Applied Incremental Dynamic Analysis [1]

Dimitrios Vamvatsikos,[a)] **M.EERI, and C. Allin Cornell**[b)] **M.EERI**

We are presenting a practical and detailed example of how to perform incremental dynamic analysis (IDA), interpret the results and apply them to performance-based earthquake engineering. IDA is an emerging analysis method that offers thorough seismic demand and capacity prediction capability by using a series of nonlinear dynamic analyses under a multiply scaled suite of ground motion records. Realization of its opportunities requires several steps and the use of innovative techniques at each one of them. Using a nine-story steel moment-resisting frame with fracturing connections as a testbed, the reader is guided through each step of IDA: (1) choosing suitable ground motion intensity measures and representative damage measures, (2) using appropriate algorithms to select the record scaling, (3) employing proper interpolation and (4) summarization techniques for multiple records to estimate the probability distribution of the structural demand given the seismic intensity, and (5) defining limit-states, such as the dynamic global system instability, to calculate the corresponding capacities. Finally, (6) the results can be used to gain intuition for the structural behavior, highlighting the connection between the static pushover (SPO) and the dynamic response, or (7) they can be integrated with conventional probabilistic seismic hazard analysis (PSHA) to estimate mean annual frequencies of limit-state exceedance. Building upon this detailed example based on the nine-story structure, a complete commentary is provided, discussing the choices that are available to the user, and showing their implications for each step of the IDA.

INTRODUCTION

An important issue in performance-based earthquake engineering (PBEE) is the estimation of structural performance under seismic loads, in particular the estimation of the mean annual frequency (MAF) of exceeding a specified level of structural demand (e.g., the maximum, over all stories, peak interstory drift ratio $\theta_{\max}$) or a certain limit-state capacity (e.g., global dynamic instability). A promising method that has recently risen to meet these needs is incremental dynamic analysis (IDA), which involves performing nonlinear dynamic analyses of the structural model under a suite of ground motion records, each scaled to several intensity levels designed to force the structure all the way from elasticity to final global dynamic instability (Vamvatsikos and Cornell 2002a).

Applying IDA to determine the performance of a structure requires several steps. First, a proper nonlinear structural model needs to be formed, and a suite of records must be compiled. Then, for each record, the scaling levels must be selected, the dynamic analyses run and the results postprocessed. Thus, we can generate IDA curves of the structural response, as measured by a damage measure (*DM*, e.g., peak roof drift ratio θ_{roof} or $\theta_{\max}$), versus the ground motion

[1] Based on a short paper presented at the 12th European Conference on Earthquake Engineering, London, 2002

[a)] Theagenous 11, Athens, 11634, Greece; e-mail: divamva@stanfordalumni.org

[b)] Professor, Department of Civil and Environmental Engineering, Stanford University, CA 94305-4020

Table 1. The set of twenty ground motion records used

No	Event	Station	ϕ° [1]	Soil[2]	M[3]	R[4] (km)	PGA (g)
1	Loma Prieta, 1989	Agnews State Hospital	090	C,D	6.9	28.2	0.159
2	Imperial Valley, 1979	Plaster City	135	C,D	6.5	31.7	0.057
3	Loma Prieta, 1989	Hollister Diff. Array	255	–,D	6.9	25.8	0.279
4	Loma Prieta, 1989	Anderson Dam Downstream	270	B,D	6.9	21.4	0.244
5	Loma Prieta, 1989	Coyote Lake Dam Downstream	285	B,D	6.9	22.3	0.179
6	Imperial Valley, 1979	Cucapah	085	C,D	6.5	23.6	0.309
7	Loma Prieta, 1989	Sunnyvale Colton Ave	270	C,D	6.9	28.8	0.207
8	Imperial Valley, 1979	El Centro Array #13	140	C,D	6.5	21.9	0.117
9	Imperial Valley, 1979	Westmoreland Fire Station	090	C,D	6.5	15.1	0.074
10	Loma Prieta, 1989	Hollister South & Pine	000	–,D	6.9	28.8	0.371
11	Loma Prieta, 1989	Sunnyvale Colton Ave	360	C,D	6.9	28.8	0.209
12	Superstition Hills, 1987	Wildlife Liquefaction Array	090	C,D	6.7	24.4	0.180
13	Imperial Valley, 1979	Chihuahua	282	C,D	6.5	28.7	0.254
14	Imperial Valley, 1979	El Centro Array #13	230	C,D	6.5	21.9	0.139
15	Imperial Valley, 1979	Westmoreland Fire Station	180	C,D	6.5	15.1	0.110
16	Loma Prieta, 1989	WAHO	000	-,D	6.9	16.9	0.370
17	Superstition Hills, 1987	Wildlife Liquefaction Array	360	C,D	6.7	24.4	0.200
18	Imperial Valley, 1979	Plaster City	045	C,D	6.5	31.7	0.042
19	Loma Prieta, 1989	Hollister Diff. Array	165	–,D	6.9	25.8	0.269
20	Loma Prieta, 1989	WAHO	090	–,D	6.9	16.9	0.638

[1] Component
[2] USGS, Geomatrix soil class
[3] Moment magnitude
[4] Closest distance to fault rupture

intensity level, measured by an intensity measure (*IM*, e.g., peak ground acceleration, PGA, or the 5%-damped first-mode spectral acceleration $S_a(T_1, 5\%)$). In turn these are interpolated for each record and summarized over all records to estimate the distribution of demand *DM* given intensity *IM*. Subsequently, limit-states (e.g., immediate occupancy or collapse prevention in SAC 2000a,b) can be defined on each IDA curve and summarized to produce the probability of exceeding a specified limit-state given the *IM* level. The final results are in a suitable format to be conveniently integrated with a conventional PSHA hazard curve in order to calculate MAFs of exceeding a certain limit-state capacity, or a certain demand.

Building upon this foundation, we will discuss several topics of practical interest, showing in detail the reasons behind the choices made in our example and the advantages or disadvantages of each. In particular, subjects like the number of runs, the algorithms used for scaling-level selection, and possible approximations used for the probabilistic calculations are going to be presented showing their impact upon the accuracy of PBEE calculations.

MODEL AND GROUND MOTION RECORDS

To illustrate our methodology, we will use a centerline model of a nine-story steel moment-resisting frame designed for Los Angeles according to the 1997 NEHRP provisions (Lee and Foutch 2002). The model has a first-mode period of $T_1 = 2.37$ sec and it incorporates ductile members, shear panels and realistically fracturing reduced beam section connections, while it includes the influence of interior gravity columns and a first-order treatment of global geometric nonlinearities (P-Δ effects).

In addition we need a suite of ground motion records. Previous studies (Shome and Cornell

1999) have shown that for mid-rise buildings, ten to twenty records are usually enough to provide sufficient accuracy in the estimation of seismic demands, assuming a relatively efficient *IM*, like $S_a(T_1, 5\%)$, is used. Consequently, we have selected a set of twenty ground motion records, listed in Table 1, that belong to a bin of relatively large magnitudes of 6.5 – 6.9 and moderate distances, all recorded on firm soil and bearing no marks of directivity; effectively they represent a scenario earthquake.

PERFORMING THE ANALYSIS

Once the model has been formed and the ground motion records have been selected, we need a fast and automated way to perform the actual dynamic analyses required for IDA. This entails appropriately scaling each record to cover the entire range of structural response, from elasticity, to yielding, and finally global dynamic instability. Our task is made significantly easier by using an advanced algorithm, like *hunt & fill* (Vamvatsikos and Cornell 2002a). This ensures that the record scaling levels are appropriately selected to minimize the number of required runs: Analyses are performed at rapidly increasing levels of *IM* until numerical non-convergence is encountered (signaling global dynamic instability), while additional analyses are run at intermediate *IM*-levels to sufficiently bracket the global collapse and increase the accuracy at lower *IM*s. The user only needs to specify the desired accuracy for demand and capacity, select the maximum tolerable number of dynamic analyses, and then wait for a few hours to get the results. Since the algorithm has been implemented in software (Vamvatsikos and Cornell 2002b) able to wrap around most existing analysis programs (e.g., DRAIN-2DX, Prakhash et al. 1992) it renders IDA almost effortless, needing no human supervision.

As an example, we will show in detail the computations resulting to the *IM*-levels selected by hunt & fill when tracing record #14 from Table 1. To express the scaling level we need an initial, temporary choice of *IM*, and we have chosen $S_a(T_1, 5\%)$, a decision that need not restrict us in any way: scaling can be re-expressed in any other scalable *IM* (Vamvatsikos and Cornell 2002a) that we wish after the runs are performed. Hence, in $S_a(T_1, 5\%)$ terms, the algorithm was configured to use an initial step of 0.1 g, a step increment of 0.05 g and a designated first elastic run at 0.005 g, while a maximum of 12 runs was allowed for each record. Additionally, we specified a resolution of 10% on the global collapse capacity, i.e., we expect the model to develop numerical nonconvergence and show practically infinite $\theta_{\max}$ at some high intensity level, and we wish this level to be known within 10% of its *IM*-value. Finally, we allowed the demand resolution, i.e., the maximum difference between successive *IM*-values, to run to its best attainable value by expending all the 12 runs. Alternatively, we could have designated some minimum satisfactory *IM*-gap below which we do not wish to proceed, thus saving some runs.

Using the above settings we get the sequence of runs shown in Table 2. The first run is meant to be in the elastic region. In the subsequent five runs, 2–6, we are hunting upwards till the first numerical nonconvergence appears in the form of "infinite" $\theta_{\max}$. Then, the dynamic analysis algorithm does not converge thus either failing to complete the dynamic run (as happened for this record) or producing extreme values of $\theta_{\max}$, say 200%. The next two runs, 7–8, are used to better bracket the first appearance of nonconvergence, closing within 10% of its *IM*-value ($(0.872 - 0.805)/0.805 = 8.32\% < 10\%$) so that the gap between highest converging and lowest nonconverging run is less that 10% of the former. Notice that instead of placing each new run in the middle of the gap, the algorithm places it closer to the converging run, only one third of the way to the nonconverging one. This ensures that the search will be somewhat

Table 2. Sequence of runs generated by the hunt & fill tracing algorithm for record #14

No.	calculations	$S_a(T_1,5\%)$ (g)	$\theta_{\max}$
1		0.005	0.05%
2	$0.005+0.10$	0.105	0.79%
3	$0.105+0.10+1\times0.05$	0.255	2.02%
4	$0.255+0.10+2\times0.05$	0.455	3.01%
5	$0.455+0.10+3\times0.05$	0.705	5.85%
6	$0.705+0.10+4\times0.05$	1.005	$+\infty$
7	$0.705+(1.005-0.705)/3$	0.805	18.83%
8	$0.805+(1.005-0.805)/3$	0.872	$+\infty$
9	$(0.805+0.705)/2$	0.755	9.18%
10	$(0.705+0.455)/2$	0.580	3.27%
11	$(0.455+0.255)/2$	0.355	2.96%
12	$(0.255+0.105)/2$	0.180	1.34%

biased towards converging runs, which are more informative than nonconverging ones (which are essentially discarded). The rest of the runs, up to the maximum of 12, are used to fill in the IDA at lower levels, being sequentially placed in the middle of the largest *IM*-gaps. Thus, the large gaps left by the initial increasing steps to the flatline (runs 2–6), are filled in; this step increases the demand resolution and, given enough runs, it ensures that the algorithm has not missed an earlier collapse. Although it is a rare phenomenon in multi-degree-of-freedom structural models, certain records may cause them to collapse for a range of *IM*-values, but not for some higher *IM*, an event we call *structural resurrection* (Vamvatsikos and Cornell 2002a). By reducing the *IM*-gaps with runs 9–12, we are making sure that we have not missed such an earlier (in *IM* terms) global collapse and the flatline we have found is the first one to occur.

Notice that the maximum *IM*-gap, i.e., the demand resolution, is about 0.13 g (but less than half on average), while the *IM*-difference between the highest converging and lowest nonconverging run (the capacity resolution) is much less than 10% of the highest converging *IM*, about 0.06 g. Naturally, if we knew a priori the approximate *IM*-height of the flatline, we could use a stepping algorithm with 12 runs and constant step of 0.1 g to achieve similar results with a homogeneous distribution of the accuracy, but this scheme would fail with the next records, producing either too few or too many runs, due to the large record-to-record variability.

Assuming that the computational cost for each run is the same, then, the more the analyses per record, the longer for IDA to complete but the better the accuracy. Still, with the use of such an advanced algorithm no runs are wasted, thus 12 runs per record will suffice to strike a good compromise between speed and accuracy. Nevertheless, it may be pointed out that performing 240 dynamic runs for a model with thousands of degrees-of-freedom is a daunting task. Yet, even for such a complicated model, it took less than 12 hours on two 1999-era Pentium-class processors running independently. The process is completely automated and so easily performed overnight that actually setting up the structural model can now be expected to take substantially more (human) time than doing the analysis, while computer time is becoming an ever-cheaper commodity.

POSTPROCESSING

Equally important to the analysis is the postprocessing of the resulting data and perhaps the most important issue is selecting a suitable *IM* and *DM*. There are several issues of efficiency and sufficiency associated with the *IM* selection (Luco 2002). Since there are no directivity-influenced records in our suite and the building is of medium height (hence first-mode-dominated), the 5%-damped first-mode spectral acceleration $S_a(T_1, 5\%)$ will be our choice; it has been proven to be both efficient, by minimizing the scatter in the results, requiring only a few ground motion records to provide good demand and capacity estimates, and sufficient, as it provides a complete characterization of the response without the need for magnitude or source-to-site distance information (Shome and Cornell 1999). Similarly, selecting a *DM* can be application-specific; for example, the peak floor accelerations are correlated with contents' damage and many types of non-structural elements' damage, while the maximum peak interstory drift ratio $\theta_{\max}$ (the maximum over time and over all stories of the interstory drift ratios recorded during the timehistory analysis) is known to relate well (SAC 2000a) to global dynamic instability and several structural performance limit-states upon which we intend to focus. Therefore, $\theta_{\max}$ will be our *DM*-choice. Still, it must be emphasized that these *IM* and *DM* choices are by no means limiting. Assuming that additional *DM*s have been recorded from the analyses, they can be substituted instead of $\theta_{\max}$, and by employing the postprocessing techniques presented, the IDA data can be expressed in a different scalable *IM*, without any need to rerun the dynamic analyses.

Having selected our *IM* and *DM*, we are still faced with an abundance of IDA-generated data that need to be sorted out and presented in meaningful ways. It is a time-consuming and challenging task that we are going to step our way through, but it can be rendered totally effortless with the proper software. Actually, most of what follows is a direct description of the inner workings of an automated postprocessing program (Vamvatsikos and Cornell 2002b), whose graphical output appears in the accompanying figures.

GENERATING THE IDA CURVES BY INTERPOLATION

Once the desired *IM* and *DM* values (in our case $S_a(T_1, 5\%)$ and $\theta_{\max}$) are extracted from each of the dynamic analyses, we are left with a set of discrete points for each record that reside in the *IM*-*DM* plane and lie on its IDA curve, as in Figure 1. By interpolating them, the entire IDA curve can be approximated without performing additional analyses. To do so, we may use a basic piecewise linear approximation, or the superior spline interpolation. Based on the concept of natural, coordinate-transformed, parametric splines with a centripetal scheme for knot-selection (Lee 1989, Farin 1990), a realistic interpolation can be generated that accurately represents the real IDA curve, as shown in Figure 1 for our example of record #14 in Table 2. Having the complete curve available, it is now possible to calculate *DM* values at arbitrary levels of *IM*, allowing the extraction of more (IM, DM) points with a minimum of computation.

The spline comes in n cubic polynomial pieces and is parameterized on a single non-negative parameter, $t \in [0, t_1] \bigcup \ldots \bigcup [t_{n-1}, t_n]$, where n is the number of convergent runs/points including the default (0,0) point, i.e., $n = 10 + 1 = 11$ for record #14, Table 2. For each value of the parameter t, and depending on the interval $[t_{i-1}, t_i]$ where it lies, we get two polynomials, one for the *IM* (the x-variable) and one for the *DM* (the y-variable):

$$\begin{cases} x_i(t) = a_{xi}\, t^3 + b_{xi}\, t^2 + c_{xi}\, t + d_{xi} \\ y_i(t) = a_{yi}\, t^3 + b_{yi}\, t^2 + c_{yi}\, t + d_{yi} \end{cases} \qquad t \in [t_{i-1}, t_i], \quad i = 1, \ldots, n \tag{1}$$

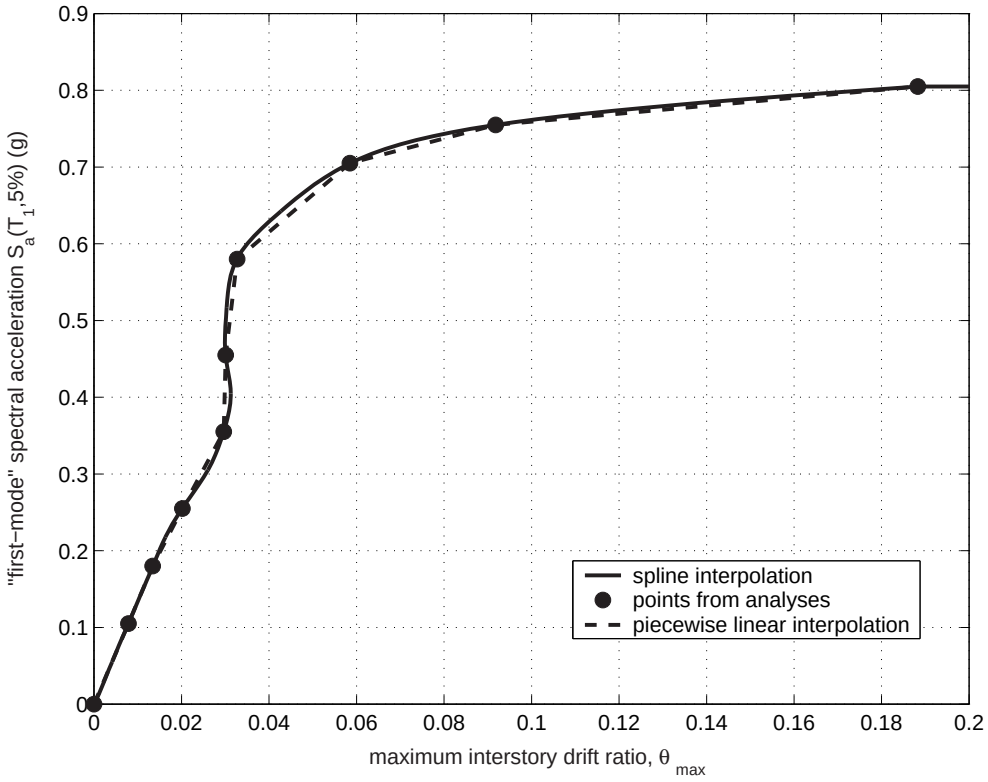

Figure 1. The numerically-converging dynamic analysis points for record #14, Table 2, are interpolated, using both a spline and a piecewise linear approximation.

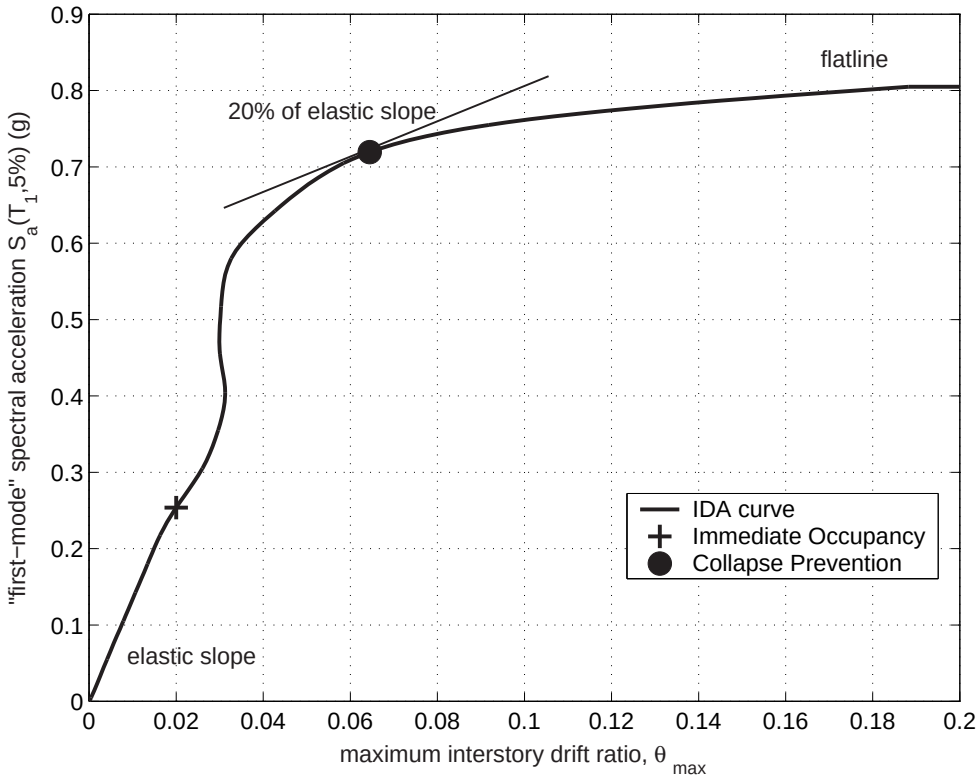

Figure 2. The limit-states, as defined on the IDA curve of record #14.

With the help of Equation 1 we can approximate the *DM*-value at arbitrary levels of *IM* and vice versa. All we need is to solve the appropriate $x_i(t)$ polynomial piece given the value of x to get the parameter t and then replace at the corresponding $y_i(t)$ piece to get the appropriate y-value (*DM*), i.e.,

$$DM = y\left(x^{-1}(IM)\right), \tag{2}$$

$$IM = x\left(y^{-1}(DM)\right), \tag{3}$$

where the -1 superscript denotes the inverse of a function. All these operations only involve polynomials, hence they are trivial to perform, especially if properly coded in a program.

The smooth IDA curve provided by the interpolation scheme offers much to observe. Even for the single record depicted in Figure 1 the IDA curve is not at all simple. It starts as a straight line in the elastic range but then shows the effect of yielding and slightly "softens" at 0.3 g by displaying a tangent slope less than the elastic. Subsequently, it "hardens," having a local slope higher that the elastic, and the building apparently responds with almost the same $\theta_{\max} \approx 3\%$ for $S_a(T_1, 5\%)$ in the range of 0.35 g – 0.55 g. Finally, the IDA curve starts softening again, showing ever decreasing slopes, i.e., greater rates of *DM* accumulation as *IM* increases, reaching the "flatline" at $S_a(T_1, 5\%) \approx 0.81$ g, where the structure responds with practically "infinite" $\theta_{\max}$ values and numerical nonconvergence has been encountered during the analysis. That is when the building has reached global dynamic instability, when a small increment in the *IM*-level results in unlimited increase of the *DM*-response.

Observing Figure 1, it becomes apparent that the relation of *IM* (or x) and t in Equation 1 should always be monotonically increasing. The formulation presented does not strictly enforce this property, but a properly fitted spline will always observe this restriction. Consequently, Equation 2 will always return only one *DM* for a given *IM*. On the other hand, the relation of *DM* and t is often nonmonotonic, due to the occasional hardening of IDA curves, hence Equation 3 may generate more than one *IM* solutions that produce a given *DM*.

DEFINING LIMIT-STATES ON AN IDA CURVE

In order to be able to do the performance calculations needed for PBEE, we need to define limit-states on the IDA curves. For our case study, we chose to demonstrate three: immediate occupancy (IO), collapse prevention (CP), both defined in SAC (2000a,b), and global dynamic instability (GI). For a steel moment-resisting frame with reduced beam section connections, IO is violated at $\theta_{\max}$ = 2% according to SAC (2000a). On the other hand, CP is not exceeded on the IDA curve until the final point where the local tangent reaches 20% of the elastic slope (Figure 2) or $\theta_{\max}$ = 10%, whichever occurs first in *IM* terms (SAC 2000a). The main idea is to place the CP limit-state at a point where the IDA curve is softening towards the flatline but at low enough values of $\theta_{\max}$ so that we still trust the structural model. Finally, GI happens when the flatline is reached and any increase in the *IM* results in practically infinite *DM* response.

Calculating the *IM*-value of the flatline capacity is trivial, as our best estimate is actually somewhere between the highest numerically-converging run and the lowest nonconverging one, as produced by the hunt & fill algorithm. We choose to use the *IM*-value of the highest numerically-converging run as the estimate, e.g., $S_a(T_1, 5\%) = 0.81$ g for record #14. We could have used, for example, the average of the highest converging and lowest nonconverging run, $(0.81 + 0.87)/2 = 0.84$ g, but the difference is negligible and gets smaller and smaller as we increase our capacity resolution in the hunt & fill tracing algorithm.

It is equally easy to calculate the *IM*-values for the IO limit-state; all we need to do is use Equation 3 for $DM \equiv \theta_{\max} = 2\%$, calculate all the *IM*-values that produce $\theta_{\max} = 2\%$ and, if more than one, select the lowest. This is the one that signals the very first exceedance of the limit-state for the given record. For our example of record #14 in Figure 2, IO is violated for $S_a(T_1, 5\%) \geq 0.26$ g or $\theta_{\max} \geq 2\%$.

On the other hand, the CP points are harder to generate, as we need the tangent slope (i.e., the first-order derivative) of the IDA curve to find points where the local stiffness is 20% of the elastic. We also need the curvature of the IDA curve, to discard candidate points that lie on a hardening part of the curve, rather than the desired softening. The cubic spline interpolation is by definition twice differentiable everywhere, so if we use the prime to denote differentiation by the interpolation-parameter t and apply the chain-rule, we can generate the first two derivatives of *IM* (or x) given *DM* (or y):

$$\frac{\mathrm{d}x}{\mathrm{d}y} = \frac{x'}{y'} \tag{4}$$

$$\frac{\mathrm{d}^2x}{\mathrm{d}y^2} = \frac{x''y' - y''x'}{(y')^3} \tag{5}$$

According to the CP limit-state concept, we need to find the highest (in *IM*-value) point where the IDA slope is equal to 20% of the elastic while the point also belongs to a softening branch. Additionally, another candidate point is at $\theta_{\max} = 10\%$; therefore whichever comes first (in *IM*), the slope or the $\theta_{\max}$ limit, decides capacity. Hence, we specify:

$$\left.\frac{\mathrm{d}x}{\mathrm{d}y}\right|_t = 0.20 \left.\frac{\mathrm{d}x}{\mathrm{d}y}\right|_{t=0} \tag{6}$$

$$\left.\frac{\mathrm{d}^2x}{\mathrm{d}y^2}\right|_t < 0 \tag{7}$$

$$t = y^{-1}(10\%), \tag{8}$$

All we need to do is solve for all t satisfying Equation 6 and select the maximum such t (corresponding to the maximum *IM*) that still satisfies Equation 7. Then, we compare against the minimum t that satisfies Equation 8. Whichever is the smallest is the t that defines the CP point. Following this procedure with record #14 we get: $S_a(T_1, 5\%) = 0.72$ g, $\theta_{\max} = 6.4\%$ from Equations 6 and 7, and $S_a(T_1, 5\%) = 0.76$ g, $\theta_{\max} = 10\%$ from Equation 8. By choosing the smallest t, or equivalently the smallest *IM*, we end up with the first of the two points, i.e., in this case the slope limit defines CP (Figure 2).

SUMMARIZING THE IDAS

By generating the IDA curve for each record and subsequently defining the limit-state capacities, a large amount of data can be gathered, only part of which is seen in Figure 3. There, the IDA curves display a wide range of behavior, showing large record-to-record variability, thus making it essential to summarize such data and quantify the randomness introduced by the records. We need to employ appropriate summarization techniques that will reduce this data to the distribution of *DM* given *IM* and to the probability of exceeding any specific limit-state given the *IM* level.

The limit-state capacities can be easily summarized into some central value (e.g., the mean or the median) and a measure of dispersion (e.g., the standard deviation, or the difference

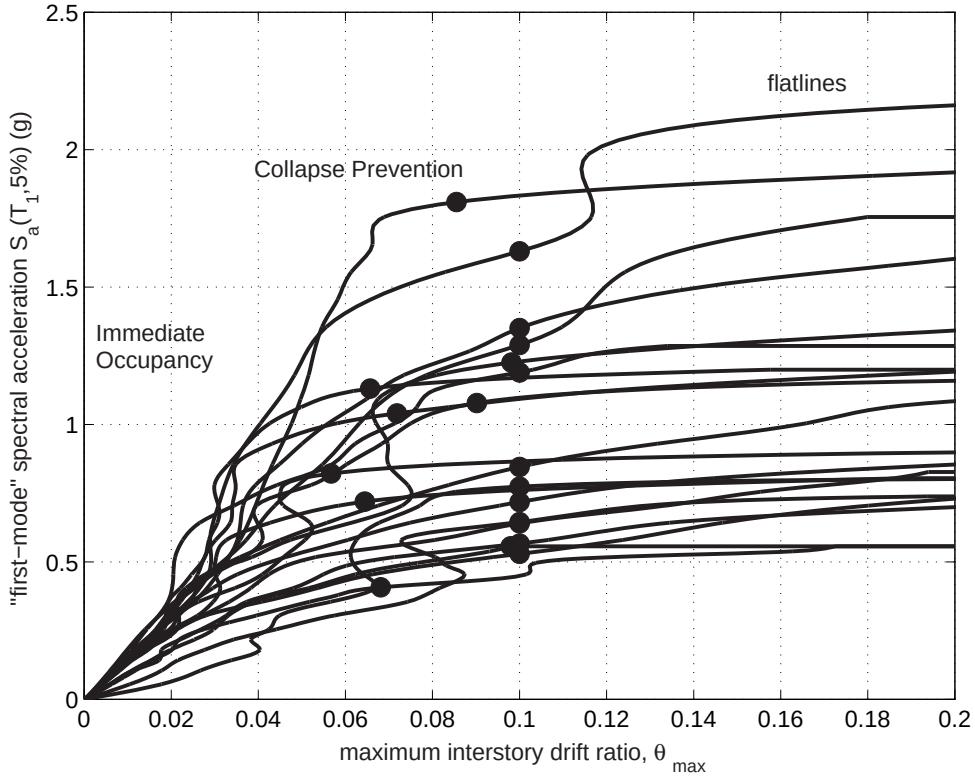

Figure 3. All twenty IDA curves and the associated limit-state capacities. The IO limit is at the intersection of each IDA with the $\theta_{\max} = 2\%$ line, the CP limit is represented by the dots, while GI occurs at the flatlines.

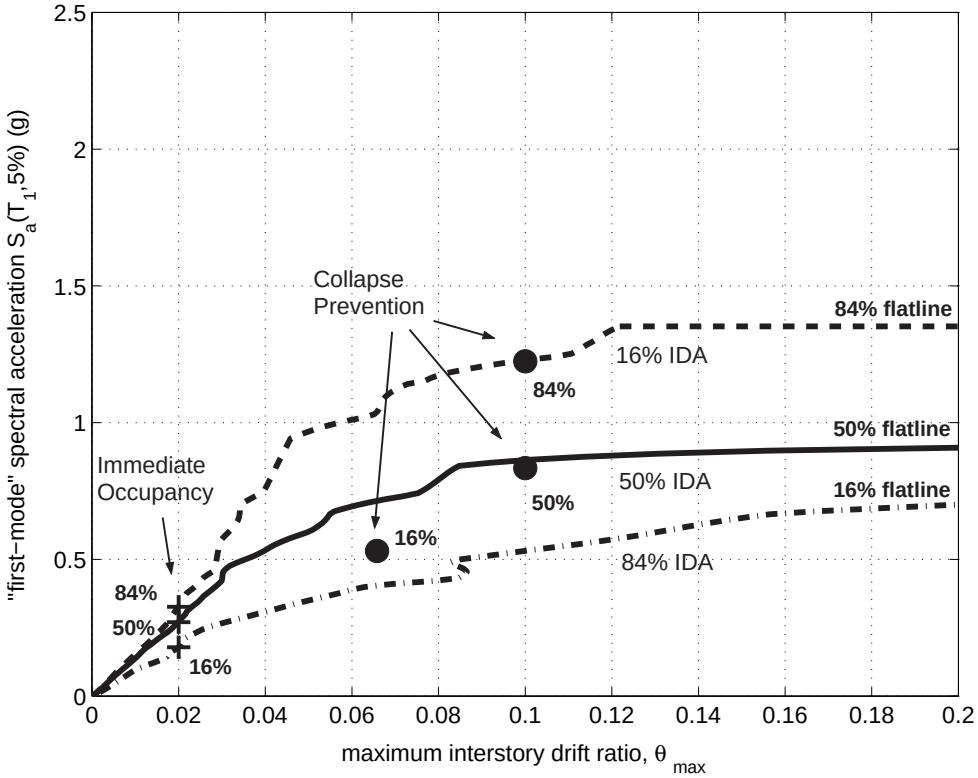

Figure 4. The summary of the IDA curves and corresponding limit-state capacities into their 16%, 50%, and 84% fractiles.

between two fractiles). Consequently, we have chosen to calculate the 16%, 50%, and 84% fractile values of DM ($DM^c_{16\%}$, $DM^c_{50\%}$, and $DM^c_{84\%}$, respectively) and IM ($IM^c_{16\%}$, $IM^c_{50\%}$, and $IM^c_{84\%}$, respectively) for each limit-state, as shown in Table 3, and also graphically depicted in Figure 4. For example, reading off Table 3, at $S_a(T_1,5\%) = 0.83$ g or equivalently at $\theta_{\max} = 0.10$, 50% of the ground motion records have forced the nine-story structure to violate CP.

Table 3. Summarized capacities for each limit-state

	$S_a(T_1,5\%)$ (g)			$\theta_{\max}$		
	$IM^c_{16\%}$	$IM^c_{50\%}$	$IM^c_{84\%}$	$DM^c_{16\%}$	$DM^c_{50\%}$	$DM^c_{84\%}$
IO	0.18	0.27	0.33	0.02	0.02	0.02
CP	0.57	0.83	1.29	0.07	0.10	0.10
GI	0.74	0.91	1.35	$+\infty$	$+\infty$	$+\infty$

There are several methods to summarize the IDA curves, but the cross-sectional fractiles are arguably the most flexible and robust with respect to the infinite DMs introduced by the flatlines (Vamvatsikos and Cornell 2002a). Using the spline interpolation we can generate stripes of DM-values at arbitrary levels of the IM; each stripe contains twenty DM-values, one for each record, that may be finite or even infinite when a record has already reached its flatline at a lower IM-level. By summarizing the DM-values for each stripe into their 16%, 50%, and 84% percentiles, we get fractile values of DM given IM that are in turn interpolated for each fractile to generate the 16%, 50%, and 84% fractile IDA curves, shown in Figure 4. For example, given $S_a(T_1,5\%) = 0.4$ g, 16% of the records produce $\theta_{\max} \leq 2.3\%$, 50% of the records $\theta_{\max} \leq 2.5\%$, and 84% $\theta_{\max} \leq 6.5\%$. Under suitable assumptions of continuity and monotonicity of the IDA curves (as shown at a later section), the fractiles can also be used in the inverse way, e.g., in order to generate demand $\theta_{\max} = 4\%$, 84% of the records need to be scaled at levels $S_a(T_1,5\%) \geq 0.31$ g, 50% of the records at $S_a(T_1,5\%) \geq 0.52$ g, and 16% at $S_a(T_1,5\%) \geq 0.76$ g. Consequently, the 16%, 50%, and 84% IO points and GI flatlines actually reside on the 84%, 50%, and 16% IDA curves respectively, a direct result of the definition of these limit-states. On the other hand, no such general property exists for the CP points, but experience has shown that they usually lie quite close and often on top of their corresponding fractile IDAs, just like the other limit-state points.

PBEE CALCULATIONS

One of the goals of PBEE is producing mean annual frequencies (MAFs) of exceedance for the limit-states. This can be easily accomplished with the summarized results that have been calculated so far, especially if one considers the formats proposed by SAC/FEMA (SAC 2000a,b) or by the Pacific Earthquake Engineering Research Center (Cornell and Krawinkler 2000). The process invariably involves calculating the MAF of exceeding values of the chosen IM, readily available for $S_a(T_1,5\%)$ from conventional PSHA, and appropriately integrating with the conditional probabilities of exceeding each limit-state (given the IM or DM level) to produce the desired MAFs of limit-state exceedance. It is a relatively straightforward method that has been described in extent, for example, by Cornell et al. (2002). However, the complete process calls for the inclusion of the effects of epistemic uncertainties, e.g., due to the imperfect model used (e.g., see Cornell et al. 2002, Baker and Cornell 2003). In the interests of simplicity we will leave this step out, thus the final numerical results should be interpreted with this fact in mind.

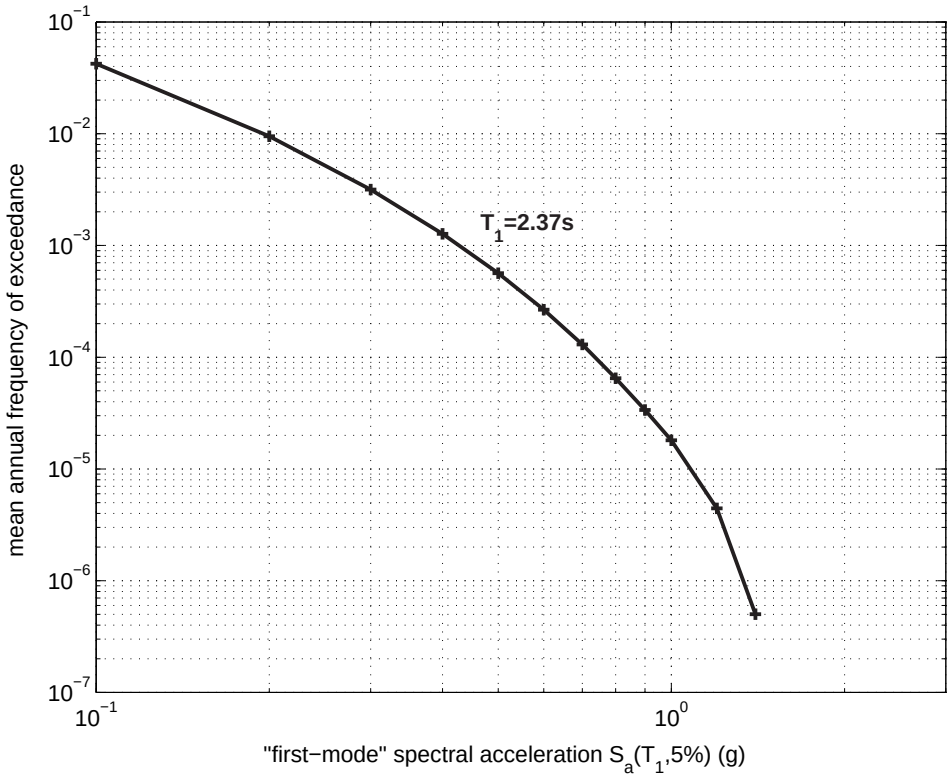

Figure 5. Hazard curve for the Van Nuys Los Angeles site, for $S_a(2.37\text{ s}, 5\%)$.

Here we will perform MAF calculations using a form of the framing equation adopted by the Pacific Earthquake Engineering Research Center (Cornell and Krawinkler 2000, Vamvatsikos and Cornell 2002a),

$$\lambda(DV) = \iint G(DV|DM)\ |\mathrm{d}G(DM|IM)|\ |\mathrm{d}\lambda(IM)| \tag{9}$$

To simplify the above equation and the ones to follow, we will loosely use $\lambda(X)$, $F(X)$, and $G(X)$ to denote the MAF function, cumulative distribution function (CDF), and the complementary CDF (CCDF), respectively, of their arguments. For example, $\lambda(X)$ actually means $\lambda_X(x)$ and is a different function from $\lambda(Y) \equiv \lambda_Y(y)$.

In this paper we have generally used $S_a(T_1, 5\%)$ for the *IM* and $\theta_{\max}$ as *DM* for the limit-states of interest. The decision variable, *DV*, here is simply a scalar "indicator variable": $DV = 1$ if the limit-state is exceeded (and zero otherwise). $\lambda(IM) \equiv \lambda_{IM}(x)$ is the conventional hazard curve, i.e., the MAF of *IM* exceeding, say, x. $|\mathrm{d}G(DM|IM)|$ is the differential of the (conditional) CCDF of *DM* given *IM*, or $f_{DM|IM}(y|x)\,\mathrm{d}y$, i.e., it is the probabilistic characterization of the distribution of *DM* given *IM*, offered by the fractile IDAs. Finally in the limit-state (LS) case, when on the left-hand side of Equation 9 we seek the MAF of exceeding the limit-state, $\lambda(DV{=}1) = \lambda(0) = \lambda_{\mathrm{LS}}$ and $G(0|DM)$ becomes simply the probability that the capacity DM^c is less than some level of the *DM*; so $G(0|DM) = F(DM^c|DM)$, where $F(DM^c|DM)$ is the CDF of DM^c, i.e., the statistical characterization of the *DM*-value of capacity, as offered, e.g., by the fractiles of *DM*-capacity.

Thus, for our purposes, we can modify Equation 9 to become:

$$\begin{aligned}\lambda_{LS} &= \iint G(0|DM)\ \left|\mathrm{d}G(DM|IM)\right|\ \left|\mathrm{d}\lambda(IM)\right| \\ &= \int_{DM=0}^{DM=+\infty} F(DM^c|DM)\left\{\int_{IM=0}^{IM=+\infty}\left|\frac{\mathrm{d}G(DM|IM)}{\mathrm{d}DM}\right|\left|\frac{\mathrm{d}\lambda(IM)}{\mathrm{d}IM}\right|\ \mathrm{d}IM\right\}\mathrm{d}DM \\ &= \int_{DM=0}^{DM=+\infty} F(DM^c|DM)\left|\frac{\mathrm{d}\lambda(DM)}{\mathrm{d}DM}\right|\ \mathrm{d}DM \end{aligned} \tag{10}$$

where the integration over IM in the braces needs to be carried out either numerically or by an appropriate analytic approximation (Cornell et al. 2002) to produce the absolute value of the DM hazard gradient $|\mathrm{d}\lambda(DM)/\mathrm{d}DM|$. Then we can proceed to integrate over DM and estimate λ_{LS}. If, on the other hand, we first integrate-out the DM, then we can rewrite the above equation to use the IM-value of capacity:

$$\begin{aligned}\lambda_{LS} &= \int G(0|IM)\ \left|\mathrm{d}\lambda(IM)\right| \\ &= \int_{IM=0}^{IM=+\infty} F(IM^c|IM)\left|\frac{\mathrm{d}\lambda(IM)}{\mathrm{d}IM}\right|\ \mathrm{d}IM \end{aligned} \tag{11}$$

where the quantity in the absolute value is the IM hazard gradient and $F(IM^c|IM)$ is the CDF of the IM-value of limit-state capacity. In this case, all quantities in Equation 11 are known, and only one integration is needed to calculate λ_{LS}.

We can proceed to the MAF calculations using either the DM-form (Equation 10) or the pure IM-form (Equation 11). There are several issues of compatibility with current guidelines (e.g., SAC 2000a) that may dictate the use of the DM-approach, otherwise the IM-form is more attractive, as it needs only one integration rather than two; hence, it will be our method of choice. Still, it must be emphasized that either of the two approaches should provide the exact same results if the integrations are performed with sufficient accuracy (see also Jalayer and Cornell 2002). These are just two ways to the same goal, and the choice lies with the user.

Table 4. MAFs of exceedance for each limit-state, calculated both numerically from Equation 11 and with the approximate analytical form of Equation 12, using either a global or a local fit to the IM-hazard curve

	IO	CP	GI
numerical	0.019	0.0004	0.00010
analytical (global fit)	0.017	0.0002	0.00003
analytical (local fit)	0.008	0.0005	0.00040

The MAF calculations for any of the two approaches can be carried out either numerically or with an analytical approximation. If a high degree of accuracy is desired, a trapezoidal rule can be employed to directly integrate Equation 11. All we need to do is assign $1/20$ probability to each of the 20 records, then derive the empirical CDF of the IM-value of capacity and numerically integrate with values of the hazard curve slope, calculated either by differentiating a smooth interpolation or by simply reading them off Figure 5. On the other hand, if we make some reasonable approximations, Equation 11 can be analytically integrated (Shome and Cornell 1999, Cornell et al. 2002). We only need to assume that the IM-values of capacity are

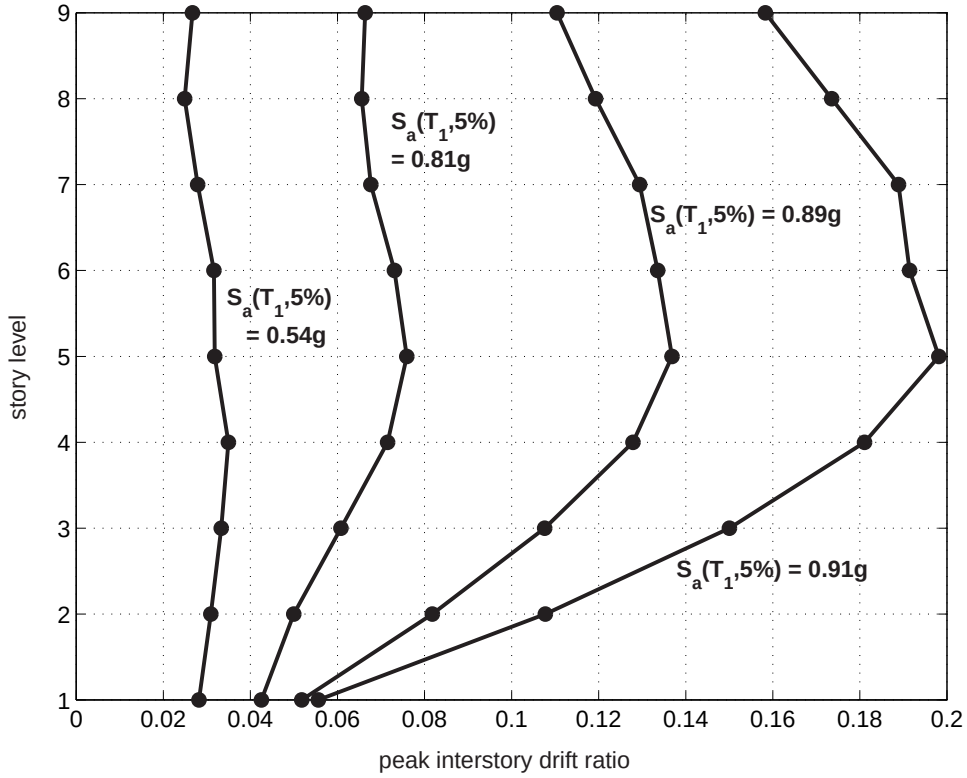

Figure 6. The median peak interstory drift ratios for all stories at several specified $S_a(T_1, 5\%)$ levels.

lognormally distributed and then approximate the *IM*-hazard curve by fitting a straight line in the log-log space, $\lambda(IM) = k_0 IM^{-k}$, either by a global regression, same for all limit-states, or by a local fit at the median *IM*-capacity for each limit-state. Then we arrive at the equation

$$\lambda_{\mathrm{LS}} = \lambda(IM^c_{50\%}) \cdot \exp\left(\frac{1}{2}(k \cdot S_{\ln IM^c})^2\right) \tag{12}$$

where $S_{\ln IM^c} = \left(\ln IM^c_{50\%} - \ln IM^c_{16\%}\right)$ is (approximately) the standard deviation of the natural logarithm of the *IM*-capacity.

As an example, the MAFs of exceeding each of the three limit-states (IO, CP, and GI) were calculated using both the approximate analytic approach (with either the global or the local fit to the hazard curve) and the "exact" numerical integration (Table 4). In general, it seems that by approximating the hazard curve with a global fit, the MAFs are consistently underestimated. On the other hand, the local fit seems to cause overestimation for all limit-states but IO. The approximations may sometimes miss the MAFs by a factor of three or get as close as 10%. Still, the large record-to-record variability coupled with the limited size of our suite of twenty records may generate considerable standard errors around these estimates, possibly making the approximate results statistically indistinguishable from the exact MAF for some limit-states. This is an issue that is going to be investigated in a later section.

TAKING ADVANTAGE OF THE DATA: SPO VERSUS IDA

Beyond the essential calculations needed for PBEE, there is much more information that we could easily glean out of the IDA by taking a closer look at the results and plotting them in new ways. For example, Figure 6 displays a story-to-story profile of the median peak interstory drift

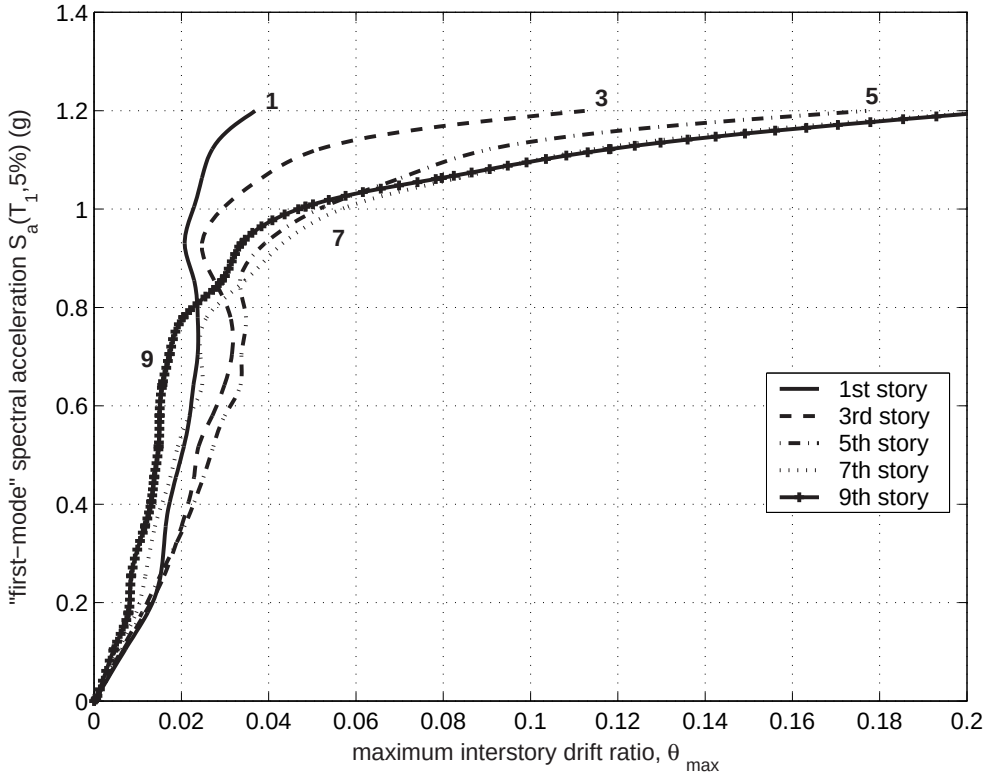

Figure 7. The IDA curves of the odd stories for record #1.

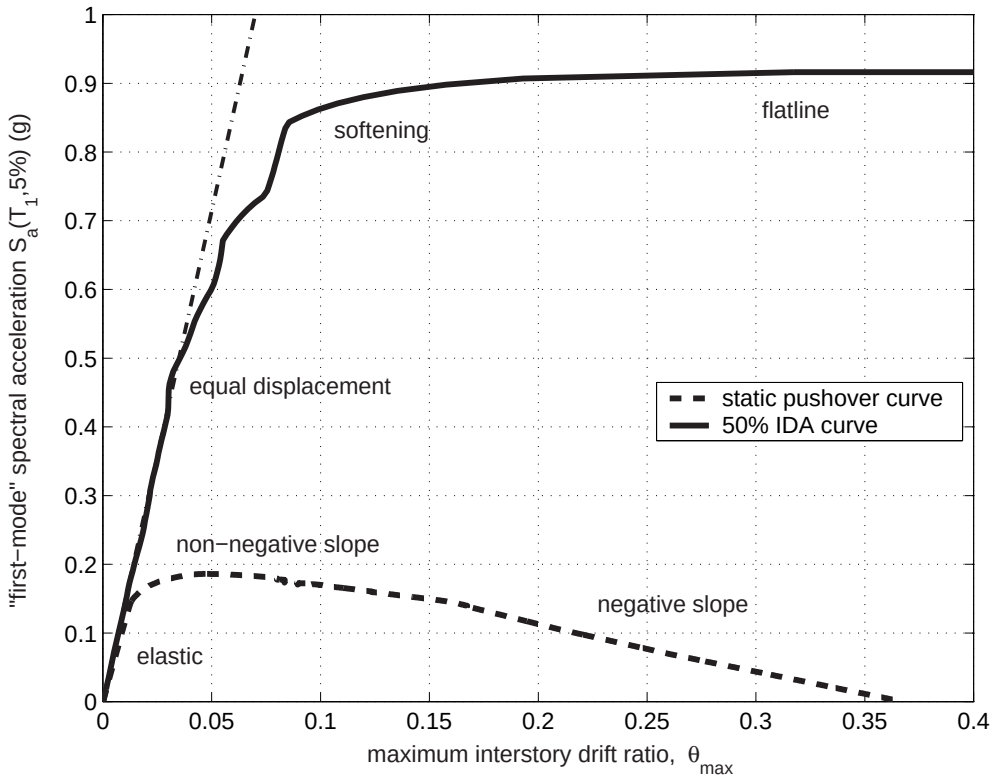

Figure 8. The SPO curve generated from an inverse-triangle (maximum force at roof) load pattern versus the median IDA.

ratios at several $S_a(T_1, 5\%)$-levels. As the intensity increases, then, in a median sense across all records, the fifth floor seems to accumulate most of the deformation. On the other hand, in Figure 7 the individual story drift IDA curves are plotted for record #1, showing a record-specific picture of the odd-numbered stories. Most interesting for this record is the sudden change of behavior that occurs around $S_a(T_1, 5\%) = 0.82$ g, when the top floors suddenly start accumulating more and more deformation as *IM* increases, while the previously leading lower floors are held back, displaying almost constant peak interstory drifts.

It is also very informative to visually compare on the same figure the static pushover (SPO) curve (also known as the nonlinear static procedure curve, e.g., ATC 1997) versus the median (50%-fractile) IDA. Since the SPO curve usually comes in base shear versus θ_{roof} (peak roof drift ratio) coordinates, it needs to be transformed into *IM* and *DM* axes. In our case, the $\theta_{\max}$ response can be easily extracted from the SPO analysis results, while the base shear can be converted to acceleration units by dividing with the building mass times some (ad hoc) factor chosen to make the curves match in their elastic range. This can be achieved for our structure by dividing the base shear with 85% of the total building mass (which is very close to the ratio of the first modal over the total mass). By thus plotting the two curves together, as pictured in Figure 8, we see that they correspond to each other. The elastic region of the IDA matches the SPO by construction, and the post-yield non-negative SPO segment corresponds to a continuation of the elastic region in the IDA, where the IDA is following the familiar "equal displacement" rule for moderate period structures (Veletsos and Newmark 1960). When the SPO turns into a negative slope, the IDA softens and acquires a local slope less than the initial elastic that gradually decreases till the IDA becomes flat. Essentially, the ending of the SPO at zero strength signals the end of the IDA by the flatline.

The question then arises as to why this relationship exists. Some light can be shed on this issue if we simplify the problem and think in terms of a single-degree-of-freedom system with a force-deformation backbone that has the shape of the building's SPO curve. Then we need to realize that in terms of dynamics, where the IDA is concerned, an ascending part of the "backbone" generally means a "dynamically stable" part while a descending branch corresponds to a "dynamically unstable" part (e.g., Macrae and Kawashima 1997). For each dynamic run the $\theta_{\max}$ value serves as an indicator of whether the building has remained completely in the ascending parts (approximately $\theta_{\max} < 5\%$ in Figure 8) or it has ventured into the descending branch as well. So, for lower *IM*s, approximately $S_a(T_1, 5\%) < 0.6$ g in Figure 8, the building (in a median sense, i.e., for at least 50% of the records) oscillates along the ascending part of its "SPO backbone" thus the increase in $\theta_{\max}$ is controlled and stable in the median. But for higher *IM*s the building (in a median sense again) also sustains more and more cycles in the descending part of the "SPO backbone," thus the median $\theta_{\max}$ increases uncontrollably towards infinity. This can help us understand why the behavior of the median IDA changes so drastically when the median $\theta_{\max}$ is higher than 5%, as it shifts from an ascending to a descending branch. On the contrary, the median IDA remains virtually indifferent when this moderate period structure passes from the elastic part to the non-negative post-yield segment of the SPO, since both are ascending branches.

Observing these facts, one could stipulate that some more direct, perhaps quantitative rules may be devised to connect the two curves. Actually, one such attempt has been tried out both for single (Vamvatsikos and Cornell 2004b) and multi-degree-of-freedom systems (Vamvatsikos and Cornell 2004a) with encouraging results.

DISCUSSION OF CHOICES AND THEIR INFLUENCE ON IDA

We took the reader through a direct, hands-on example of how to perform IDA and apply it for the purposes of PBEE. At each step we had to make some choices, e.g., how to set up the dynamic analysis algorithm, what tracing algorithm and interpolation scheme to use, how to summarize the IDAs (using stripes given the *IM* instead of stripes given the *DM*) or how many records and how many runs per record to allow. Still, we chose not to focus on such details; instead we proceeded by making seemingly ad hoc choices. Now, armed with the knowledge of the complete IDA process, we can discuss such choices, explain the reasons behind them and understand their influence to the final results.

NUMERICAL CONVERGENCE

The details of the analysis and the structural model play an important role in generating accurate IDA curves. Especially in the region of global dynamic instability, the very existence of the flatline and the associated numerical nonconvergence may often generate several accuracy problems. Ideally, the structural model would be composed of (numerically) robust and well-tested elements, while the dynamic analysis algorithm should be able to accurately track the structural response through, e.g., yielding events, sharp strength drops, load redistribution and geometric nonlinearities; it would fail to converge only when the structure has exhausted its reserves to become dynamically unstable, thus correctly matching global dynamic instability with numerical nonconvergence. Unfortunately, most algorithms and element models have not really been designed or tested to operate in such extreme ranges of behavior. As a result, some records may cause a premature nonconvergence, creating a characteristic halting of the IDA curve which does not resemble a flatline.

All the flatlines in our model normally occur beyond $\theta_{\max} = 12\%$ (Figure 3), meaning that the model can remain stable at least up to such $\theta_{\max}$-values. Still, in our initial attempt to trace the IDA curves, two of the twenty records failed prematurely, at $\theta_{\max} \approx 2\%$, barely past the end-of-elasticity value of $\theta_{\max} \approx 1\%$. The main reason is the use of a large, complex model with many degrees of freedom, plus the adoption of the fracturing connection model (Shi and Foutch 1997) with sharp strength drops that probably tend to destabilize the solution algorithm of DRAIN-2DX (Prakhash et al. 1992). Further confirmation is provided by the SPO-to-IDA connection, as $\theta_{\max}$ values in the order of 2% are still on the ascending branch of the SPO in Figure 8, thus deemed unable to cause collapse. Actually, each IDA curve should be able to behave stably at least up to the start of the SPO's negative slope, at about $\theta_{\max} = 7\%$. Still, this comparison should not be carried too far; while the SPO ends at $\theta_{\max} = 37\%$, the post-peak part of the SPO is often very load-pattern dependent, and an arbitrary load pattern may result in very optimistic $\theta_{\max}$-values that do not reflect the dynamic behavior (Vamvatsikos and Cornell 2004a).

Such illegitimate and premature collapses are thus relatively easily identified, either by looking at the SPO or at the IDAs of other records, but how are they to be fixed? Of course, if the model or the elements or the algorithm are deficient in the first place, the situation is hopeless. Experience has shown that this is not the case with the well-tested DRAIN-2DX; it is more a problem of correctly setting up the analysis parameters, rather than anything else. The best strategy is to tweak the analysis knobs, e.g., reduce the integration time-step, adopt a variable-step solution or experiment with the parameters of the event-to-event solver. In our case study, the dynamic analyses of the two problematic records had to be repeated at a reduced time-step, one-fourth instead of one-half of the acceleration timehistory time-step, thus easily

resolving convergence issues. Note that after such false, premature collapses are dealt with, then further (reasonable) changes in the parameters of the solution algorithm will make only small arbitrary changes to the IDA results. For example, it has been found empirically that changing the integration time-step can incur arbitrary changes of up to 10% to the flatline heights, where, surprisingly, smaller steps do not necessarily mean more stability (i.e., higher flatline heights). This is simply the effect of small errors piling up on each step of the time-integration that may affect convergence when the structure is close to the flatline, sometimes causing it to collapse a bit earlier and sometimes not. This is the reason why when tracing each record, we specified a capacity resolution of only 10%; a better accuracy does not have much meaning in the presence of these analysis uncertainties.

Such inaccuracies remain relatively insignificant when good analysis software is used. Actually, we cannot stress enough the need for reliable, bug-free algorithms and well-tested, robust element models. Such tools are exactly what makes the difference in such analyses, especially for the limit-states close to global dynamic instability, and when available, with only a little attention to the analysis details allow us to easily obtain accurate IDA curves.

CHOICE OF TRACING ALGORITHM

When tracing the IDA curve for each record, the choice of the *IM*-level for each run is a decision left to the automated tracing algorithm that we use. We have theoretically argued about the superiority of the hunt & fill algorithm versus the use of a constant *IM*-step (i.e., the stepping algorithm) in Vamvatsikos and Cornell (2002a), so it is time to see in detail what the true differences really are when both are applied to the nine-story structure.

Before we proceed, keep in mind that given the same structural model, analysis program and computing platform, still not all runs are equal in computational cost. In general, the closer the run is to the flatline (either at a lower or a higher *IM*) the longer it takes to complete the analysis. On the other hand, both converging and nonconverging runs that are far away from the flatline will be significantly faster, as convergence or nonconvergence will be achieved within a minimum of iterations. Still, when comparing the tracing algorithms, we will assume that the intent is to trace the whole IDA curve and a similar amount of runs will be spent both high and low in the curve (in *IM* terms). Thus, looking at each record as a whole, the total amount of runs (converging and nonconverging alike) spent for it provide a very accurate idea of the computational time needed, while the number of converging runs accurately describes the accuracy achieved.

The most important task that a user faces when applying either of the two algorithms is setting up their parameters correctly. For the stepping algorithm, the only parameter is the step size, while for the hunt & fill the most important ones are the initial step, the step increment and the allowed number of runs per record. Both algorithms were used to trace the IDAs of the nine-story structure for the suite of 20 records, using various settings for their parameters, the results shown in Table 5. Obviously, changing the step size of the stepping algorithm generates huge differences in the total number of runs. Still, if we let the minimum number of converging runs generated for any of the 20 records be our standard for accuracy, we need at least a step size of 0.05 g, or 475 runs to get at least 11 runs per record and reach the standards of hunt & fill. Also, if we do not set the stepping size correctly, we either get too few or too many runs, the resolution easily dropping to 1 or 2 runs in the worst case, if we happen to set a step size of 0.2 g or 0.3 g. On the other hand, we can change the initial step or the step increment for the hunt & fill within a wide range, increasing them or reducing them by two or four times, and the

Table 5. Comparing the sensitivity to parameters of the stepping versus the hunt & fill algorithm

Algorithm	Parameter (g)	Total C+NC [1]	min C [2]	max C [2]	average C [2]
stepping	0.05	475	11	45	22.8
(step-size sensitivity)	0.075	318	7	30	14.9
	0.1	244	5	22	11.2
	0.2	128	2	11	5.4
	0.3	87	1	7	3.4
hunt & fill	0.05	280	12	13	12.2
(initial-step sensitivity)	0.1	280	12	13	12.2
	0.2	280	11	13	12.0
	0.3	280	11	12	11.8
hunt & fill	0.025	280	11	13	12.4
(step-increment sensitivity)	0.05	280	12	13	12.2
	0.1	280	11	13	12.0
	0.2	280	11	13	11.9

[1] Converging and nonconverging runs for all records
[2] Converging runs per record

hunting algorithm remains practically unchanged, constantly providing at least 11 converging runs per record. In essence, it has the right knobs to be tuned to the tolerance limits that we wish and allows us to do the runs the way we want, not the way nature decides through the records.

Still, one may notice that if we over-increase the initial step or the step increment, then the accuracy starts to slowly drop, as the algorithm overshoots the flatline by a lot and spends many nonconverging runs to find its way down. But still the effect is minor, not overwhelming. Notice also that keeping both parameters relatively small seems to improve accuracy both on average and in the minimum. Still, we should not decrease them too much because as the steps become smaller we are risking expending all the allotted runs before reaching the flatline.

Coming back to our example, in Table 2, we used $S_a(T_1, 5\%)$ to measure the *IM*-value for our runs. Why not another *IM*? We could have used pretty much any monotonic and scalable *IM* (Vamvatsikos and Cornell 2002a) that we might want, but the less efficient it is, the further dispersed the IDA flatlines would be, and we would start having some resolution discrepancies within tracing, i.e., a greater difference between the observed number of minimum and maximum convergent runs per record in our suite. By using at least $S_a(T_1, 5\%)$, we are assured that our algorithm, be it hunt & fill or stepping, will be efficient for a wide range of conditions. If another, more efficient *IM* appears that can drastically reduce the record-to-record flatline variability, then the hunt & fill would only marginally benefit, but the stepping algorithm would significantly improve. In conclusion, the hunt & fill procedure desensitizes IDA from the *IM* selection and the setting of the algorithm's parameters, easily achieving the desired resolution, in contrast to the very sensitive stepping algorithm. Additionally, it fixes the number of total runs performed, so we can plan ahead and assign enough computers to run in parallel so the IDA is computed in time.

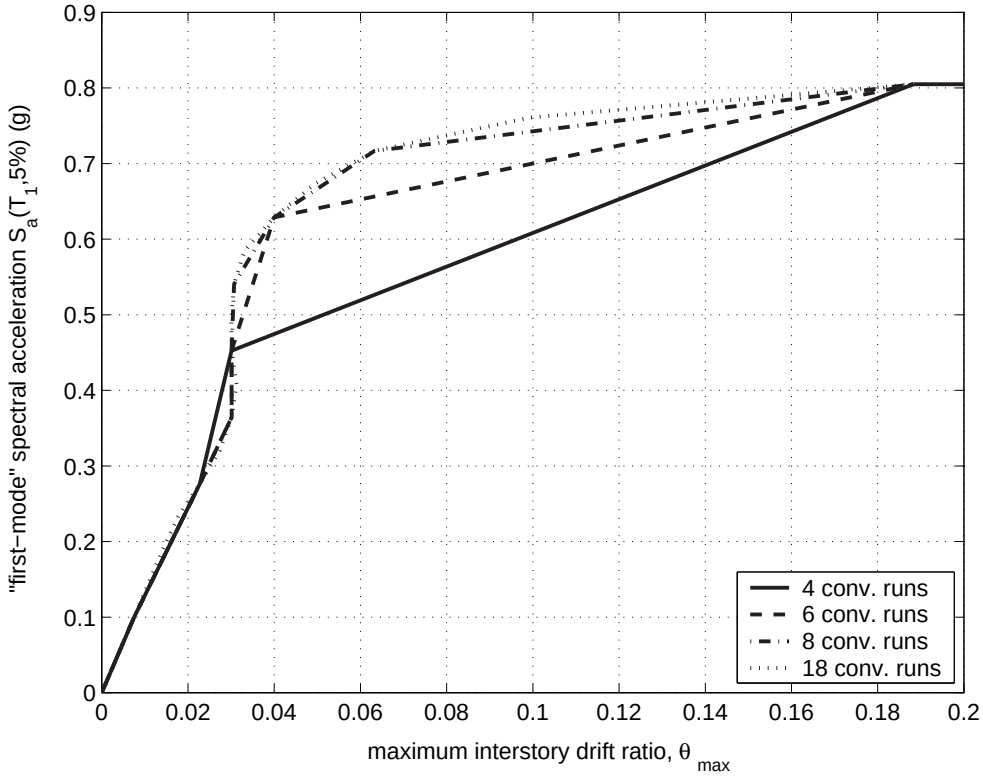

Figure 9. Linearly interpolated IDA curve for record #14, traced with a different total number of converging runs.

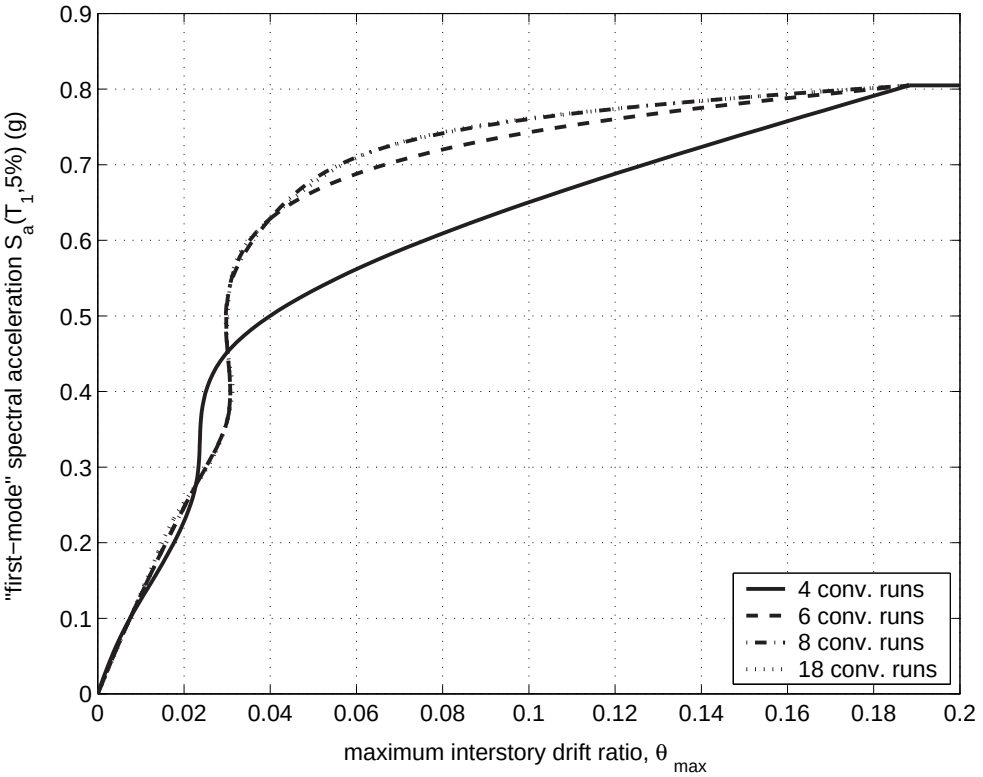

Figure 10. Spline-interpolated IDA curve for record #14, traced with a different total number of converging runs.

INTERPOLATION ISSUES

By interpolating the discrete points to generate each record's IDA curve we are gaining one enormous advantage: we do not need to have our runs in stripes of the same *IM*-level. The consequences are very important. First, this allows us to use the hunt & fill algorithm instead of the stepping one, thus gaining in all the aspects described previously. Second, it allows us to express the IDA results in any *IM*. All we need is to calculate the new *IM* for each run, replot the IDA curves and re-interpolate versus the new *IM*. In this way, IDA becomes truly independent of the *IM* used for tracing, allowing us to reuse the same data again and again, without needing to run any more analyses.

But why use complex spline schemes instead of the simpler linear interpolation? In Figures 9 and 10, we present a comparison of the linear and the spline interpolation scheme, pitted against each other for a given number of converging runs. We have tweaked the hunt & fill tracing so that in all cases the flatline is accurately reached with the same number of runs, and then the algorithm is allowed to fill in the gaps using up to a total of 4, 6, 8, or 18 runs. Unless we are only interested in the flatline, 4 converging runs are just too few to capture the IDA adequately, regardless of the interpolation scheme used. Clever postprocessing cannot make up for gross data gaps. On the other hand, if we allow only 2 more runs, for a total of 6, the results are markedly better, but only if we are using a spline scheme. Had we used 6 linearly interpolated runs we would be grossly underestimating the CP limit-state capacity, finding a capacity point at only $S_a(T_1, 5\%) = 0.63$ g, $\theta_{\max} = 4\%$ instead of the correct 0.72 g and 6.4%. At 8 and 18 runs, the spline interpolations are practically indistinguishable, while the linear ones are close enough but still can be told apart. In conclusion, if we allow enough runs, the interpolation scheme doesn't really matter, both schemes will provide good results. On the other hand, if we use too few runs, it doesn't really matter again because both schemes are going to give us bad results. But there is a gray area in between, where using a better and smarter interpolation can make the difference to increase the accuracy in our final IDA curve. In retrospect, this is precisely what gives us confidence to reduce the allotted number of runs and save on computational resources.

SENSITIVITY OF THE LIMIT-STATE CAPACITIES TO THEIR DEFINITION

Several limit-states were defined on the IDA curves, often through the use of ad hoc rules. For example, the IO limit-state was defined at $\theta_{\max} = 2\%$, while the CP limit-state was based on the arbitrary 20% fraction of the elastic slope (or stiffness) and the additional $\theta_{\max} = 10\%$ limit (SAC 2000a). On the other hand, the GI limit-state was unambiguously defined to be on the flatline of the IDA, being subject to no such arbitrary rules. Therefore, it is of interest to investigate the sensitivity of the summarized *IM*, *DM* capacities to these choices, both for the IO and the CP limit-state.

For the IO limit-state, the simplicity of the definition makes it easy to understand what is happening. If we look at Figure 4, it is obvious that IO is occurring in the "equal displacement" region of the fractile IDAs, i.e., the fractiles are almost straight lines resembling a continuation of the elastic segment. In turn, this means that moderate changes to the defining value for IO, i.e., from $\theta_{\max} = 2\%$ to 1% or 3% will proportionately increase or decrease the *IM* and *DM* values of capacity.

On the other hand, the definition of the CP limit-state is quite more complicated. The elastic stiffness fraction controls how much the IDA has to soften, i.e., how close to the flatline it can come, before CP is reached. Hence, increasing this fraction will force the CP points

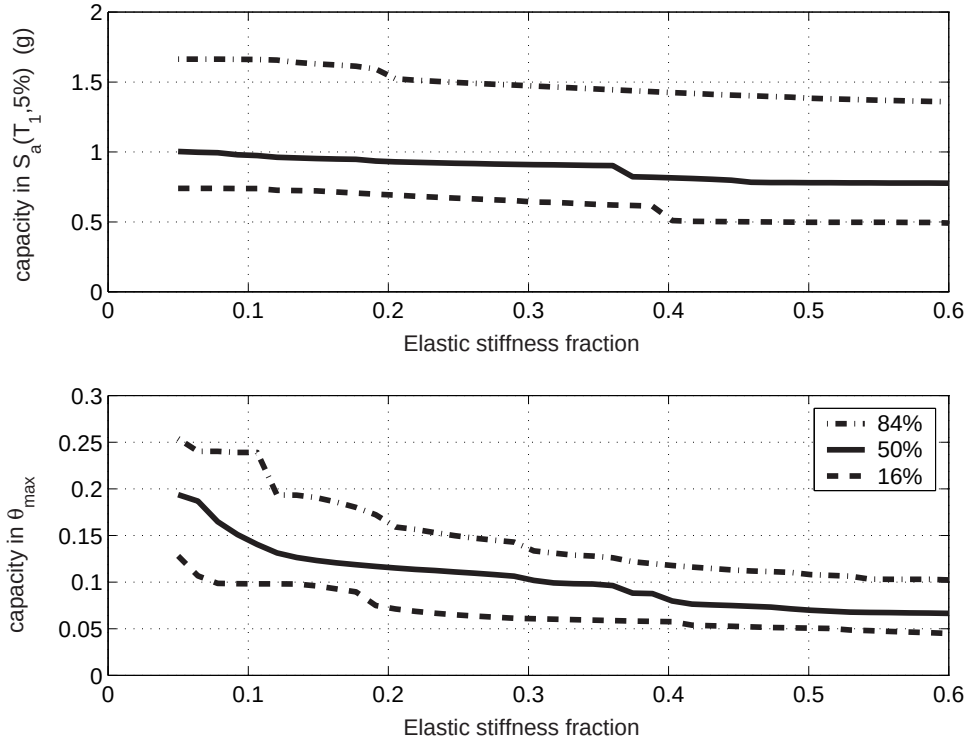

(a) Resulting $S_a(T_1,5\%)$ and $\theta_{\max}$ capacities when the $\theta_{\max} = 10\%$ limit *is not imposed.*

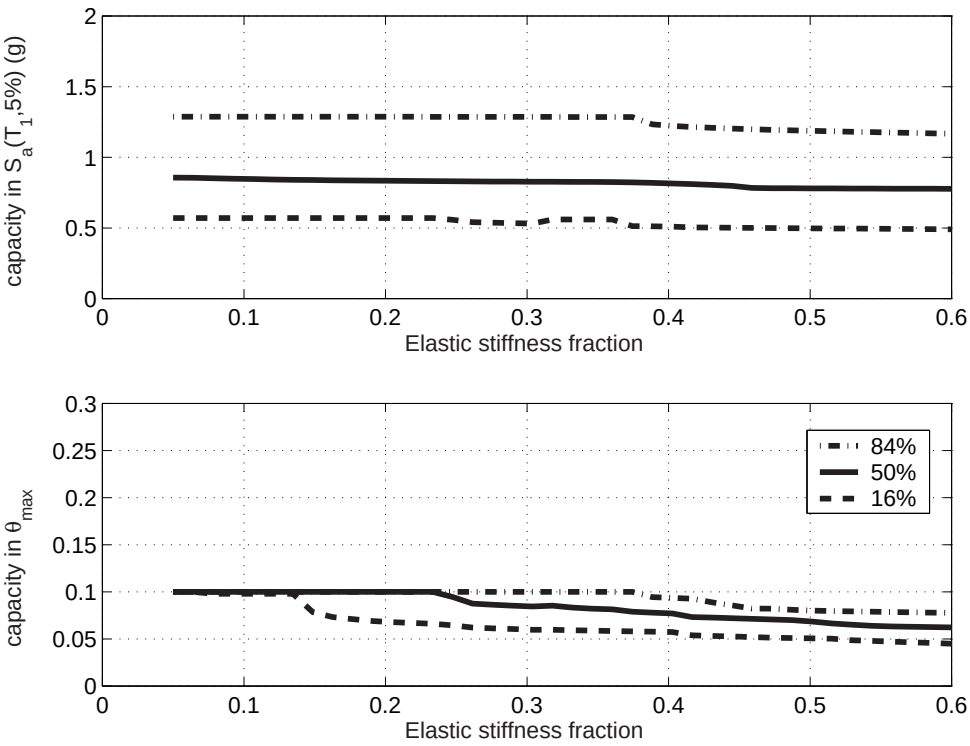

(b) Resulting $S_a(T_1,5\%)$ and $\theta_{\max}$ capacities when the $\theta_{\max} = 10\%$ limit *is imposed.*

Figure 11. The sensitivity of the fractile (16%, 50%, and 84%) $S_a(T_1,5\%)$ and $\theta_{\max}$ capacities for the CP limit-state to the elastic stiffness fraction used (20% is the standard by SAC 2000a). The results are less sensitive if the $\theta_{\max} = 10\%$ limit is used.

(e.g., Figure 3) to move to lower *IM*s and *DM*s. The influence of the $\theta_{\max} = 10\%$ limit is more straightforward. It enforces a rigid limit on the capacity points, restricting the $\theta_{\max}$ value they can reach, i.e., it is another way to restrict the CP points from coming close to the flatline. Actually, in our case of the nine-story building, Figure 3, it becomes obvious that by changing the $\theta_{\max} = 10\%$ limit to, say, 8% or 12%, the *IM*-value of capacity will only slightly change, but the *DM*-value will be highly influenced, the 50% and 84% $\theta_{\max}$ capacities actually becoming 8% or 12%, respectively.

To show the combined influence of the two rules on the CP limit-state, the fraction of the elastic stiffness has been varied from 10% to 60% and the resulting fractile $S_a(T_1, 5\%)$, $\theta_{\max}$ capacities have been plotted, both when the $\theta_{\max} = 10\%$ rule is imposed (Figure 11b) and when it is not (Figure 11a). In the latter case, the *IM* capacity becomes relatively sensitive, almost linearly, to the elastic stiffness fraction. The *DM* capacity is even more sensitive, decreasing in a geometric fashion as the fraction increases. This makes absolute sense given the shape of the IDAs (Figure 3); close to global collapse, each IDA softens towards the flatline, hence, as the slope-fraction decreases, the CP *IM* capacity approaches the flatline *IM*-height. On the other hand, the *DM* capacity is destabilized by the same flattening, since by definition, in the vicinity of the flatline, small changes in the elastic stiffness fraction result to large changes of the *DM*-value.

If we include the $\theta_{\max} = 10\%$ limit, as in Figure 11b, both the *IM* and especially the *DM* capacity are stabilized, as this hard upper limit simply cuts off all higher values. Furthermore, this limit seems to drastically reduce the *DM*-capacity dispersion, at all levels of the elastic stiffness fraction. Obviously, several records now have the same CP limit-state *DM* capacity, namely $\theta_{\max} = 10\%$. Therefore, the 10% limit makes the CP capacity more stable, but no less arbitrary, as the $\theta_{\max} = 10\%$ limit is often the governing rule. Actually, looking at the tables in SAC (2000a,b) it becomes obvious that 10% is often the quoted median $\theta_{\max}$-capacity for all but the tallest buildings. Is, then, this arbitrarily imposed $\theta_{\max} = 10\%$ a problem? From an MAF-sensitivity point-of-view, the answer is negative. In Equation 11 it becomes apparent that it is only the *IM*-value of capacity that truly matters. As we have observed, at least for this structure, the *IM*-value of CP-capacity is only mildly sensitive to the definition of the rules, thus yielding similarly mildly sensitive MAFs. Even if the calculation is done using the *DM*-form in Equation 10, assuming that the integrations are accurately performed, the conclusions will still be the same.

There are also several other details and corresponding sensitivity issues in the implementation of the CP limit-state definition, that may or may not make a difference. For example, in Yun et al. (2002) the 20% fraction is applied to the median elastic stiffness of all records and the resulting reduced stiffness is used for the capacity point estimation. On the other hand, we have used the 20% fraction on the elastic stiffness of each individual record to define its CP capacity. In this case, the summarized capacities show negligible difference between the two approaches. On the other hand, in Yun et al. (2002) CP is defined to occur at the first point where the IDA curve softens to 20% of the (median) elastic slope, while we use the last point where it reaches the reduced stiffness. This may make a large difference for some records that alternately harden and soften before global collapse, and may be interpreted as another sign of sensitivity to the CP definition. Still, for reasons explained in Vamvatsikos and Cornell (2002a), we believe it is more consistent with the CP limit-state concept to use the last rather than the first such point, thus resolving this problem.

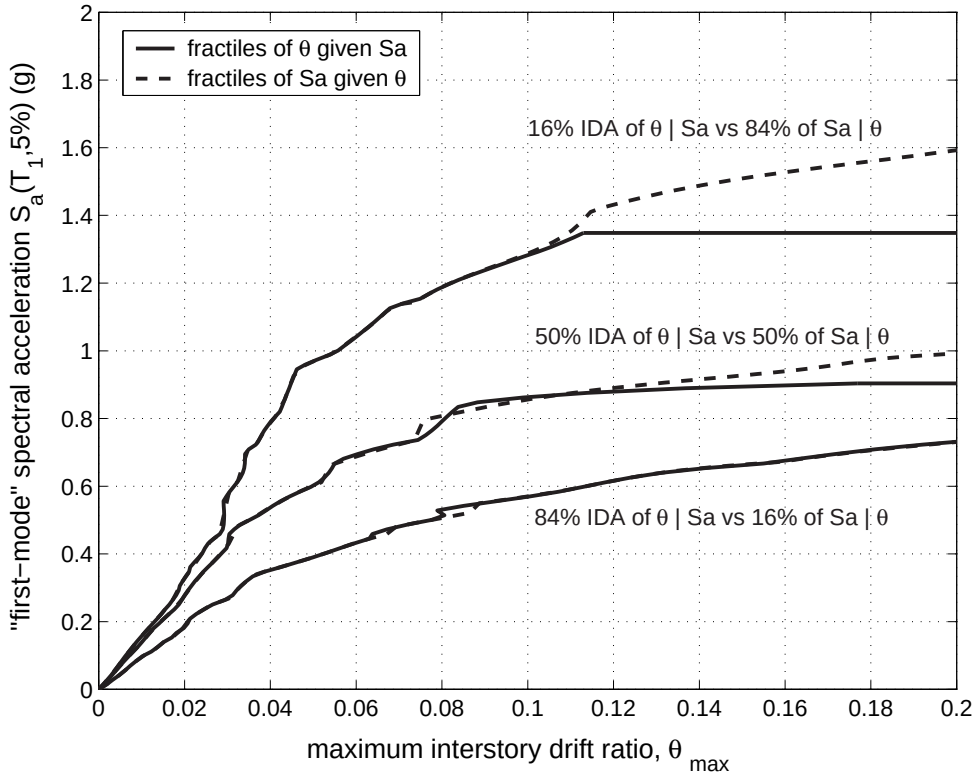

Figure 12. Summarization into fractiles of *IM* given *DM* versus fractiles of *DM* given *IM*.

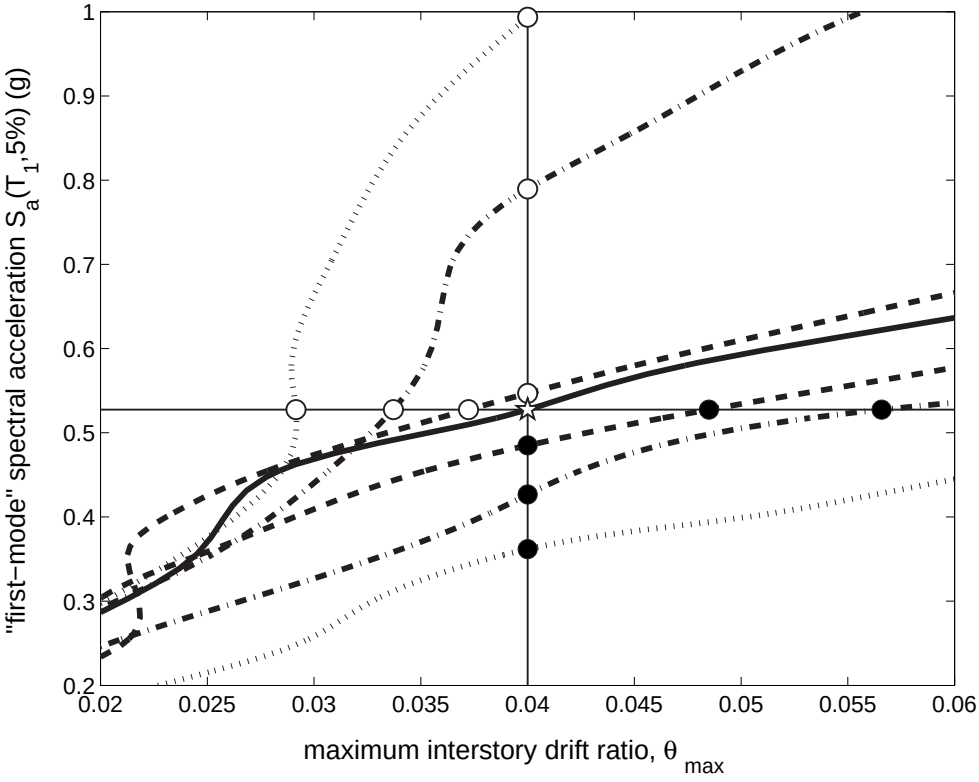

Figure 13. Two stripes, one given *DM* and one given *IM*, converging on the same median (the star). The records are No.18,5,19,6,9,10,13 from the highest to the lowest *IM* at $\theta_{\max} = 4\%$.

SUMMARIZATION GIVEN *IM* OR *DM*

When summarizing the IDA curves, we decided to use stripes of *DM* given levels of *IM*, instead of stripes of *IM* given *DM*. It often becomes an issue in the literature (e.g., Miranda 2001), whether one should summarize given *IM* or *DM*. The first approach can be thought of providing the distribution of demand *DM* that a given level of intensity *IM* can generate in the structure, while the latter is the distribution of intensities *IM* that are required to produce a given level of damage *DM*. Clearly, if we use the mean and standard deviation to summarize such stripes, the results will be very different between the two approaches (e.g., Miranda 2001). When fractiles are employed, though, this is not so; as shown in Figure 12, the 16%, 50%, and 84% fractiles given *IM* ($S_a(T_1, 5\%)$) almost perfectly match the 84%, 50%, and 16% fractiles respectively, given *DM* ($\theta_{\max}$).

The reasons behind this surprising fact become apparent in Figure 13. There, we have selected a subset of only seven records and have generated a (vertical) stripe of *IM*s given $DM \equiv \theta_{\max} = 4\%$. The median falls on the fourth, the middle of the seven curves, and is estimated to be $S_a(T_1, 5\%) = 0.53$ g (represented by a star). A (horizontal) stripe given *IM* is generated at this precise level and, remarkably, the median *DM* given $S_a(T_1, 5\%) = 0.53$ g is found to lie on the same IDA curve, right at the star, at $\theta_{\max} = 4\%$. To better illustrate this, we use white dots for IDA crossings on the left of the horizontal stripe and on the top of the vertical, but black dots at the bottom of the vertical or to the right of the horizontal. Local continuity and monotonicity assure that any IDA curve can only have two dots of the same color, i.e., each IDA curve will remain on the same side of the median curve.

Of course, it often happens that IDA curves are neither continuous, nor monotonic as due to hardening increased *IM*s may sometimes produce lower or the same *DM*-response (Figure 1). But even then, significant discrepancies (e.g., serious hardening in several curves at the same time) must occur to influence the robust fractiles, thus only slightly disturbing the matching of the fractiles given *DM* and given *IM*, and only in isolated places.

Why then are the 50% and 84% flatlines in Figure 12 not exactly matching? In the case of the seven curves in Figure 13, the median is conveniently falling right on the fourth of the seven curves. Since in Figure 12 a sample of 20 records is used, none of the three fractiles matches one of 20 curves. In that case, there are several ways to approximate the fractiles, and the one that we use involves linear interpolation between the closest two curves. For example, for 20 records, the median is calculated as the average of the 10th and the 11th record, as ordered on the relevant stripe. Obviously such interpolation generates different results given *IM* or *DM*. This problem becomes more apparent close to the flatline, where for summarization given *DM* we always have finite values of *IM* to interpolate, while for the summarization given *IM*, one of the closest two records produces infinite *DM* (which cannot be used for interpolation). If we use a larger sample, such discrepancies are reduced and eventually eliminated. Similarly, we could use another method to approximate the fractiles, e.g., select the lower of the two points that we use for the interpolation and similarly eliminate the problem. In any case, given the record-to-record variability, the fractiles are close enough and increasing the sample size they will actually converge to the same curves, no matter what method we use to estimate them.

SENSITIVITY TO THE RECORD SUITE SIZE

The IDA curves display significant record-to-record variability, as becomes obvious in Figure 3. It is only natural to investigate the accuracy of the results given the limited sample size of twenty records. Traditional analytical forms are difficult to implement for the fractiles or

the MAFs, hence we turn to the bootstrap method (Efron and Tibshirani 1993) to fill this gap. Application of the bootstrap involves sampling with replacement from the twenty records to generate an arbitrary number of alternate record suites and a corresponding number of summarized capacities or MAF estimates. From such samples of estimates, one can easily calculate the standard error or confidence intervals of the desired coverage (e.g., percentile bootstrap confidence intervals) for both fractile *IM*, *DM* capacities and MAFs.

The bootstrap estimate of the standard error, plus a 90% bootstrap confidence interval on the median *IM* and *DM* limit-state capacities appear on Table 6. It becomes obvious that using only 20 records provides a relatively accurate estimate of the capacity values for this structure; the median *IM* capacities show very small dispersion that predictably increases for limit-states closer to global dynamic instability. We should expect comparable, albeit higher, standard errors (and wider confidence intervals) for the 16% and 84% fractiles, as they are closer to the edges of the sample and thus relatively more variable. On the other hand, the fractile *DM* capacities have practically negligible standard error. In the case of IO and GI, this is a direct result of their definition, as they both lie at fixed values of $\theta_{\max}$ (2% and $+\infty$, respectively). Similarly, the median *DM* capacity for CP is almost always dominated by the $\theta_{\max} = 10\%$ rule, drastically reducing its dispersion. Again, this difference in the standard errors does not imply that using the *DM*-based form (Equation 10) instead of the *IM*-based (Equation 11), will result in higher confidence (less dispersion) in the MAFs estimate. The results should be identical even in this aspect when using any of the two approaches.

The influence of the number of records becomes more apparent if we realize that the standard error of the mean estimate (and approximately of the median as well) tends to fall of with a rate of $1/\sqrt{n}$ where n is the number of records (e.g., Benjamin and Cornell 1970). Hence, quadrupling the number of records to use a total of $n = 80$, results in only half the dispersion, while decreasing it by a factor of four, to use only $n = 5$, will (approximately) double the dispersion.

Table 6. Median *IM* and *DM* capacities for each limit-state, shown versus the bootstrapped standard error and the 90% confidence interval on the median estimate

	$S_a(T_1, 5\%)$ (g)			$\theta_{\max}$		
	$IM^c_{50\%}$	SE[1]	90% CI[2]	$DM^c_{50\%}$	SE[1]	90% CI[2]
IO	0.27	0.02	[0.24, 0.30]	0.02	-	-
CP	0.83	0.14	[0.72, 1.13]	0.10	0.004	[0.09, 0.10]
GI	0.91	0.17	[0.75, 1.20]	$+\infty$	-	-

[1] Standard Error

[2] Confidence Interval

How do the standard errors in the fractile capacities translate to the estimates of the MAFs? By applying the bootstrap to both the "exact" numerical (Equation 11) and the approximate analytic form (Equation 12) with either a local or a global fit to the hazard curve, we get the results shown in Table 7. As seen from the "exact" results, the limited sample of 20 records causes standard errors in the order of 50% in the estimates of the λ_{LS} for all limit-states. On the other hand, the approximation through Equation 12 considerably increases the standard error; in some cases it is in the order of 200% but sometimes the approximation totally fails and considerably overestimates the MAF. For the IO limit-state, it is the approximation with a global fit that may be destabilized, while at the CP and GI limit-state, it is the local fit that

Table 7. MAFs for each limit-state, calculated both numerically and with the approximate analytical form (global or local fit). The bootstrapped standard error and 90% confidence interval on the MAF estimate are also presented. Additionally, we test the hypothesis that the approximate λ_{LS} is equal to the exact at the 95% confidence level

limit-state	method	λ_{LS}	SE[1]	90% CI[2]	"equal" to exact?
IO	exact	0.019	0.011	[0.007, 0.04]	
	global	0.017	> 1000	[0.005, 0.33]	yes
	local	0.008	2.5	[0.004, 0.04]	yes
CP	exact	0.0004	0.0002	[0.0002, 0.0009]	
	global	0.0002	0.0008	[0.0001, 0.0004]	yes
	local	0.0005	> 1000	[0.0001, 0.7]	yes
GI	exact	0.0001	0.00007	[0.0001, 0.0003]	
	global	0.00003	0.00006	[0.00001, 0.00002]	yes
	local	0.0004	> 1000	[0.00003, 160]	yes

[1] Standard Error

[2] Confidence Interval

may become highly inaccurate. What happens is that individual bootstrap samples violate the assumptions needed to derive Equation 12; in some cases the *IM*-capacities are not nearly-lognormally distributed and in other cases either the global or the local fit fail to capture the shape of the hazard curve.

The bootstrap also offers us a way to investigate the accuracy of the approximate versus the "exact" calculation of the MAFs, given that we have only used 20 records. By bootstrapping the difference of the "exact" minus the approximate MAFs, a 95% confidence interval can be generated for each limit-state. If the interval contains zero, then, at the 95% confidence level, we cannot reject the hypothesis that the analytical and the numerical method produce the same results. As seen in Table 7, given the record-to-record variability and the limited sample size, the approximate results cannot be distinguished from the exact ones for any limit-state. In general, as long as we take care not to violate the stated assumptions, Equation 12 will provide good estimates.

CONCLUSIONS

The step-by-step practical application of incremental dynamic analysis has been demonstrated for a nine-story steel moment-resisting frame. By using publicly available software it has become almost trivial to perform the analysis, interpolate the IDA curves, estimate limit-state capacities and summarize the results into a format that can be easily integrated with modern PBEE frameworks. IDA offers a complete methodology to handle the abundant data from numerous analyses and extract useful conclusions. Still, the attention to detail is important: How many records, how many runs per record, how well interpolated, the use of approximations, are just some of the issues that can make a difference in the accuracy of the final IDA results. The methods that have been presented are designed to strike a favorable compromise between speed and accuracy and thus resolve such issues. Perhaps, the single most important thing to remember is the wealth of information that can be found in IDA if only we take advantage of ever-cheaper computing power and automated methods to investigate the structure's behavior.

ACKNOWLEDGEMENTS

Financial support for this research was provided by the sponsors of the Reliability of Marine Structures Affiliates Program of Stanford University.

REFERENCES CITED

Applied Technology Council (ATC), 1997. NEHRP guidelines for the seismic rehabilitation of buildings, prepared for the Building Seismic Safety Council, published by the Federal Emergency Management Agency, *Report No. FEMA-273*, Washington DC.

Baker, J. W., and Cornell, C. A., 2003. Uncertainty specification and propagation for loss estimation using FOSM methods, in *Proceedings of the 9th International Conference on Applications of Statistics and Probability in Civil Engineering (ICASP9)*, San Francisco, CA, 1–8.

Benjamin, J. R., and Cornell, C. A., 1970. *Probability, Statistics, and Decision for Civil Engineers*, McGraw-Hill, New York.

Cornell, C. A., Jalayer, F., Hamburger, R. O., and Foutch, D. A., 2002. The probabilistic basis for the 2000 SAC/FEMA steel moment frame guidelines, *ASCE Journal of Structural Engineering* **128** (4), 526–533.

Cornell, C. A., and Krawinkler, H., 2000. Progress and challenges in seismic performance assessment, *PEER Center News* **3** (2), (accessed: February 12th, 2004). http://peer.berkeley.edu/news/2000spring/index.html

Efron, B., and Tibshirani, R. J., 1993. *An Introduction to the Bootstrap*, Chapman & Hall/CRC, New York.

Farin, G., 1990. *Curves and Surfaces for Computer Aided Geometric Design: A Practical Guide*, 2nd ed., Academic Press, San Diego, CA.

Jalayer, F., and Cornell, C. A., 2002. A technical framework for probability-based demand and capacity factor (DCFD) seismic formats, *Report No. RMS-43*, RMS Program, Stanford University, Stanford, CA.

Lee, E., 1989. Choosing nodes in parametric curve interpolation, *Computer Aided Design* **21** (6), 363–370.

Lee, K., and Foutch, D. A., 2002. Performance evaluation of new steel frame buildings for seismic loads, *Earthquake Engineering and Structural Dynamics* **31** (3), 653–670.

Luco, N., 2002. Probabilistic seismic demand analysis, SMRF connection fractures, and near-source effects, PhD Dissertation, Department of Civil and Environmental Engineering, Stanford University, Stanford, CA, (accessed: February 12th, 2004). http://www.stanford.edu/group/rms/Thesis/Luco_Dissertation.zip

Macrae, G. A., and Kawashima, K., 1997. Post-earthquake residual displacements of bilinear oscillators, *Earthquake Engineering and Structural Dynamics* **26** (7), 701–716.

Miranda, E., 2001. Estimation of inelastic deformation demands of SDOF systems, *ASCE Journal of Structural Engineering* **127** (9), 1005–1012.

Prakhash, V., Powell, G. H., and Filippou, F. C., 1992. DRAIN-2DX: Base program user guide, *Report No. UCB/SEMM-92/29*, Department of Civil Engineering, University of California, Berkeley, CA.

SAC Joint Venture, 2000a. Recommended seismic design criteria for new steel moment-frame buildings, *Report No. FEMA-350*, prepared for the Federal Emergency Management Agency, Washington DC.

SAC Joint Venture, 2000b. Recommended seismic evaluation and upgrade criteria for existing welded steel moment-frame buildings, *Report No. FEMA-351*, prepared for the Federal Emergency Management Agency, Washington DC.

Shi, S., and Foutch, D. A., 1997. Connection Element (type 10) for DRAIN-2DX, *Report*, Department of Civil and Environmental Engineering, University of Illinois at Urbana-Champaign, Urbana, IL.

Shome, N., and Cornell, C. A., 1999. Probabilistic seismic demand analysis of nonlinear structures, *Report No. RMS-35*, RMS Program, Stanford University, Stanford, CA, (accessed: Feb 12th, 2004). http://www.stanford.edu/group/rms/Thesis/NileshShome.pdf

Vamvatsikos, D., and Cornell, C. A., 2002a. Incremental dynamic analysis, *Earthquake Engineering and Structural Dynamics* **31** (3), 491–514.

Vamvatsikos, D., and Cornell, C. A., 2002b. Tracing and post-processing of IDA curves: Theory and software implementation, *Report No. RMS-44*, RMS Program, Stanford University, Stanford, CA.

Vamvatsikos, D., and Cornell, C. A., 2004a. Direct estimation of the seismic demand and capacity of MDOF systems through incremental dynamic analysis of an SDOF approximation, *ASCE Journal of Structural Engineering*, in review.

Vamvatsikos, D., and Cornell, C. A., 2004b. Direct estimation of the seismic demand and capacity of oscillators with multi-linear static pushovers through incremental dynamic analysis, *Earthquake Engineering and Structural Dynamics*, in review.

Veletsos, A. S., and Newmark, N. M., 1960. Effect of inelastic behavior on the response of simple systems to earthquake motions, in *Proceedings of the 2nd World Conference on Earthquake Engineering*, Tokyo, Japan, 895–912.

Yun, S. Y., Hamburger, R. O., Cornell, C. A., and Foutch, D. A., 2002. Seismic performance evaluation for steel moment frames, *ASCE Journal of Structural Engineering* **128** (4), 534–545.