θ . After studying Criterion 2 with $0.20 \leq \bar{p} \leq 0.80$, we also found that $\bar{p} = 0.60$ worked best. Values of $\bar{p} > 0.60$ tended to make our algorithm unacceptably conservative, i.e. tended to locate the MTS frequently at schedules lower than the MTS, while $\bar{p} < 0.50$ tended to lead our algorithm to

select overly toxic schedules too often. We studied the design with a maximum sample size of 30 patients, which is feasible but sufficient to allow one to determine the optimal schedule with reasonable accuracy. Subject inter-arrival times were assumed to be uniformly distributed from 12 to 16 days.

We considered a variety of possible circumstances, and also investigated the method’s sensitivity to particular aspects of the underlying assumptions. In each simulation, the posteriors were based upon 2000 samples, using a burn-in of 500 samples from the posterior of θ via MCMC. We first examined the design’s performance using both Criterion 1 and Criterion 2 (with $\bar{p} = 0.60$), under each of six scenarios with schedule $s^{(j)}$ optimal under the j th scenario. The times to toxicity were simulated assuming that θ_1 occurs 2 days after administration and that θ_3 is either 18 or 50 days. The value of θ_2 is varied for each scenario to reflect which schedule is optimal; those values are shown in Table 1, which also gives the actual Day 100 probabilities of toxicity for each schedule under each scenario.

[Table 1 about here.]

5.2 *Simulation Results*

Table 2 summarizes the simulation results of the scenarios in terms of (1) how frequently each schedule was selected as the MTS, and (2) the mean number of patients assigned to each schedule. In these simulations, the means of the prior distributions for θ_1 and θ_3 were identical to the true values of θ_1 and θ_3 used to simulate toxicity times, while the mean for the prior distribution of θ_2 was determined by the belief that Schedule 6 is optimal. The simulation results assuming that the maximum hazard duration from a single administration of KGF was either $\theta_3 = 18$ days or 50 days are summarized in Tables 2 and 3, respectively. In interpreting the selection percentages given in these tables, it is important to bear in

mind that, considering the six schedules in turn, the smallest absolute difference between the toxicity probability 0.20 at the optimal schedule and the toxicity probabilities at the neighboring schedules were 0.16, 0.08, 0.06, 0.04, 0.04 and 0.03, respectively, so that the goal of choosing the optimal schedule was successively more difficult going from $s^{(1)}$ to $s^{(6)}$.

When $\theta_3 = 18$ days, Table 2 shows that the algorithm with Criterion 1 selected a MTS within one schedule of the optimum MTS in 99.5%, 97.7%, 84.5%, 66.7%, 64.7%, and 58.7% of simulations for scenarios 1-6, respectively. The algorithm tends to err conservatively, ie. is more likely to select schedules shorter than the MTS over schedules longer than the MTS. In addition, more patients are assigned to the MTS than any of the other schedules in all six scenarios. Similar conclusions are reached when using Criterion 2 with $\bar{p} = 0.60$. The algorithm selected a MTS within one schedule of the optimum MTS in 99.7%, 96.0%, 74.6%, 65.1%, 74.3%, and 72.4% of simulations for scenarios 1-6, respectively. Criterion 2 did slightly better than Criterion 1 when the true MTS was one of the shortest or longest schedules, and Criterion 1 outperformed Criterion 2 when the true MTS was in the middle of all the schedules. Note that θ_1 has no impact upon each course's estimated cumulative probability of toxicity. Instead, θ_1 influences the time at which each patient experiences toxicity and influences the process of escalation and de-escalation of courses during the study. This fact was supported by additional simulations (not shown) in which we found the prior distribution for θ_1 had negligible impact on the results shown in Table 2.

The prior belief was $s^{(6)}$ was optimal and was reflected in the prior mean of θ_2 . Given that the algorithm remained conservative even when the longest schedule was the *a priori* optimum MTS, the first five scenarios supply evidence that the

method reliably shifts the posterior distribution of θ_2 away from a misspecified prior mean, even with the limited sample size of 30 patients.

[Table 2 about here.]

Table 3 gives results analogous to Table 2 when the true duration of the toxicity hazard for each single administration is extended to $\theta_3 = 50$ days. Since $p(\theta_3)$ has a mean of 18 days, here toxicities can occur much later than the prior distribution anticipates. If the prior of θ_3 is not sufficiently flexible, dose escalation will move too rapidly, essentially because the algorithm places insufficient weight on any "late-onset" toxicities when computing the posterior of θ . As a result, one might expect the algorithm to be liberal, *i.e.* likely to identify the MTS at a schedule longer than the optimum MTS, as well as assigning too many patients to overly toxic schedules. However, Table 3 shows that the algorithm is relatively insensitive to the fixed value of θ_3 in all six scenarios. This is due to the fact that $p(\theta_3)$ is sufficiently flexible to accommodate actual $\theta_3 = 18$ or 50. We also found that, with either Criterion 1 or 2, the algorithm is more likely to identify the MTS when the true θ_3 is large, albeit very slightly (simulations not shown). This finding is most apparent for scenarios 5 and 6, where the algorithm correctly identifies the MTS more often when $\theta_3 = 50$ than when $\theta_3 = 18$.

[Table 3 about here.]

6. Concluding Remarks

We have developed a flexible algorithm for identifying the maximum schedule at which a fixed dose of an experimental agent can be administered without causing

undue cumulative toxicity in patients. Although the actual study will assign patients to a finite number of schedules, the data collected from the study may be used to estimate the cumulative toxicity of any schedule.

An alternate study design for this study was recently published by Braun et al. (2003), who treated each schedule as a "dose" and determined the MTS using the TiTE CRM with study-specific modifications. By forcing each schedule to be a dose, patients who received an incomplete schedule were only evaluable up to the point of their last fully completed schedule, however. Furthermore, the "doses" overlapped, leading to some ambiguity as to which "dose" contributed to a late-onset toxicity. By modelling each administration separately, our approach avoids both of these limitations.

Our method may be made more flexible to accommodate a wider range of applications. One important generalization is to model the single administration hazard so that the dose administered at each day can vary both within each patient, as well as between patients. One may also allow the hazard to vary with patient characteristics, so that schedules may be chosen more specifically for each patient. Our model parameterization makes a strong homogeneity assumption about the hazard of toxicity over the sequence of treatment times. Under our model, we could observe nothing other than patients who receive the shortest sequence, $\mathbf{s}^{(1)}$, learn about $\boldsymbol{\theta}$, and make predictions about the risk of toxicity for a patient who is scheduled to receive any $\mathbf{s}^{(j)}$, $j \geq 1$. However, this is a feature of any parametric regression model that is assumed to hold over a given domain of covariates, which in this case are the $\mathbf{s}^{(j)}$'s. If this assumption is considered invalid in a particular setting, then a more general model accounting for inhomogeneity across schedules would be needed. We currently are investigating these generalizations.

ACKNOWLEDGEMENTS

Tom Braun and Zheng Yuan's research was partially supported by the University of Michigan's NIH Cancer Center Support Grant 5-P30-CA-46592. Peter Thall's research was partially supported by NCI grant RO1-CA-83932.

REFERENCES

- Babb, J., Rogatko, A. and Zacks, S. (1998). Cancer Phase I clinical trials: efficient dose escalation with overdose control. *Statistics in Medicine* **17**, 1103–1120.
- Braun, T. M., Levine, J. E. and Ferrara, J. L. M. (2003). Determining a maximum tolerated cumulative dose: Dose reassignment within the titration-crm. *Controlled Clinical Trials* **24**, 669–681.
- Cheung, Y. and Chappell, R. (2000). Sequential designs for phase I clinical trials with late-onset toxicities. *Biometrics* **56**, 1177–1182.
- Farrell, C., Bready, J., Rex, K. and et al (1998). Keratinocyte growth factor protects mice from chemotherapy and radiation-induced gastrointestinal injury and mortality. *Cancer Research* **58**, 933–939.
- Goodman, S., Zahurak, M. and Piantadosi, S. (1995). Some practical improvements in the continual reassessment method for Phase I studies. *Statistics in Medicine* **14**, 1149–1161.
- Krijanovski, O. I., Hill, G. R., Cooke, K. R. and et al (1999). Keratinocyte growth factor separates graft-versus-leukemia effects from graft-versus-host disease. *Blood* **94**, 825–831.
- Meropol, N., Somer, R., Gutheil, J. and et al (2003). Randomized phase I trial of recombinant human keratinocyte growth factor plus chemotherapy: potential

- role as mucosal protectant. *Journal of Clinical Oncology* **21**, 1452–1458.
- O’Quigley, J., Pepe, M. and Fisher, L. (1990). Continual reassessment method: A practical design for Phase I clinical trials in cancer. *Biometrics* **46**, 33–48.
- Panoskaltis-Mortari, A., Lacey, D., Vallera, D. and Blazar, B. (1998). Keratinocyte growth factor administered before conditioning ameliorates graft-versus-host disease after allogeneic bone marrow transplantation in mice. *Blood* **92**, 3960–3967.
- Robert, C. P. (2001). *The Bayesian Choice, 2nd edition*. New York: Springer.
- Robert, C. P. and Casella, G. (1999). *Monte Carlo Statistical Methods*. New York: Springer.
- Storer, B. E. (1989). Design and analysis of Phase I clinical trials. *Biometrics* **45**, 925–937.
- Thall, P. F. and Russell, K. T. (1998). A strategy for dose-finding and safety monitoring based on efficacy and adverse outcomes in phase I/II clinical trials. *Biometrics* **54**, 251–264.




Figure 1. Parametric hazard function for a single administration of an agent.

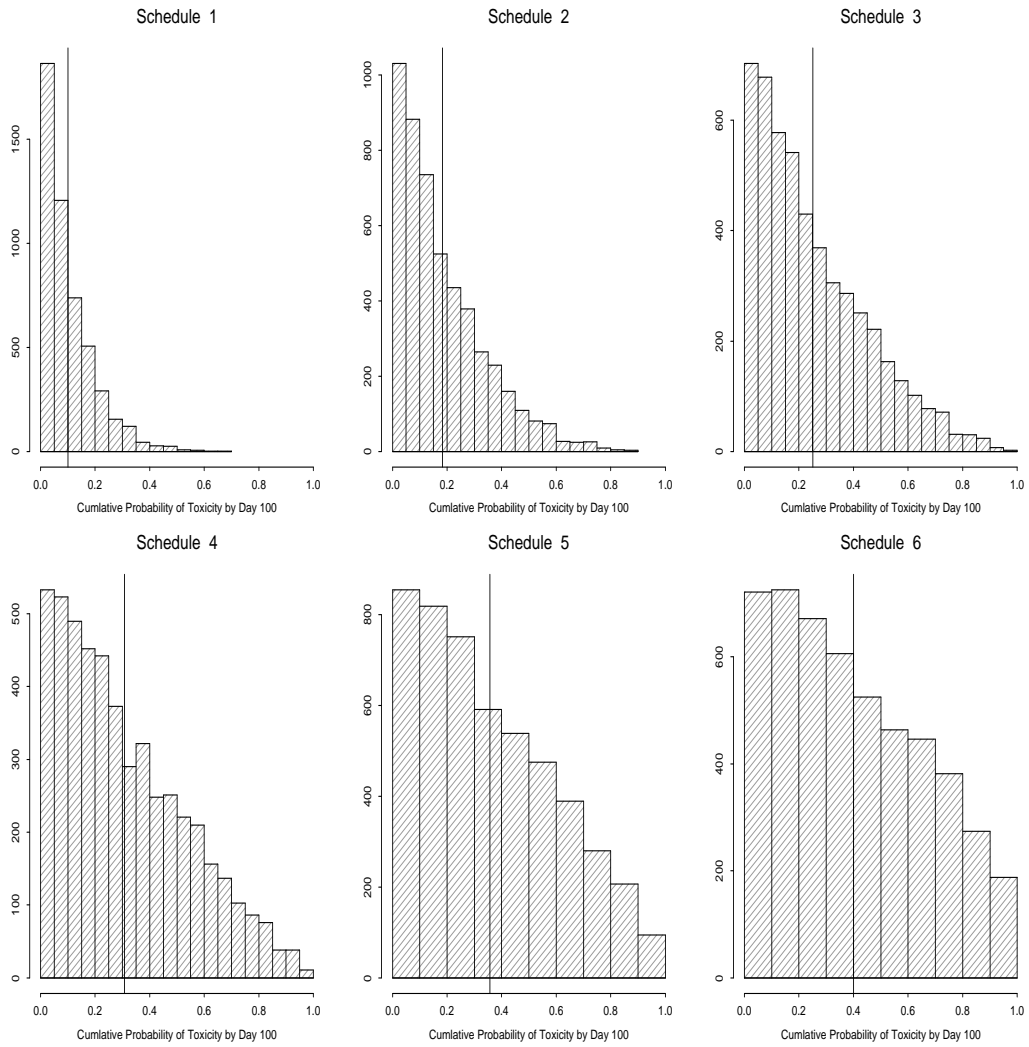


Figure 2. Empirical prior distributions for cumulative probability of toxicity by Day 100 for each schedule. Each histogram is based upon 5000 observations. The solid vertical line represents the mean of the distribution.

Table 1

Probability of toxicity within 100 days for each schedule under each of the simulation scenarios. The hazard of toxicity per administration lasts 18 days for each of the first six rows and 50 days for each of the last six rows.

Duration of Hazard	Optimal # Courses	1000 θ_2	Schedule (Weeks)					
			1(2)	2 (4)	3 (6)	4 (8)	5 (10)	6 (12)
18 days	1	4.13	0.20	0.36	0.49	0.59	0.67	0.74
	2	2.07	0.11	0.20	0.28	0.36	0.43	0.49
	3	1.38	0.07	0.14	0.20	0.26	0.31	0.36
	4	1.03	0.05	0.11	0.15	0.20	0.24	0.28
	5	0.83	0.04	0.09	0.13	0.16	0.20	0.24
	6	0.69	0.04	0.07	0.11	0.14	0.17	0.20
50 days	1	1.49	0.20	0.36	0.49	0.59	0.67	0.72
	2	0.74	0.11	0.20	0.28	0.36	0.42	0.47
	3	0.50	0.07	0.14	0.20	0.26	0.31	0.35
	4	0.37	0.05	0.11	0.15	0.20	0.24	0.27
	5	0.30	0.04	0.09	0.13	0.17	0.20	0.23
	6	0.26	0.04	0.08	0.11	0.15	0.18	0.20

Table 2

Performance of the design with 30 patients, assuming the single administration hazard duration is 18 days. Each entry is the schedule's selection percentage, with number of patients assigned to that schedule given in parentheses. Values within one schedule of the true MTS are given in italics.

Optimal # Courses		Schedule (Weeks)					
		1(2)	2 (4)	3 (6)	4 (8)	5 (10)	6 (12)
Criterion 1	1	86.6 <i>(20.9)</i>	13.3 <i>(6.4)</i>	0.1 (1.4)	0.0 (0.6)	0.0 (0.3)	0.0 (0.3)
	2	27.8 <i>(8.9)</i>	53.5 <i>(11.8)</i>	16.4 <i>(4.8)</i>	2.0 (1.9)	0.2 (1.5)	0.1 (1.1)
	3	7.9 (4.3)	33.0 <i>(7.0)</i>	37.1 <i>(9.4)</i>	14.4 <i>(3.8)</i>	4.9 (3.4)	2.7 (2.1)
	4	1.7 (2.7)	21.0 (4.7)	23.2 <i>(6.6)</i>	32.6 <i>(6.9)</i>	10.9 <i>(6.2)</i>	10.6 (3.0)
	5	0.8 (2.1)	10.8 (4.8)	23.7 (4.7)	24.2 <i>(5.9)</i>	26.3 <i>(9.1)</i>	14.2 <i>(3.4)</i>
	6	0.1 (1.7)	5.7 (3.7)	16.9 (4.5)	18.6 (3.8)	20.1 <i>(5.0)</i>	38.6 <i>(11.3)</i>
Criterion 2 ($\bar{p} = 0.60$)	1	93.3 <i>(21.9)</i>	6.4 <i>(4.6)</i>	0.3 (1.7)	0.0 (0.8)	0.0 (0.5)	0.0 (0.5)
	2	40.1 <i>(8.9)</i>	41.9 <i>(10.4)</i>	14.0 <i>(4.8)</i>	3.0 (2.3)	0.8 (2.2)	0.2 (1.4)
	3	14.2 (5.0)	28.6 <i>(6.5)</i>	31.4 <i>(7.5)</i>	14.6 <i>(5.6)</i>	5.6 (3.0)	5.6 (2.4)
	4	4.7 (2.8)	16.6 (4.7)	20.1 <i>(5.2)</i>	24.8 <i>(7.4)</i>	20.2 <i>(6.0)</i>	13.6 (3.9)
	5	1.9 (2.3)	9.8 (3.2)	14.0 (3.6)	18.4 <i>(5.0)</i>	39.1 <i>(10.8)</i>	16.8 <i>(5.1)</i>
	6	0.4 (1.6)	3.9 (2.6)	9.9 (3.4)	13.4 (3.3)	19.0 <i>(5.3)</i>	53.4 <i>(13.8)</i>

Table 3

Performance of the design with 30 patients, assuming the single administration hazard duration is 50 days. Each entry is the schedule's selection percentage, with number of patients assigned to that schedule given in parentheses. Values within one schedule of the true MTS are given in italics.

		Optimal # Courses	Schedule (# Weeks)					
			1(2)	2 (4)	3 (6)	4 (8)	5 (10)	6 (12)
Criterion 1	1	83.3 (20.4)	11.4 (6.0)	0.3 (1.8)	0.0 (0.8)	0.0 (0.5)	0.0 (0.5)	
	2	25.8 (7.6)	50.9 (10.9)	18.1 (5.5)	4.2 (2.5)	0.9 (2.2)	0.1 (1.3)	
	3	6.3 (3.5)	31.9 (6.9)	34.0 (8.1)	17.9 (4.3)	6.1 (4.7)	3.8 (2.5)	
	4	1.7 (2.4)	13.4 (4.9)	26.7 (5.3)	30.4 (6.5)	15.6 (7.6)	12.2 (3.3)	
	5	0.9 (1.9)	7.9 (3.9)	20.1 (4.6)	21.0 (5.1)	32.2 (11.0)	17.9 (3.5)	
	6	0.3 (1.6)	4.9 (3.4)	14.6 (4.4)	15.6 (3.7)	19.9 (4.4)	44.7 (12.5)	
Criterion 2 ($\bar{p} = 0.60$)	1	96.4 (21.7)	3.4 (4.1)	0.2 (1.9)	0.0 (1.0)	0.0 (0.7)	0.0 (0.6)	
	2	39.2 (8.2)	44.0 (10.1)	11.8 (4.5)	3.1 (3.0)	1.4 (2.6)	0.5 (1.6)	
	3	13.0 (4.6)	29.1 (5.9)	30.7 (7.7)	13.5 (5.3)	6.7 (3.9)	7.0 (2.6)	
	4	5.5 (3.1)	18.7 (4.5)	21.6 (5.8)	23.2 (8.3)	18.8 (4.8)	12.2 (3.5)	
	5	2.6 (2.1)	6.4 (3.1)	15.7 (3.8)	22.6 (5.8)	37.9 (11.0)	14.8 (4.2)	
	6	0.9 (1.8)	5.2 (2.6)	11.8 (3.1)	12.8 (3.3)	15.1 (4.4)	54.2 (14.7)	