Proceedings of the 2011 Winter Simulation Conference
S. Jain, R.R. Creasey, J. Himmelspach, K.P. White, and M. Fu, eds.

COMBINING SIMULATION ALLOCATION AND OPTIMAL SPLITTING FOR RARE-EVENT SIMULATION OPTIMIZATION

Ben Crain

George Mason University
Department of Systems Engineering & Operations Research
Fairfax, VA, USA

Chun-Hung Chen

National Taiwan University
Department of Electrical Engineering
College of EECS
Taipei, TAIWAN

John F. Shortle

George Mason University
Department of Systems Engineering & Operations Research
Fairfax, VA, USA

ABSTRACT

This paper presents research toward generalizing the optimization of the allocation of simulation replications to an arbitrary number of designs, when the problem is to maximize the Probability of Correct Selection among designs, the best design being the one with the smallest probability of a rare event. The simulation technique within each design is an optimized version of the splitting method. An earlier work solved this problem for the special case of two designs. In this paper an alternative two-stage approach is examined in which, at the first stage, allocations are made to the designs by a modified version of the Optimal Computing Budget Allocation. At the second stage the allocation among the splitting levels within each design is optimized. Our approach is shown to work well on a two-tandem queuing model.

1 INTRODUCTION

It is well known that standard Monte Carlo (MC) techniques do not efficiently estimate the probability of rare events. Their inadequacy is exacerbated when simulations, subject to a computational budget constraint, are made over several designs, for the purpose of selecting the design with the smallest rare event probability. This paper explores methods to improve performance by optimizing the allocation of the computational budget among designs and, within each design, among the levels of a fixed effort splitting technique.

The computational budget is defined in terms of time. The problem of allocating this time optimally, in order to maximize the probability of correct selection (the probability of selecting the “best” design, however best is defined) is addressed by the Optimal Computing Budget Constraint (OCBA) (Chen et al. 2000; He et al. 2007; Fu et al. 2007). OCBA implicitly assumes that standard MC is the simulation technique used within each design. It thus suffers, in the context of rare events, from the inefficiencies of MC. Applying OCBA, by itself, to optimize the allocation of the budget among designs offers only slight improvement over equal budget allocations, when MC is used to estimate rare event probabilities within the designs.

There are several approaches to the problem of efficiently estimating rare events, within a single system or design. L’Ecuyer, Demers, and Tuffin (2006), L’Ecuyer et al. (2009), and Asmussen and Glynn

(2007) provide overviews of such approaches. Splitting techniques may be generally characterized as *fixed* splitting or *fixed-effort* splitting. *Fixed* splitting refers to methods in which each run that hits its splitting level launches a fixed number of runs at the next level. See Villen-Altamirano and Villen-Altamirano (2006) for an example of a fixed splitting method. Lagnoux-Renaudie (2008) develops techniques to improve the effectiveness of fixed splitting methods. *Fixed-effort* splitting refers to methods in which the number of runs at each level is pre-specified (L'Ecuyer, Demers, and Tuffin 2006). We adopt a version of the fixed effort splitting method which optimizes the allocation of the budget, for a single design, among its splitting levels (Shortle & Chen 2008; Fischer et al. 2010; Shortle et al. 2011). It is called the Optimal Splitting Technique for Rare Event Simulation, or OSTRE. (The optimization minimizes the variance of the design's probability estimator.) Success at selecting the best of several designs is considerably enhanced by employing OSTRE within each design, even if no attempt is made to optimize the budget allocation between designs. But we need not forego optimization between designs. This paper proposes combining OCBA, which optimizes the allocation among designs, with OSTRE, which optimizes within designs. It is thus a two-stage approach, henceforth called OCBA+OSTRE.

Without dwelling on the details of implementing the optimizing algorithm(s), one feature should be made clear. The optimization problem is posed in the standard form of an objective function to be minimized, subject to a constraint. But in practice that problem is not explicitly solved. The solution is found incrementally by (approximately) satisfying, within the implementation of the algorithm, its first-order optimality conditions, and by stopping the implementation when the budget constraint is reached. The algorithm is implemented in stages. At each stage the task is to make a new (small) allocation to the appropriate design, and then to the appropriate splitting level, in order to nudge the optimality conditions toward equality. The first stage is initialization, where modest allocations are made to all designs and levels to obtain preliminary estimates of all the parameters. Using those estimates the optimality equations are evaluated, and the next (small) allocation made to the design/ level which best moves those equations toward equality. The simulation then advances to the next update, a new allocation is made, etc. The allocations at each stage being small, the updating is frequent and tends, as our testing shows, to establish approximate equality of the optimality equations at the termination of the algorithm.

Let n_{kj} denote the number of runs for design k, at level j, and b_{kj} the average time required for each such run. The decision variables are n_{kj}. At each step of the algorithm the next (small) set of runs will be allocated to one of the splitting levels of one of the designs. Thus the optimization should, theoretically, be over all the n_{kj} at once. We call that the global optimization. It is the gold standard – the "optimal" optimization. This paper proposes, instead, a two-stage optimization (OCBA+OSTRE) in which, at each updating step, an allocation is first made to a design, and then to a level. The first decision (which design?) is made without knowledge of where it will go within the selected design (which level?).

This two-step tango might seem more convoluted than a straightforward global optimization, and it might be mathematically inferior to global optimization, but it promises an attractive alternative. Global optimization becomes intractable as the number of designs/ levels increases. For anything beyond, say, 4 designs and 4 levels global optimization becomes quite ponderous. A two-stage approach promises to be more tractable, at least in higher dimensions, as it is broken into two bites, each more digestible than the whole.

In this paper we derive the optimality conditions for the OCBA allocation between designs, and wed it to the OSTRE allocation within designs, yielding a two-stage method for the simplest case, two designs and two levels. We prove the mathematical equivalence of this two-stage method and global optimization, in this simplest case. We then present tests of the method, on a queuing model, for two and three designs.

2 DERIVATION OF A TWO-STAGE METHOD FOR TWO DESIGNS, TWO LEVELS

The two-stage method may be thought of as an OCBA step, then an OSTRE step. The OCBA step allocates runs between designs. The OSTRE step then allocates the runs received by a design to its splitting levels. This section derives the OCBA allocation, in conjunction with the OSTRE allocation derived in Shortle et al. (2011).

OCBA was originally developed on the assumption of MC within the designs. Its results are not directly applicable, without adjustment, in the context of splitting. We derive the appropriate OCBA, and merge it with OSTRE, for two designs and two splitting levels.

Notation

$p_k =$ probability of rare event for design k, $k = 1, 2$

$\hat{p}_k =$ estimator of p_k

$\sigma_k^2 =$ variance of $\hat{p}_k$

$N_k =$ number of simulation replications (runs) for design k

$n_{kj} =$ number of runs in level j, design k

$H_{kj} =$ number of hits in level j, starting from level $j-1$, in n_{kj} trials, in design k

$p_{kj} =$ probability of hitting level j, starting from level $j-1$, in design k

$\hat{p}_{kj} =$ estimator of p_{kj}

$b_k =$ average time consumed by one run in design k

$b_{kj} =$ average time for one run in level j, design k

$T =$ total computational budget

$\Phi =$ cdf of the standard normal distribution

Without loss of generality we assume $p_1 < p_2$. To maximize the probability of correct selection:

$$Max\left[P\left(\hat{p}_1 < \hat{p}_2 \mid p_1 < p_2\right)\right] = Min\left[P\left(\hat{p}_1 - \hat{p}_2\right) > 0 \mid p_1 < p_2\right]$$

$$\text{s/t} \quad b_{11}n_{11} + b_{12}n_{12} + b_{21}n_{21} + b_{22}n_{22} = T, \qquad n_{kj} \geq 1$$

$H_{kj} \sim Binomial\left(p_{kj}, n_{kj}\right)$. Assuming n_{kj} is sufficiently large, H_{kj} is approximately normal, as is

$$\hat{p}_{kj} = \frac{H_{kj}}{n_{kj}} \sim N\left[p_{kj}, \frac{p_{kj}\left(1-p_{kj}\right)}{n_{kj}}\right]$$

$$p_1 = p_{11}\, p_{12} \qquad p_2 = p_{21}\, p_{22}$$
$$\hat{p}_1 = \hat{p}_{11}\, \hat{p}_{12} \qquad \hat{p}_2 = \hat{p}_{21}\, \hat{p}_{22}$$

It can be shown that $\hat{p}_k$ is an unbiased estimator of p_k even though $\hat{p}_{k1}$ and $\hat{p}_{k2}$ are not necessarily independent (see L'Ecuyer, Demers, and Tuffin 2006).

Assuming that the probability of reaching level j, starting from any entrance state to level $j-1$, does not depend on the entrance state, the variance of $\hat{p}_k$ can be derived as

$$\sigma_k^2\left(n_{k1}, n_{k2}\right) = \frac{p_{k1}^2\, p_{k2}\left(1-p_{k2}\right)}{n_{k2}} + \frac{p_{k2}^2\, p_{k1}\left(1-p_{k1}\right)}{n_{k1}} + \frac{p_{k1}\, p_{k2}\left(1-p_{k1}\right)\left(1-p_{k2}\right)}{n_{k1}\, n_{k2}}, \quad k = 1, 2$$

Given that the $\hat{p}_{kj}$ are (approximately) normal, we assume that for large n_{kj} their products are also approximately normal:

$$\hat{p}_k = \hat{p}_{k1}\hat{p}_{k2} \sim N\left[p_k, \sigma_k^2\left(n_{k1}, n_{k2}\right)\right]$$

Define

$$g\left(n_{11}, n_{12}, n_{21}, n_{22}\right) \equiv \frac{p_2 - p_1}{\sqrt{\sigma_1^2\left(n_{11}, n_{12}\right) + \sigma_2^2\left(n_{21}, n_{22}\right)}}$$

Then

$$P\left(\hat{p}_1 - \hat{p}_2 > 0\right) = 1 - \Phi\left[g\left(n_{11}, n_{12}, n_{21}, n_{22}\right)\right]$$

Minimizing that probability is equivalent to maximizing $\Phi\left[g\right]$, which is equivalent to

$$\begin{aligned} &Min \sqrt{\sigma_1^2\left(n_{11}, n_{12}\right) + \sigma_2^2\left(n_{21}, n_{22}\right)} \\ &\text{s/t} \quad b_{11}n_{11} + b_{12}n_{12} + b_{21}n_{21} + b_{22}n_{22} = T, \quad n_{kj} \geq 1 \end{aligned}$$

This is the global optimization problem. To derive the OCBA step of OCBA+OSTRE we must transform it into a minimization over N_1 and N_2, where $N_1 = n_{11} + n_{12}$, $N_2 = n_{21} + n_{22}$.

To obtain a simplified approximation of the optimal relationship between the n_{kj} within each design, Shortle et al. (2011) consider an asymptotic solution, as $n_{kj} \to \infty$. In addition, the p_{kj} being typically small, they set $1 - p_{kj} \approx 1$. These assumptions will be made in all subsequent derivations. Under them the OSTRE relationship they derive simplifies to

$$n_{k1}\sqrt{b_{k1}p_{k1}} = n_{k2}\sqrt{b_{k2}p_{k2}} \qquad k = 1, 2$$

Define $\qquad \alpha_k \equiv \sqrt{\dfrac{b_{k1}p_{k1}}{b_{k2}p_{k2}}}, \qquad$ implying $\qquad n_{k1} = \dfrac{1}{\alpha_k + 1}N_k, \quad n_{k2} = \dfrac{\alpha_k}{\alpha_k + 1}N_k.$

For large n_{kj} the last term in the expansion of $\sigma_k^2\left(n_{k1}, n_{k2}\right)$ is negligible, compared to the first two. Ignoring it and substituting for n_{k1} and n_{k2} reduces the variance to

$$\sigma_k^2\left(N_k\right) = \frac{\left(\alpha_k + 1\right)p_{k1}p_{k2}}{\alpha_k}\left(\frac{p_{k1} + \alpha_k\, p_{k2}}{N_k}\right)$$

The requirement that $N_k \geq 2$, and that they be integers, is automatically satisfied by the implementation of the algorithm, where only positive integers are allocated, and there is, in the initialization phase, an allocation to each level of each design. For theoretical purposes we may treat the N_k, and the n_{kj} as continuous variables, the rounding, in practice, being inconsequential. Now the problem is

$$\begin{aligned} &Min\ \sqrt{\sigma_1^2(N_1)+\sigma_2^2(N_2)} \\ &\text{s/t} \quad T = b_1 N_1 + b_2 N_2, \quad N_k \geq 2. \end{aligned}$$

where $b_k \equiv$ the average time consumed by an allocation of one run to design *k*.

From $b_k N_k = b_{k1} n_{k1} + b_{k2} n_{k2}$ and the relationship between the n_{kj} and N_k stated above we have

$$b_k = \frac{b_{k1} + \alpha_k b_{k2}}{\alpha_k + 1}.$$

The optimality conditions of the minimization problem imply

$$\frac{1}{b_1}\frac{\partial \sigma_1^2(N_1)}{\partial N_1} = \frac{1}{b_2}\frac{\partial \sigma_2^2(N_2)}{\partial N_2}$$

$$\frac{(\alpha_1+1)p_{11}p_{12}}{\alpha_1 b_1}\left(\frac{p_{11}+\alpha_1 p_{12}}{N_1^2}\right) = \frac{(\alpha_2+1)p_{21}p_{22}}{\alpha_2 b_2}\left(\frac{p_{21}+\alpha_2 p_{22}}{N_2^2}\right)$$

$$\frac{N_1^2 b_1 \alpha_1}{p_1(\alpha_1+1)(p_{11}+\alpha_1 p_{12})} = \frac{N_2^2 b_2 \alpha_2}{p_2(\alpha_2+1)(p_{21}+\alpha_2 p_{22})} \tag{1}$$

Equation (1) is the OCBA relationship upon which runs are allocated to designs. If, upon updating, the left-hand side is less than the right, design 1 gets the next (small) allocation of runs. Otherwise it goes to design 2. It is then assigned to level 1 or 2 of the receiving design, depending on the intra-design OSTRE relationship. An allocation to design *k* flows to n_{k1} if $n_{k1}\sqrt{b_{k1}p_{k1}} \leq n_{k2}\sqrt{b_{k2}p_{k2}}$, n_{k2} otherwise.

Comparing this OCBA result, equation (1), with the original reveals how OCBA changes when OSTRE instead of MC is employed within the designs. The original result (see Chen & Lee 2011) is

$$N_1^2 b_1 \nu_2^2 = N_2^2 b_2 \nu_1^2$$

where $\nu_k^2 =$ the variance of <u>one</u> MC run in design *k*. A run hits the rare event with probability p_k or misses with $1-p_k$, so its variance is $p_k(1-p_k)$:

$$N_1^2 b_1 p_2(1-p_2) = N_2^2 b_2 p_1(1-p_1) \tag{2}$$

The OCBA in (2) differs from the OCBA in (1) because there is no OSTRE counterpart to <u>one</u> MC run generating a rare event with probability p_k.

3 EQUIVALENCE OF OCBA+OSTRE & GLOBAL OPTIMIZATION

In this section we show that the two-stage method is mathematically equivalent to global optimization. For two designs and an arbitrary number of splitting levels the global optimization problem has been solved by Shortle et al. (2011). (They call it OSTRE2, meaning OSTRE applied to 2 designs.) For two levels the optimality equations, after making the simplifying assumptions noted above, are

$$\frac{n_{11}\sqrt{b_{11}p_{11}}}{p_{11}p_{12}} = \frac{n_{12}\sqrt{b_{12}p_{12}}}{p_{11}p_{12}}$$

$$\frac{n_{21}\sqrt{b_{21}p_{21}}}{p_{21}p_{22}} = \frac{n_{22}\sqrt{b_{22}p_{22}}}{p_{21}p_{22}}$$

$$\frac{n_{11}\sqrt{b_{11}p_{11}}}{p_{11}p_{12}} = \frac{n_{21}\sqrt{b_{21}\,p_{21}}}{p_{21}p_{22}}$$

The first two are the intra-design OSTRE relationships. They obviously hold in the two-stage method because we employed them in developing the OCBA step. Thus the equivalence of OCBA+OSTRE and global optimization depends on the third equation. The methods are equivalent if substituting

$$N_1 = n_{11} + n_{12} \quad \text{and} \quad N_2 = n_{21} + n_{22}$$

in the OCBA+OSTRE optimality equation, and utilizing the intra-design OSTRE relationships, reduces it to the third global optimization equation. It does so reduce, proving the theorem below.

Theorem: For 2 designs, 2 levels, the OCBA+OSTRE two-stage method is mathematically equivalent to global optimization (OSTRE2), under the same set of simplifying approximations.

Proof: See Appendix A.

Proving the general case is more challenging, and is ongoing research.

4 TESTS

We compare, for 2 designs, 2 levels, the performances of OCBA+OSTRE, global optimization (OSTRE2), and MC, and then apply OCBA+OSTRE to 3 designs, 2 levels. The tests are on a queuing model with two servers in tandem. The process begins with a customer in server 1, both queues and server 2 empty. It terminates when both queues are empty, or there are 100 customers in either queue. Arrival and service rates are chosen to make the second terminating state – 100 customers in either queue – a (relatively) rare event. (The probabilities of the rare event in each design are approximately 3.6e-6, 3.9e-6, and 4.3e-6.) Designs are differentiated by arrival and service rates, with design 1 the best.

In these tests an "experiment" is the simulation of all designs, subject to a given budget constraint (in terms of time, for the set of designs), generating probability estimates for each design. From these estimates the minimum is selected. The percentages of correct selections (design 1), over a large number of experiments at several budget levels, are plotted for each method. Those percentages should, in general, increase with rising budgets. We compare the methods over ranges of fairly small budgets, where choosing the best design is, for any method, a considerable challenge.

4.1 2 Designs, 2 Levels

The budget for the first test (2 designs, 2 levels) begins at 30 seconds and increases, by 10 second increments, to 5 minutes, with 1000 experiments at each budget level. At these small budgets MC is hopeless, worth little more than coin tossing. But both global optimization (OSTRE2) and OCBA+OSTRE become quite reliable. After 5 minutes they both achieve 95% correct selection. And they are essentially indistinguishable, as expected from their mathematical equivalence. Figure 1 shows the performance.

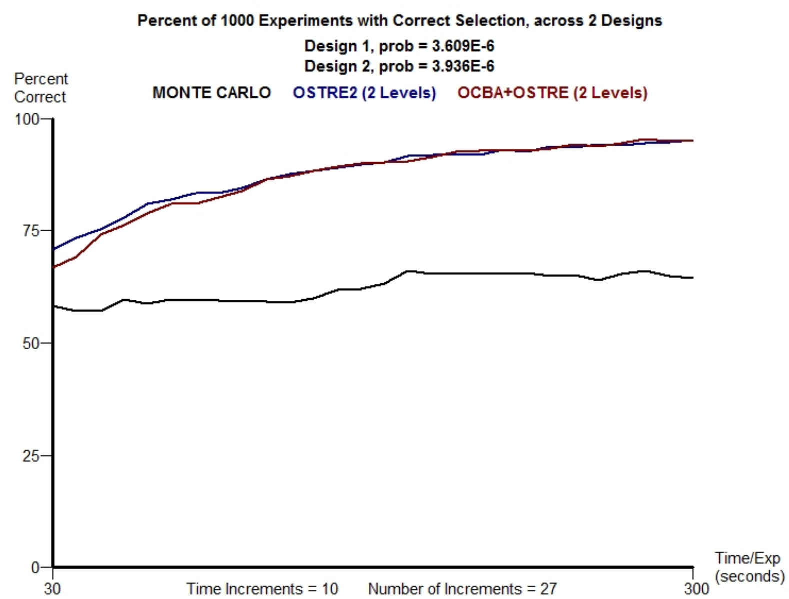

Figure 1: Performance of OCBA+OSTRE fore 2 Designs, 2 Levels

4.2 3 Designs, 2 Levels

For 3 designs we do not (yet) have a good implementation of global optimization, but we do have OCBA+OSTRE. (The derivation is similar to that for 2 designs.) Nor do we have proof that OCBA+OSTRE is equivalent to global optimization for 3 or more designs. But OCBA+OSTRE performs well in practice.

The next test pits OCBA+OSTRE against OSTRE by itself and OCBA by itself, across 3 designs. The starting time is increased to a minute, to accommodate OCBA (which requires relatively long initialization), and the budgets are extended to 7 minutes. The observation that OCBA alone is only slightly better than MC suggests that most of OCBA+OSTRE's punch should come from OSTRE. That is clearly seen in this test, as OCBA+OSTRE consistently, but only modestly, outperforms OSTRE alone. Figure 2 shows the performance in this case.

5. CONCLUSION

The problem is to select the best among a set of designs, where best means the design with the lowest probability of a rare event. The probability of making the correct selection is marginally enhanced by employing OCBA to allocate runs (or time) among the designs. The probability of making the correct selection is significantly enhanced by employing OSTRE, an optimized form of splitting, to estimate the rare event probability of each design. Combining OCBA with OSTRE – a two-stage optimization – should enhance the probability of making the correct selection more than either alone can do. The best approach would be a single optimization over all decision variables – the runs (or time) allocated to each splitting level across all designs. That optimization becomes very unwieldy as the number of designs (and levels) increases. The OCBA+OSTRE approach promises to be more tractable, and deliver results almost as good as single optimization. In fact, in the simplest case of 2 designs, 2 levels, OCBA+OSTRE is mathematically equivalent to single optimization. Whether or not equivalence holds in the general case (an arbitrary number of designs and levels) is uncertain, a question for further research.

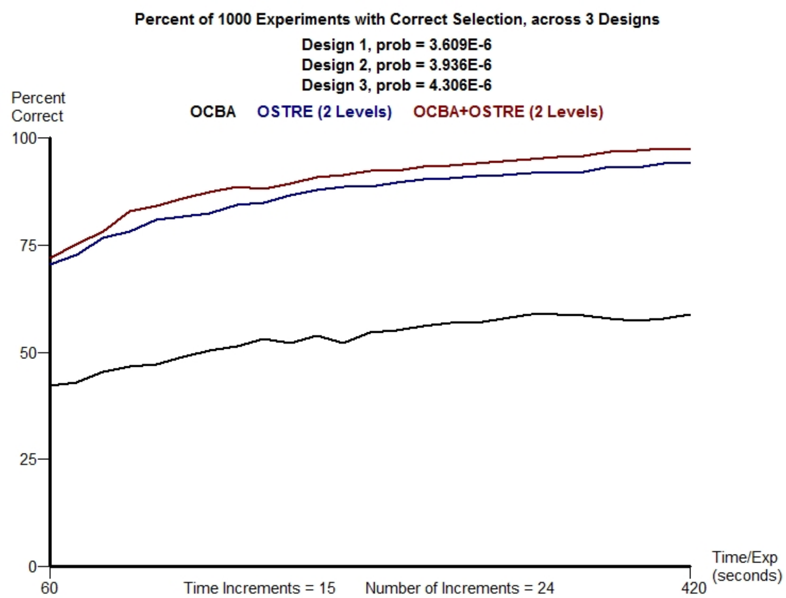

Figure 2: Performance of OCBA+OSTRE for 3 Designs

A THEOREM

Theorem:

For 2 designs, 2 levels, the OCBA+OSTRE two-stage method is mathematically equivalent to global optimization (OSTRE2), under the same simplifying approximations.

Proof:

Under the same simplifying approximations OCBA+OSTRE and OSTRE2 are equivalent if

$$\frac{N_1^2 b_1 \alpha_1}{p_1(\alpha_1+1)(p_{11}+\alpha_1 p_{12})} = \frac{N_2^2 b_2 \alpha_2}{p_2(\alpha_2+1)(p_{21}+\alpha_2 p_{22})} \quad \Rightarrow \quad \frac{n_{11}^2 b_{11} p_{11}}{(p_{11}p_{12})^2} = \frac{n_{21}^2 b_{21} p_{21}}{(p_{21}p_{22})^2}$$

The left and right-hand sides of both equations are symmetrical, differing only in the design index, so it suffices to show the LHS of OCBA+OSTRE reduces to the LHS of OSTRE2, namely

$$\frac{N_1^2 b_1 \alpha_1}{p_1(\alpha_1+1)(p_{11}+\alpha_1 p_{12})} = \frac{n_{11}^2 b_{11} p_{11}}{(p_{11}p_{12})^2}$$

Make the following substitutions:

$$p_1 = p_{11}p_{12}\,, \quad \alpha_1 \equiv \sqrt{\frac{b_{11}p_{11}}{b_{12}p_{12}}}\,, \quad b_1 = \frac{b_{11} + \alpha_1 b_{12}}{\alpha_1 + 1}\,, \quad N_1 = n_{11} + n_{12}\,, \quad n_{12} = \alpha_1 n_{11}\,.$$

Then,

$$\frac{N_1^2\, b_1\, \alpha_1}{p_1\left(\alpha_1 + 1\right)\left(p_{11} + \alpha_1\, p_{12}\right)} = \frac{\left(n_{11} + n_{12}\right)^2 \left(\frac{b_{11} + \alpha_1 b_{12}}{\alpha_1 + 1}\right)\alpha_1}{p_1\left(\alpha_1 + 1\right)\left(p_{11} + \alpha_1\, p_{12}\right)} = \frac{\left(\alpha_1 + 1\right)^2 n_{11}^2 \left(b_{11} + \alpha_1 b_{12}\right)\alpha_1}{p_1\left(\alpha_1 + 1\right)^2 \left(p_{11} + \alpha_1\, p_{12}\right)} =$$

$$\frac{n_{11}^2\left(b_{11} + \alpha_1 b_{12}\right)\alpha_1}{p_1\left(p_{11} + \alpha_1\, p_{12}\right)} = \frac{n_{11}^2\left(b_{11} + \sqrt{\frac{b_{11}p_{11}}{b_{12}p_{12}}}\, b_{12}\right)\sqrt{\frac{b_{11}p_{11}}{b_{12}p_{12}}}}{p_1\left(p_{11} + \sqrt{\frac{b_{11}p_{11}}{b_{12}p_{12}}}\, p_{12}\right)} = \frac{n_{11}^2\sqrt{b_{11}}\left(\sqrt{b_{11}} + \sqrt{\frac{p_{11}}{p_{12}}}\sqrt{b_{12}}\right)\sqrt{\frac{b_{11}p_{11}}{b_{12}p_{12}}}}{p_1\sqrt{p_{11}}\left(\sqrt{p_{11}} + \sqrt{\frac{b_{11}}{b_{12}}}\sqrt{p_{12}}\right)}\,.$$

Multiplying the numerator and denomination by $\frac{1}{\sqrt{b_{12}p_{12}}}$ and distributing the product appropriately yields:

$$\frac{n_{11}^2\sqrt{\frac{b_{11}}{p_{12}}}\,\sqrt{b_{11}p_{11}}\left(\sqrt{\frac{b_{11}}{b_{12}}} + \sqrt{\frac{p_{11}}{p_{12}}}\right)}{p_1\sqrt{p_{11}p_{12}}\left(\sqrt{\frac{b_{11}}{b_{12}}} + \sqrt{\frac{p_{11}}{p_{12}}}\right)} = \frac{n_{11}^2\, b_{11}\, \sqrt{p_{11}}}{p_{11}\, p_{12}\sqrt{p_{11}}\, p_{12}} = \frac{n_{11}^2\, b_{11}\, p_{11}}{\left(p_{11}\, p_{12}\right)^2}\,.$$

ACKNOWLEDGEMENT

This work has been supported in part by the Department of Energy under Award DE-SC0002223, and NIH under Grant 1R21DK088368-01.

REFERENCES

Asmussen, S., and P.W. Glynn. 2007. *Stochastic Simulation: Algorithms and Analysis*. Springer.

Chen, C. H., J. Lin, E. Yücesan, and S. E. Chick. 2000. “Simulation Budget Allocation for Further Enhancing the Efficiency of Ordinal Optimization.” *Journal of Discrete Event Dynamic Systems: Theory and Applications* 10:251-270.

Chen, C. H., and L. H. Lee. 2011. *Stochastic Simulation Optimization: An Optimal Computing Budget Allocatio*n. World Scientific Publishing Co.

Fischer, M., D. Masi, J. F. Shortle, and C. H. Chen. 2010. “Simulating Non-Stationary Congestion Systems Using Splitting with Applications to Cyber Security.” In *Proceedings of the 2010 Winter Simulation Conference,* edited by B. Johansson, S. Jain, J. Montoya-Torres, J. Hugan, and E. Yücesan, 2865-2875. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Fu, M. C., J. Q. Hu, C. H. Chen, and X. Xiong. 2007. “Simulation Allocation for Determining the Best Design in the Presence of Correlated Sampling.” *Informs Journal on Computing* 19(1):101-111.

He, D., S. E. Chick, C. H. Chen. 2007. "The Opportunity Cost and OCBA Selection Procedures in Ordinal Optimization," *IEEE Transactions on Systems, Man, and Cybernetics—Part C* 37(5):951-961.

Lagnoux-Renaudie, A. 2008. "Rare event simulation: Effective splitting model under cost constraint." In *Stochastic Processes and their Applications,* 1820-1851. Elsevier.

L'Ecuyer, P., V. Demers, and B. Tuffin. 2006. "Splitting for Rare-Event Simulation." In *Proceedings of the 2006 Winter Simulation Conference*, edited by L. R. Perrone, F. P. Wieland, J. Liu, B. G. Lawson, D. M. Nicol, and R. M. Fujimoto, 137-148. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

L'Ecuyer, P., F. Le Gland, P. Lezaud, and B. Tuffin. 2009. "Splitting Techniques." In *Rare Event Simulation using Monte Carlo Methods,* edited by G. Rubino and B. Tuffin. Wiley.

Shortle, J. and C. H. Chen. 2008. "A Preliminary Study of Optimal Splitting for Rare-Event Simulation." In *Proceedings of the 2008 Winter Simulation Conference*, edited by S. J. Mason, R. R. Hill, L. Mönch, O. Rose, T. Jefferson, J. W. Fowler, 266-272. Piscataway, New Jersey: Institute of Electrical and Electronics Engineers, Inc.

Shortle, J., C.H. Chen, A. Brodsky, and D. Brod. 2011. "Optimal Level Splitting for Rare-Event Simulation." to appear in *IIE Transactions*.

Villen-Altamirano, M., and J. Villen-Altamirano. 2006. "On the Efficiency of RESTART for Multidimensional State Systems." *ACM Transactions on Modeling and Computer Simulation* 16:251-279.

AUTHOR BIOGRAPHIES

BEN CRAIN is a graduate student in the Systems Engineering and Operations Research Department at George Mason University. His email address is bcrain@gmu.edu

CHUN-HUNG CHEN is a Professor of Electrical Engineering at National Taiwan University. Dr. Chen has led research projects in stochastic simulation and optimization, systems design under uncertainty, and air traffic management, which are sponsored by NSF, FAA, and NASA. He served as Co-Editor of the *Proceedings of the 2002 Winter Simulation Conference* and Program Co-Chair for *2007 INFORMS Simulation Society Workshop*. He is serving on the editorial boards of *IEEE Transactions on Automatic Control, IIE Transactions, Journal of Simulation Modeling Practice and Theory*, and *International Journal of Simulation and Process Modeling*. He received his Ph.D. degree from Harvard University in 1994. His email address is cchen9@gmu.edu.

JOHN F. SHORTLE is an Associate Professor of Systems Engineering and Operations Research at George Mason University. His research interests include simulation and queuing applications in air transportation, telecommunications, and energy. Previously, he worked at US WEST Advanced Technologies. He received his doctorate degree in industrial engineering and operations research from UC Berkeley in 1996. His email address is jshortle@gmu.edu.