Properties of Energy Storage for High Grid Penetration

This report builds upon the simulations of Solomon et al. in determining what properties would be required of an energy storage system to buffer the interaction between an electric grid and a large-scale photovoltaic power plant. In the previous paper, Solomon simulated the output of a large size range of PV plants and compared this with the power demand data from the Israel Electric Corporation. Based on the assumption that little-to-no PV energy will be dumped, he could thereby calculate constraints upon a storage system to guarantee this. These models were computed with hourly data for the year 2006, based on a hypothetical plant of non-tracking panels in Sde Boker, with their tilt-angle fixed at the latitude of the site.

Now, with the availability of more detailed data[1] from the IEC, and with the benefit of the broad examination in this previous paper, we will study these constraints on a more precise (one-minute) time scale, confirm their validity, generalize to panels with tracking systems, and focus in on the most relevant aspects of the analysis.

We should, however, adjust certain parameters of the model to suit this study. For instance, Solomon et al. plotted many of the properties of the system against the size of the plant, where this size was measured as a multiple of the “no-dump” size. (This is the maximum peak watt rating the plant could possess, without ever, at any point of the year, supplying more power than the grid could utilize). However, ND-size is not invariant under change of scale (it tends to decrease slightly as one takes smaller time bins; the larger/hourly time bins imply an averaging which smooths the output and demand curves allowing a greater sized plant). For sections of this report which involve comparing multiple time scales, the standard will be to measure system sizes in absolute quantities (ie peak watt rating of the plant), so that the movement of ND-size does not obscure the effects of scale. (See appendix for numerical details). This said, we shall begin to examine the properties required of storage for the Israeli grid.

Energy Capacity

The simplest, and most demanding, set of constraints to apply to our system is that (1) the grid always supplies whatever power is demanded of it, by drawing from solar plants, energy storage, and conventional power generation, *in that order*. So whenever the PV plant falls short, the storage and/or the conventional grid fulfill the remaining demand. (2) No PV energy is lost, other than that due to the inherent inefficiency of storage[2]. The latter constraint, another realization of “no-dump,” will be relaxed later on, but, for now, it is vital to the definition of “energy capacity.”

---

[1] The corporation provided minutely power load data from the year 2009.

[2] To keep the model simple and general, this loss will be approximated, as in Solomon et al, by a round-trip efficiency of 75%. And to be as conservative as possible in implementation, the 75% loss has been inserted entirely into the discharge of storage. (This way, we cannot understate the amount which must be stored.)

The energy capacity, of a given plant in a given year, shall be defined as the least amount of storage necessary to accommodate all the excess energy that the plant produces (to feed it into the grid when needed). Each graph below shows the energy capacity across a range of absolute plant sizes, with curves for four different flexibility factors[3]. We have graphed four data sets so we can see that this measure is relatively stable across time scales, with some yearly variation.

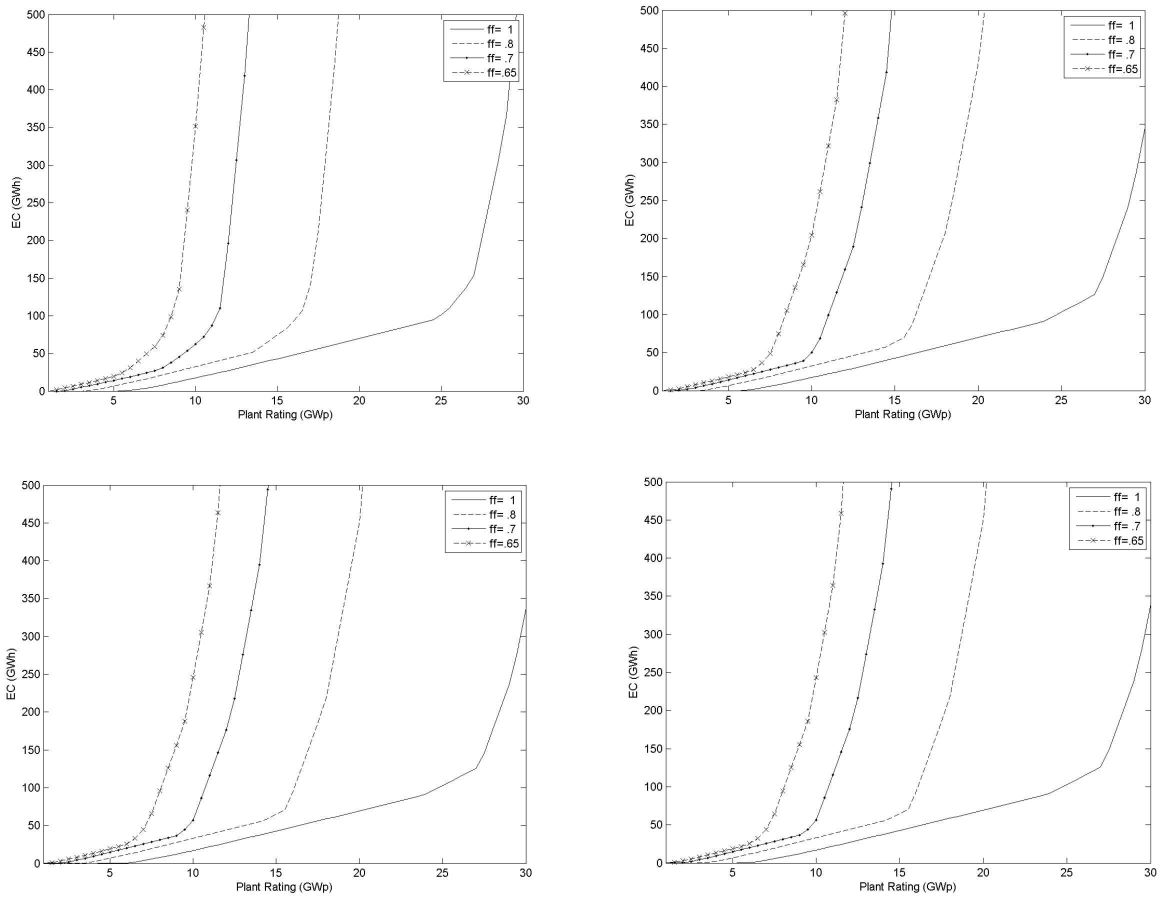

Clockwise from top left: 2006 Hourly, 2009 Hourly, 2009 10-minutely, 2009 minutely

The energy capacity graphs appear to be approximately linear as they begin, but then quickly explode. The rapid change-in-slope reflects when the plant becomes so large that the excess energy stored during the days in pre-summer is not fully used in the short term and just continues to compound[4]. For our purposes, we can see from these graphs that the important region is marked by an EC of *around* 100 GWh or less, because, beyond that point, raising the energy capacity of storage does not seem to contribute much. This claim will be made far more concrete when we discuss penetration rates. But for now, let us simply note that EC, as a function of absolute system size, is a consistent measure across time bins.

---

[3] Flexibility factor (ff), as defined by Denholm and Margolis, is a characteristic of the grid which reflects how much of its load can be solarized. The grid typically operates at ff=.65, however, the data indicates the ff could technically be increased to .8. ff=1 represents a load which could, in theory, be fully provided by solar.

[4] This trend is explained in a great deal more detail by Solomon.

Power Capacity

One may also describe the storage system by how quickly it must be required to charge and discharge. These are two *a priori* unrelated requirements, so we will refer to the charging capacity and discharging capacity as the minimum rates of the respective energy flows which the storage must be able to sustain to avoid dumping any energy.

Let us first examine charging capacity:

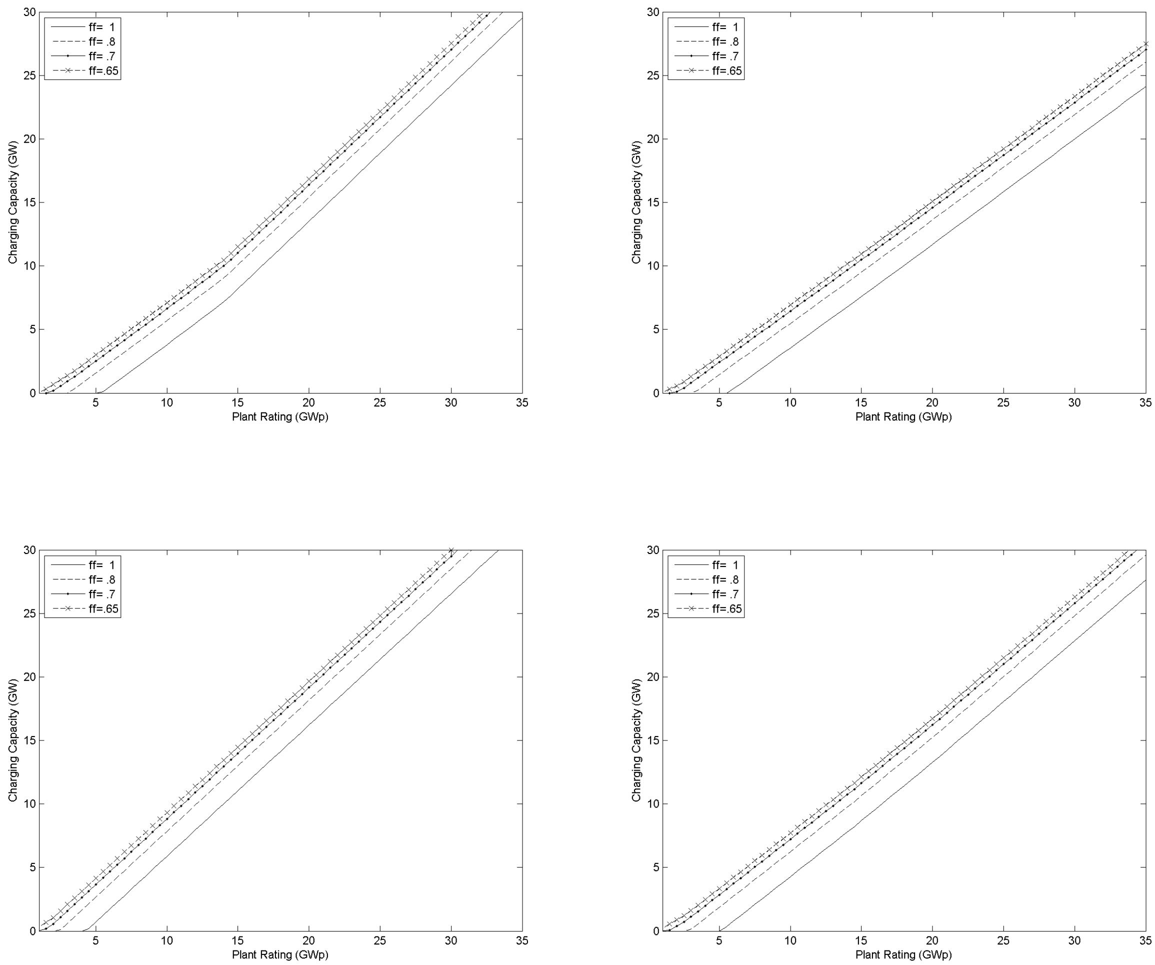

Clockwise from top left: 2006 Hourly, 2009 Hourly, 2009 Ten Minute, 2009 Minutely.

Examining the 2009 charts, we notice that the slope of the charging capacity curves increases at finer scales. Again, larger scales imply an averaging of conditions, which tends to smooth the demand and output curves. So the moments of maximum charging requirement fade into averages, which is why larger scales underrepresent the charging requirement. However, the magnitude of this effect seems to be of the same order as the change from 2006 to 2009.

Notice that, in choosing to represent these data by absolute system size, all of the slopes in a given plot agree. That is, once a PV system is large enough that it must store some energy, the charging capacity grows linearly in absolute size. However, the lines do have different zeros for different flexibility factors, because, in a more flexible grid, a plant can be relatively larger before needing to store any energy. If we were to plot the same data against no-dump multiples, the zeros corresponding to different flexibility factors would agree, but the lines would have different slopes.

Now, to consider *discharging capacity* produces a more interesting plot:

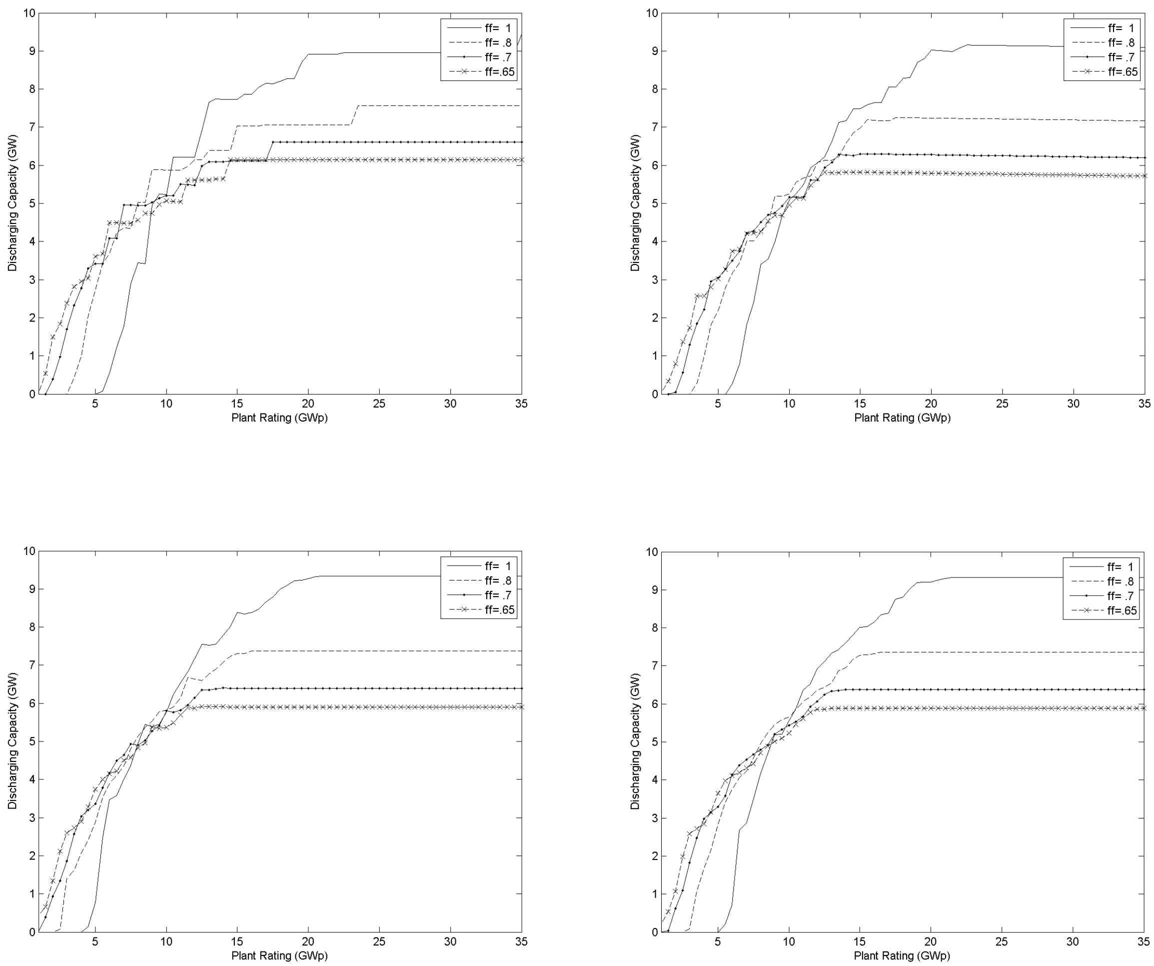

Clockwise from top left: 2006 Hourly, 2009 Hourly, 2009 10-minutely, 2009 minutely

Discharging capacity curves demonstrate three regions: (1) a rapid rise, and (2) a middle region of variable jagged growth, followed by (3) a flat mesa out to infinity. The reason for the initial rise is clear: as the size passes the no-dump threshold, it suddenly becomes necessary for storage to discharge energy, and DC must quickly shoot from zero to the capacity required to satisfy a typical day. The cause of the flat-line is also evident: storage will never need to discharge more than the maximum amount of power demanded by the grid over a given year, regardless of the size of the plant or of the storage system[5]. The reason for the often hideous

[5] Close inspection may occasionally reveal an almost unnoticeable downward trend in the flat-line region. This arises when the time-bin which happens to require the maximum amount of discharging also happens to include a small amount of PV generation. In that

segment in-between will become clear as we study an example. Let us take the cleanest serration: the rightmost in the ff=.8 plot of the 2006 Hourly graph.

This DC discontinuity occurs between plant sizes ~20GWp and ~25GWp, where the DC requirement jumps from about 7GW to about 7.5GW. Let us examine both sides of the discontinuity. According to our simulation, for a plant sized about 20GWp, the moment of maximum discharge occurs between hours 8633-8634 of the year (this is around 6PM on December 25th). For a plant sized about 25GWp, the moment of maximum discharge occurs between hours 8657-8658 (this is around 6PM on December 26th).

Let us examine a plot of the energy which is saved up in storage as we pass through the relevant time intervals, with different lines for different sized plants:

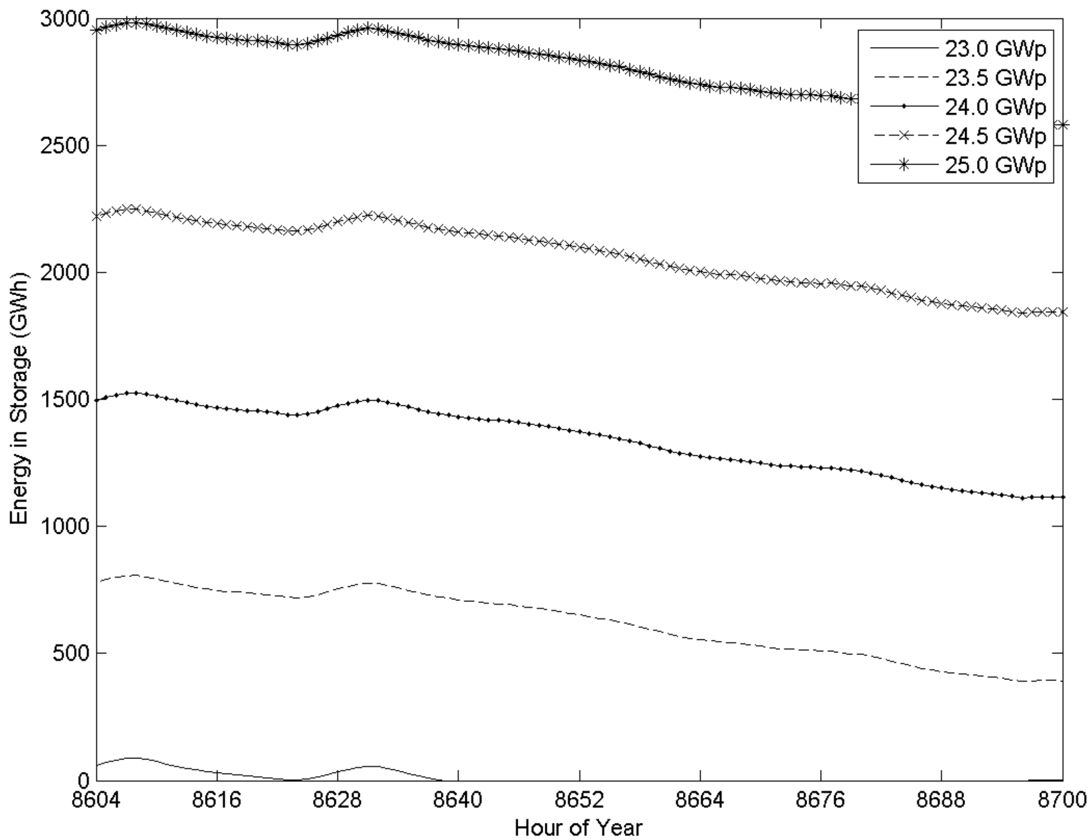

The ticks have been applied every twelve hours such that hour 8616 is the start of December 25th. Discharge, being the difference in storage between two hours, corresponds to negative slope. The first feature one should notice is that the down-sloping lines are approximately parallel when non-zero, because the amount which each storage discharges hourly is the production of the plant minus the needs of the grid. The production of the plant depends

case, a larger plant will be able to slightly increase this output, and thereby slightly decrease the contribution which storage must make. However, the region of PV size at which this effect is worth considering is far beyond any reasonable plant size.

on its size, but the range of sizes here is not too vast, and the need of the grid, of course, does not depend on size. Since, during discharge, the need of the grid is often much more than production, the downward slope at a given hour is about the same on each of these curves whenever storage is not empty.

Plants sized 23GWp and slightly smaller are all captured by the same solid curve labeled "23GWp." For these storage systems, the moment of maximum discharge is during that last downward-sloping region between hour 8628 and hour 8640. By the start of December 25th (hour 8640), these storage systems are out of energy. Meanwhile, the larger plants discharge approximately that same amount during that period, but they also have energy to discharge in the evening of December 26th, which happens to include hours 8657-8658, when the hourly need of the grid most exceeds PV production.

Now the reason for the discontinuities in discharge capacity requirement (DCR) plots should be evident. There are certain moments in the year which require large discharges, and smaller systems may not have enough energy in storage to fill those particular points, so a lesser point must be their maximum. However, a large enough system will have the energy available to supply that point, and will take that point to be its maximum discharge. This creates the sharp jumps in DCR with respect to system size. The greatest such point, of course, determines the upper bound of the DCR curve. We note again however, that the approximation (equal slopes) in this discussion relies on the output of the plant being negligible compared to the need of the grid at that moment (which is, almost by definition, the case for these maximum DCR points). When there is some power production at that moment, we do see that minute downward trend in DCR requirement with system size (as plant production increases with size)[6].

Having explained the origin of the DCR serrations in the random features of a year, it is clear that that they are highly variable (even just by looking at the graphs already shown), and the important feature is the general shape of the graph. The upper bound of DCR is a useful feature in determining what type of energy storage may be suitable, but we also notice that this bound changes with scale. Because the hourly time-binning smooths the output and the needs of the grid, it reduces the spread of discharges necessary, thereby typically lowering the maximum discharge. The difference in discharging requirement between hourly and minutely, though, is only on the order of a few percent, much like with charging capacity.

We have now confirmed that the results of Solomon et al, computed with 2006 hourly data, hold well under the new 2009 data at the full range of scales. In a moment we shall

[6] The astute reader notices that the amount of storage energy in this graph is preposterous. And correctly, examining the early plot of energy capacity reveals that plants of this size fall far into the exponential region of energy capacity, and this much would never be used. However, that is not relevant to our discussion here. This serration in the unnecessarily high plant-size domain is chosen simply because coincidental factors make it easier to plot and explain. First, that it is surrounded by flat regions (wherein the timing of the moment of maximum discharge is constant). Second, the left and right moments are extremely close to one another, such that the plot does not require a long time domain. These features make the graph straightforward and the explanation intelligible; however, they are not, in general, true, and not, in principle, necessary for any part of this discussion.

reiterate this analysis, but at a quicker pace, to show that these results further generalize to different types of photovoltaic plants. To this point, we have been generating these graphs from tilt-latitude PV panel simulations, but now let us create the same graphs for a concentrator photovoltaic (CPV) system, briefly look at them, and see that there are no qualitative changes to the previous discussion. Graphs have been arranged in the same clockwise format.

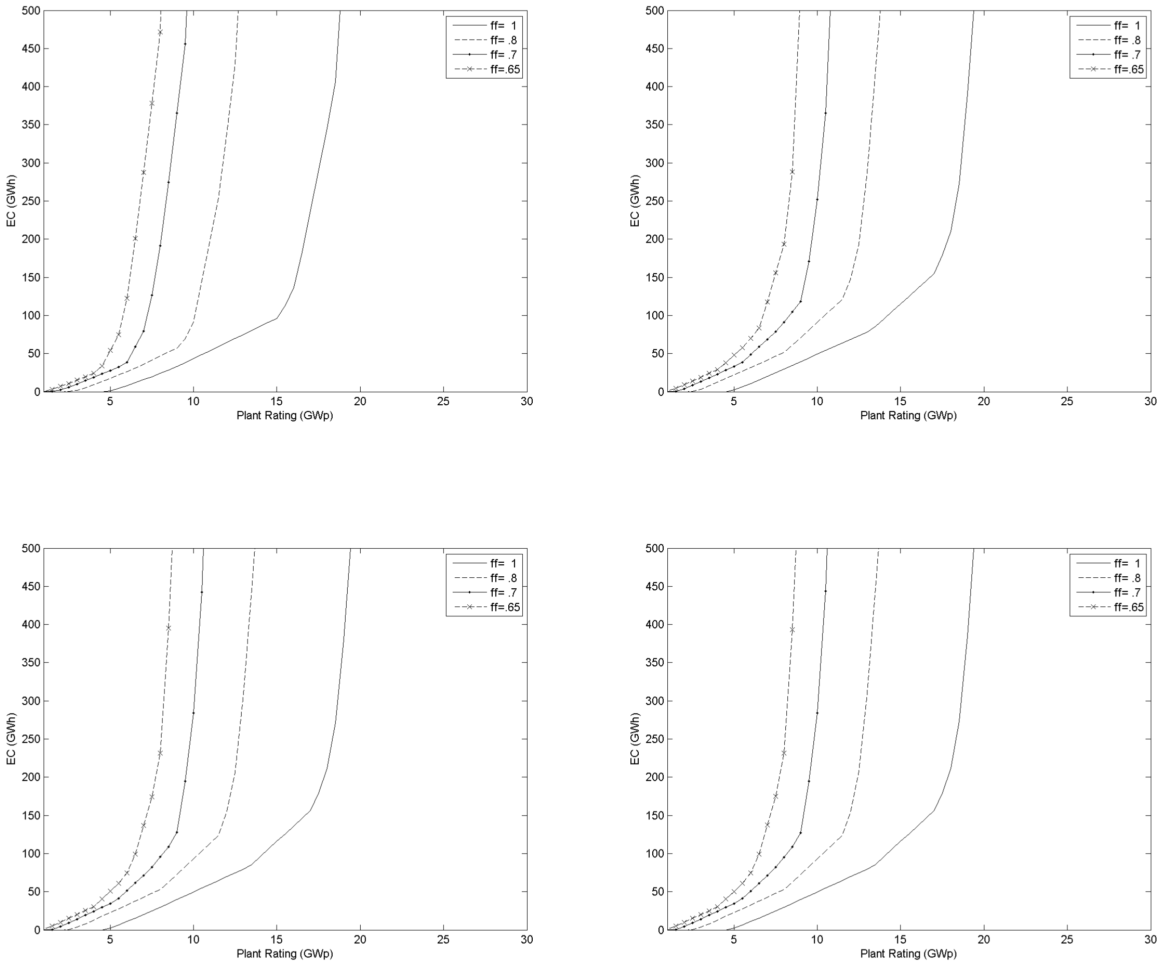

Energy Capcity

These graphs are clearly of the same form, but compressed horizontally. There is nothing to this. A CPV system, with the same rating as a tilt-latitude system, will produce more energy throughout the year because of its sun-tracking capability.

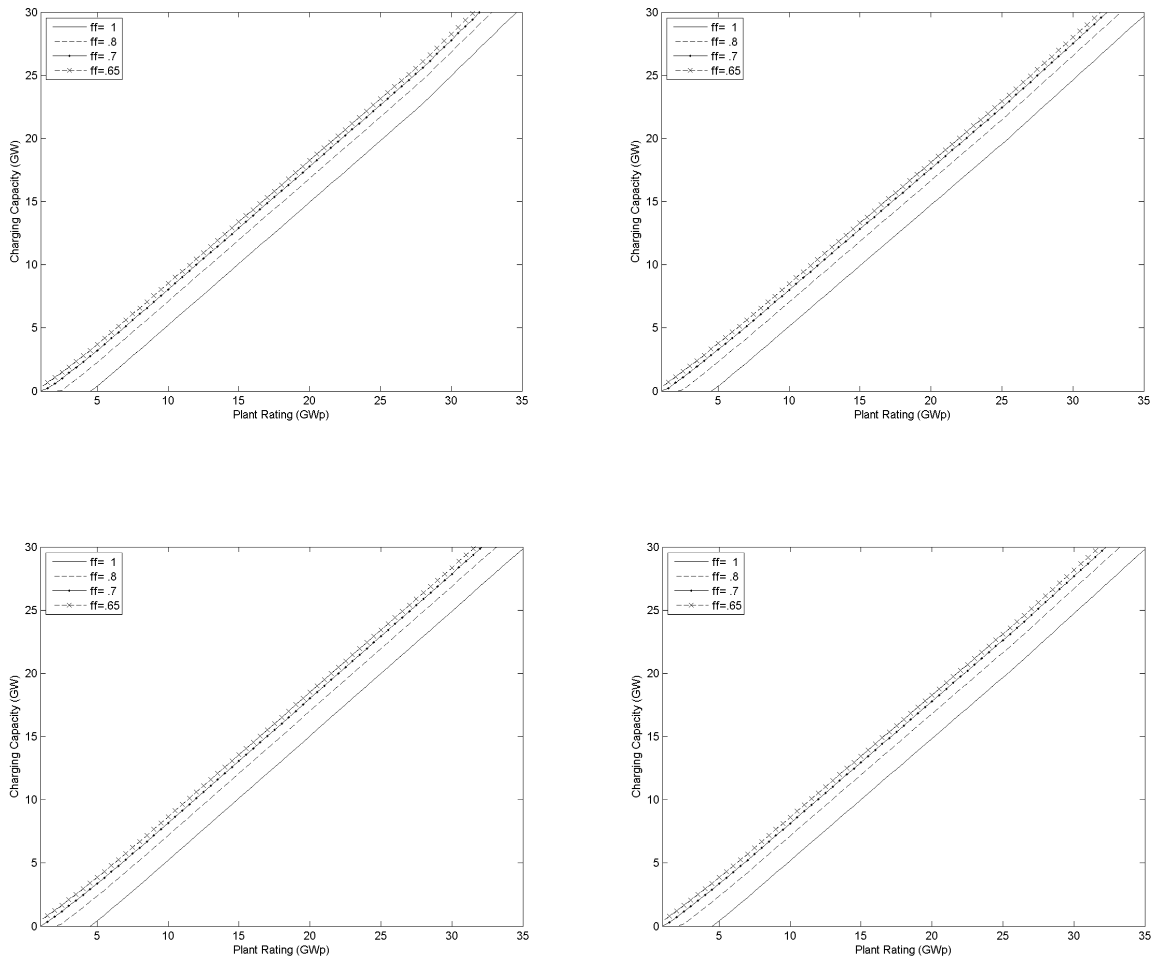

Charging Capacity

Interestingly, charging capacity seems to be even more stable for CPV systems than for tilt-latitude panels, at least looking at this particular year.

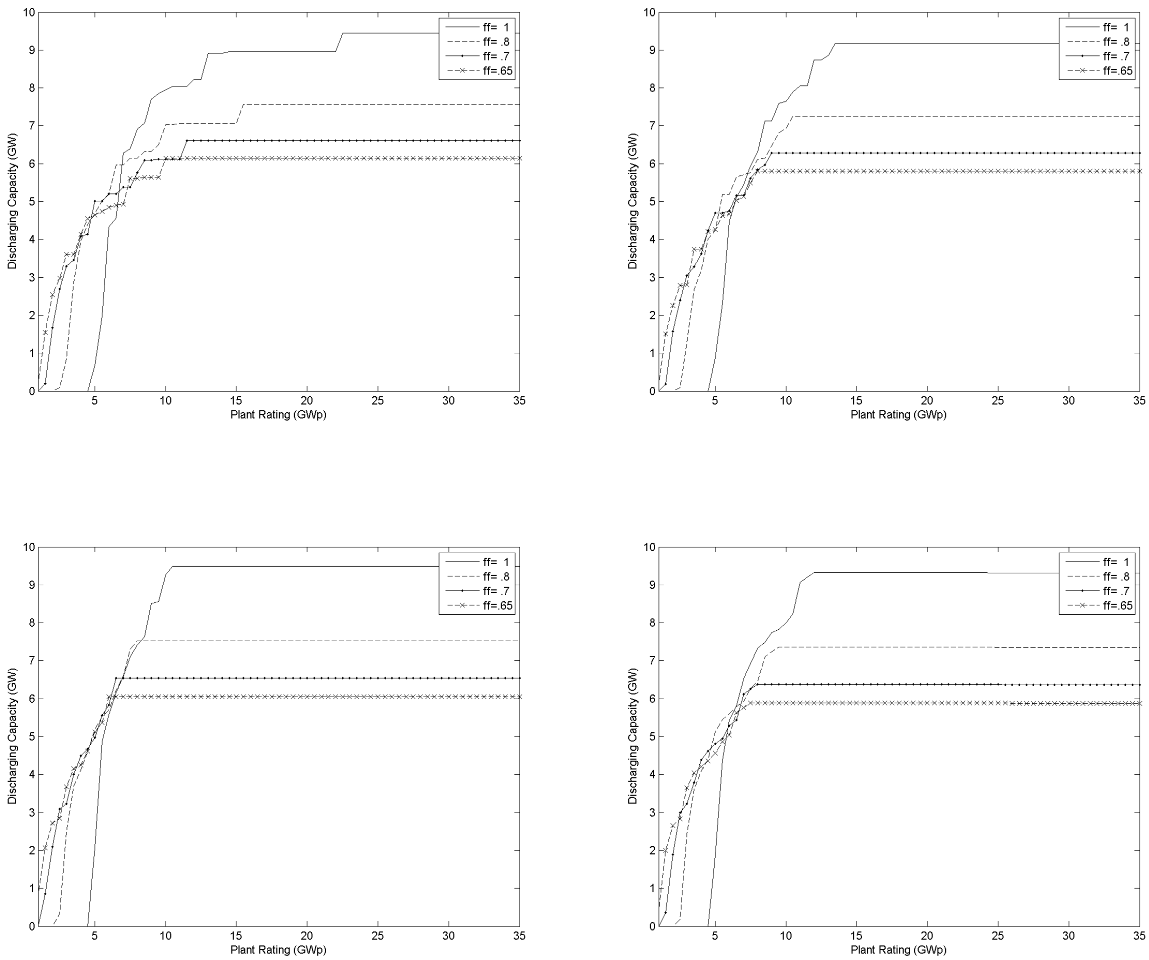

Discharging Capacity

Again, we can see evidence of this horizontal compression in that discharging capacity hits its peak a bit earlier with CPV than with tilt-latitude, but, other than that, there are no obvious differences. These sorts of changes would be mitigated by use of the ND-multiples representation, which would effectively equate to rescaling of the horizontal axis to account for the difference in effectiveness of CPV and tilt-latitude systems. However, we have no need to enter into such matters at the moment. These graphs suffice to show that the same relations hold.

Now we have validated the results of Solomon et al across different scales and different systems. We have seen that the changes in energy capacity with scale are negligible and the changes in power capacity not terribly drastic. For the rest of the paper, we will be focusing more on energy capacity, so, from here in, we will not see it necessary to include graphs for each year and scale, when a single hourly plot will suffice. The rest of the plots shown will be from CPV systems, using 2009 *hourly* data. (It is easier to compute with a data set at one-sixtieth the length, and, evidently, this will not noticeably change our coming conclusions). Given this freedom, let us apply these findings to generate a few practical graphs.

Penetration and Dump Percentages

Up until now, we have been ranging over PV system sizes, and recording what is required of storage. Now, let us put some restrictions on storage, and see how well the system performs. In other words, we are relaxing constraint #2, as promised. The following plots describe how much of the annual demand can be supplied by solar, with certain fixed sizes of energy storage systems, and how large a PV plant would be necessary to achieve this.

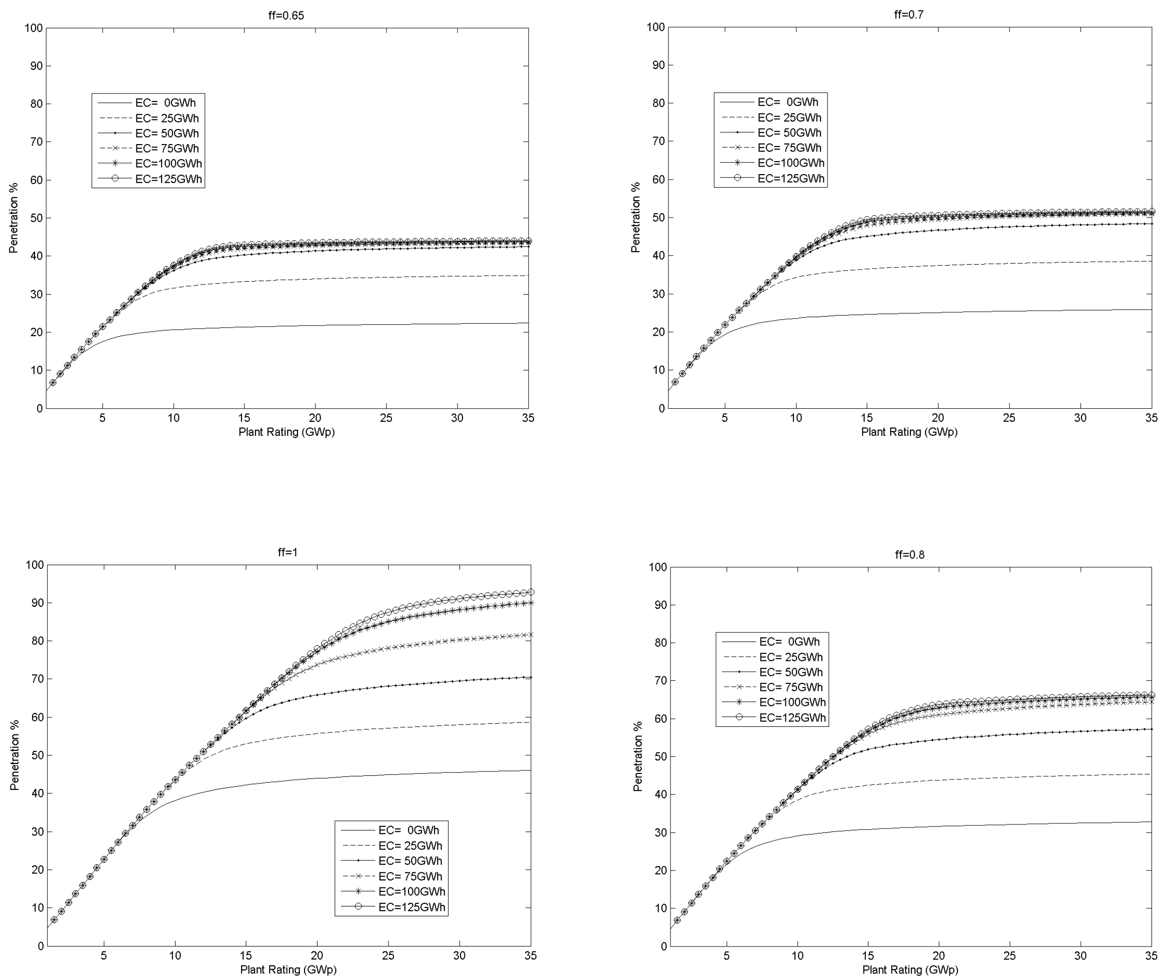

Again, the natural order for examining these graphs is clockwise.[7]

When we view fix the storage size and vary the plant size over such a wide range, it is inevitable that some energy will be dumped at times when storage cannot accommodate it. How much? Let us consider the ff=1 plot, for instance. The highest EC allowed is 125GWh. Referencing the first set of graphs in this paper, that is inadequate for PV systems past ~17GWp. Checking the ff=.65 plot, the highest relevant EC is 50GWh (after which, more storage is effectively useless). That is inadequate beyond a plant of size ~5GWp. Let us examine the percentage of energy which is simply dumped, due to insufficient storage, in order to generate these penetrations.

[7] Notice that, here, the different graphs are for different ff's, not different data scales

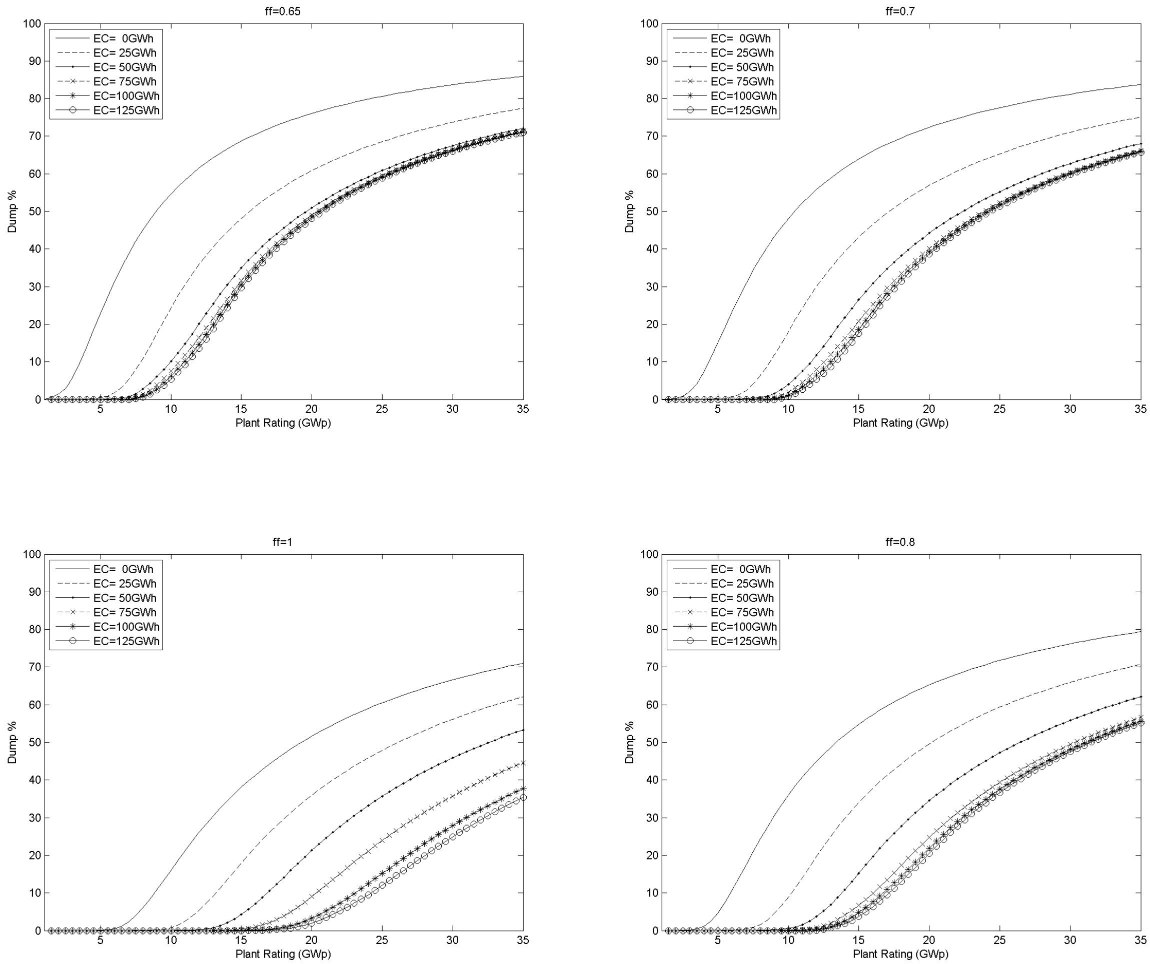

Not coincidentally, the point where the penetration curves level off is roughly where the dumpage curves reach their steepest. This qualitative observation suggests plotting penetration against dumpage. And if we do so, we arrive at the following.

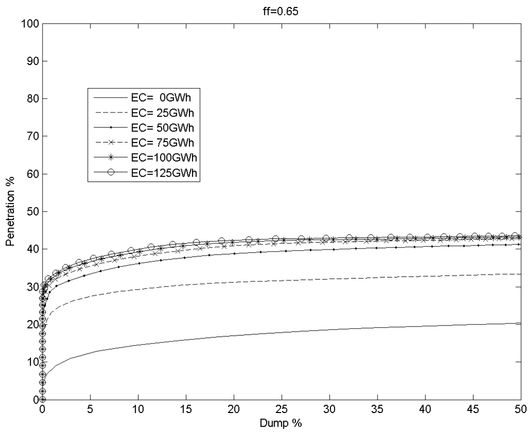

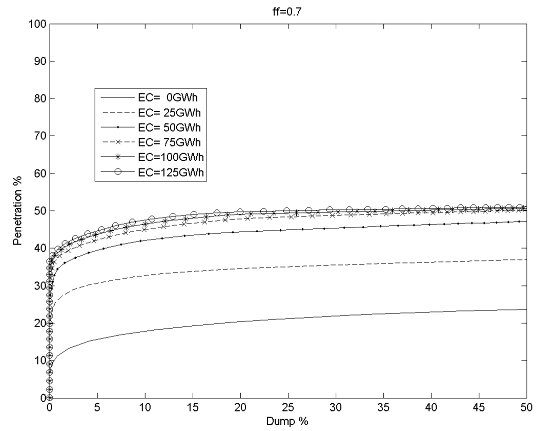

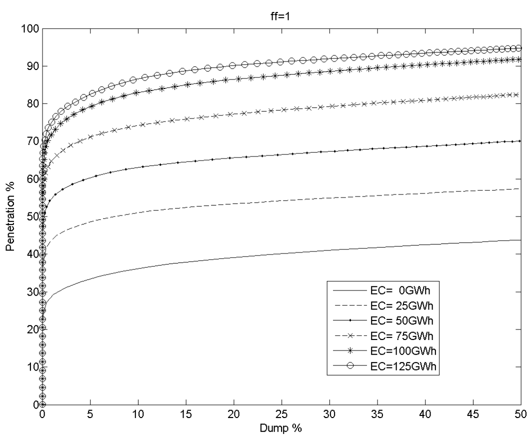

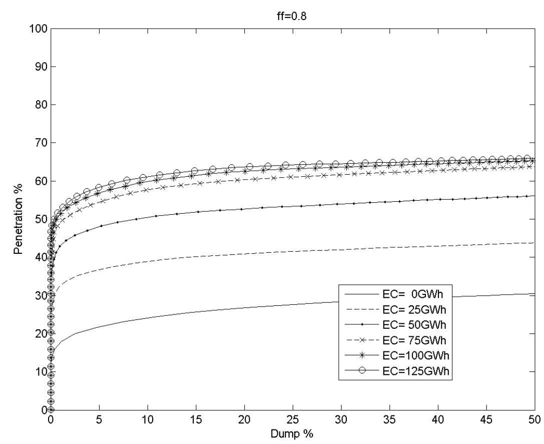

From these, we can readily observe that, for any energy storage size, a PV plant size which yields about 5% dumpage should suffice for maximizing grid penetration if that is the goal. This is not a precise, economic optimization, but simply the technical fact that each storage size yields a certain range of penetration possibilities, and almost all of that range is available with such minimal dumpage.

However, this is only the dumpage due to insufficient storage, not including the loss of inefficient storage. These plots understate the actual amount of PV energy lost, and advantage the high EC curves which make more use of storage. Taking storage inefficiency into account yields the following more comparable plots.

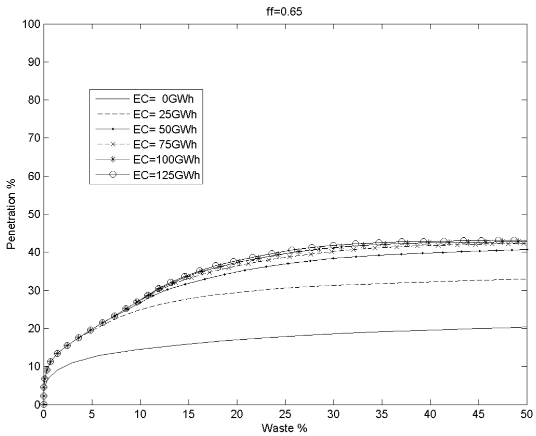

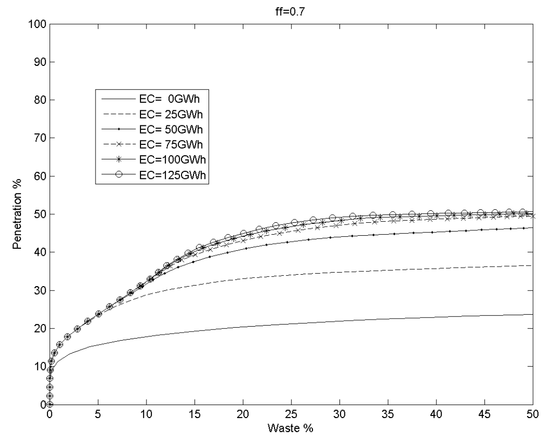

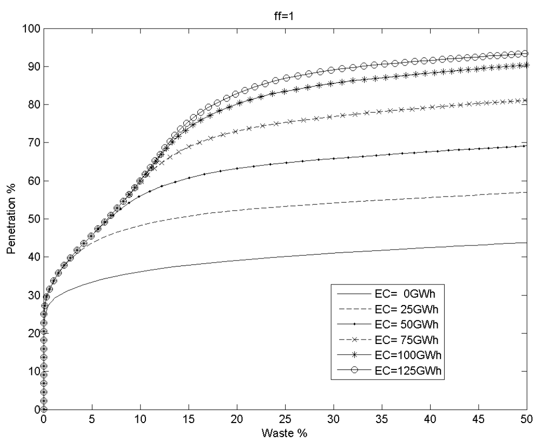

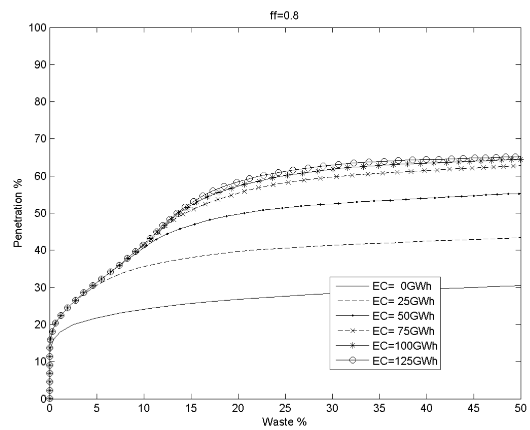

Now if we just fix a percentage of energy we are willing to lose to storage inefficiencies and insufficiencies, we can directly plot penetration against storage size (the first three plots) or against plant size (the latter three).

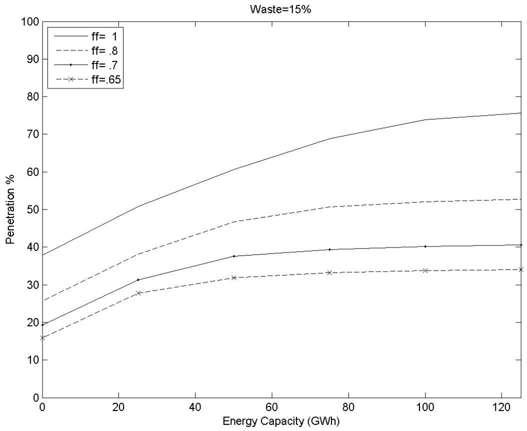

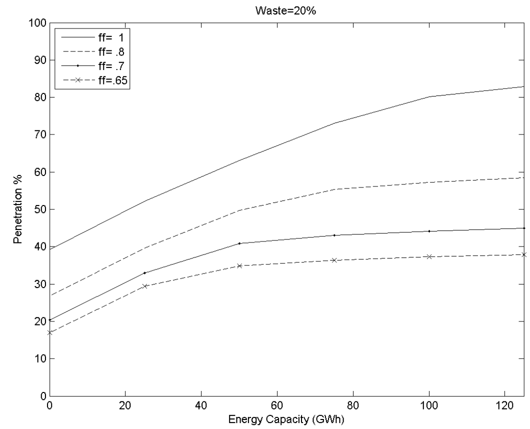

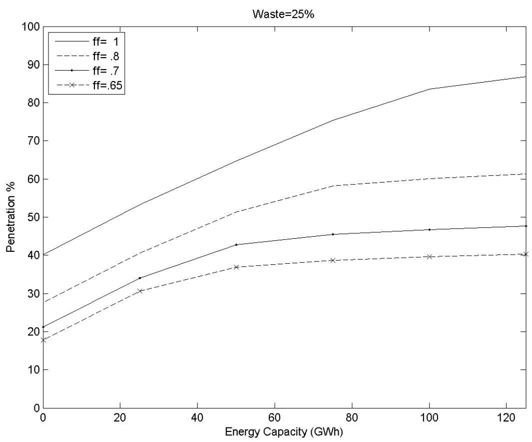

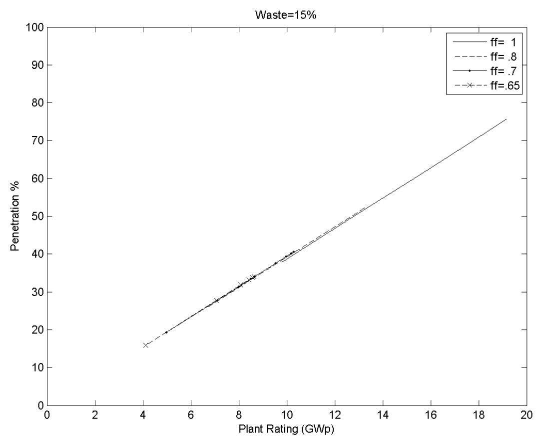

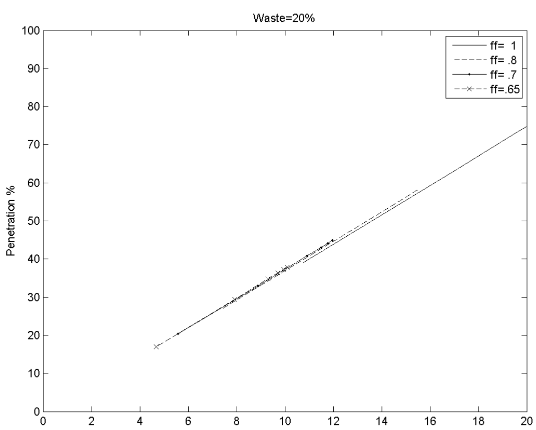

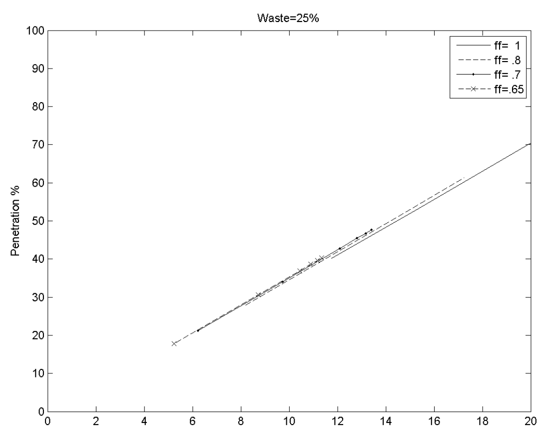

For a pratical configuration in which some dumping is allowed, these graphs, the most important in the paper, demonstrate how much grid penetration is acheivable at any size of storage or plant. Assume, as an example, that the Israeli grid will perform at ff=.8. Then, with a total waste of 15%, a PV system with a rating of 13GWp combined with a storage of 60GWh could expect to attain about 50% penetration.

For completeness, let us also compute the power capacity requirements corresponding to the above graphs. The first three below are charging capacity, and the latter three are discharging capacity. The scales have been choosen equal for the two.

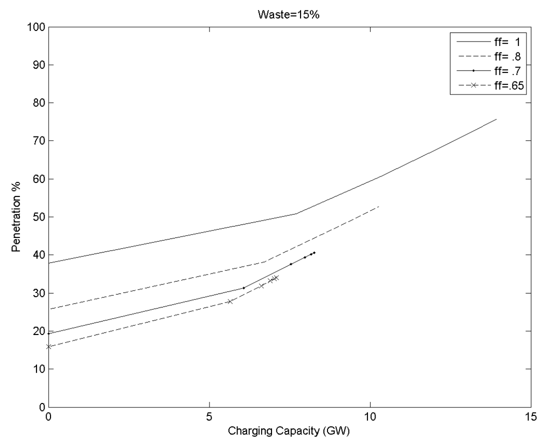

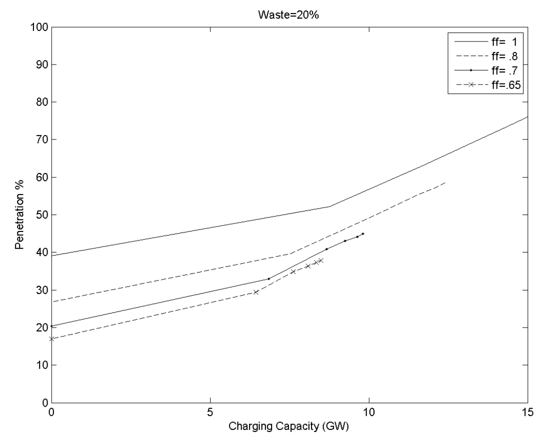

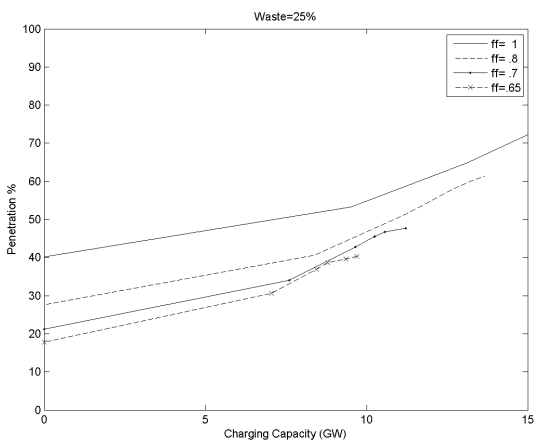

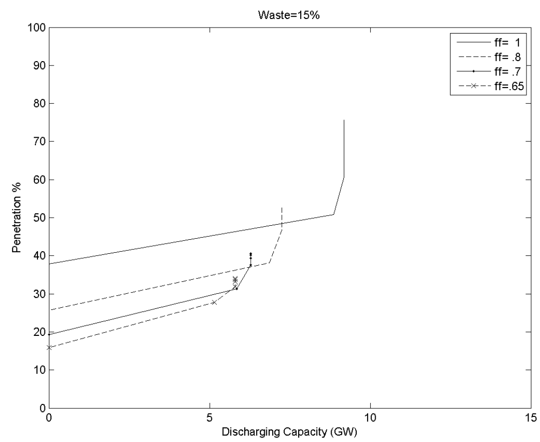

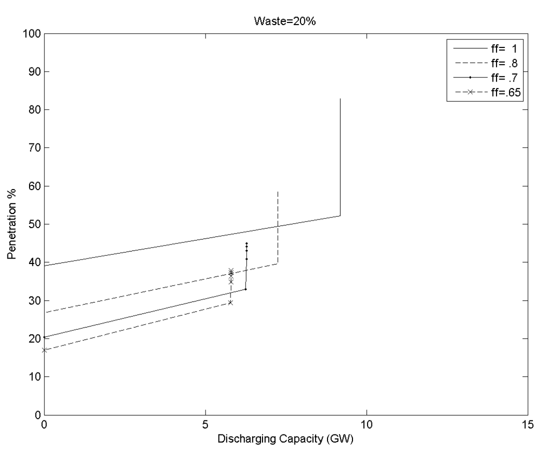

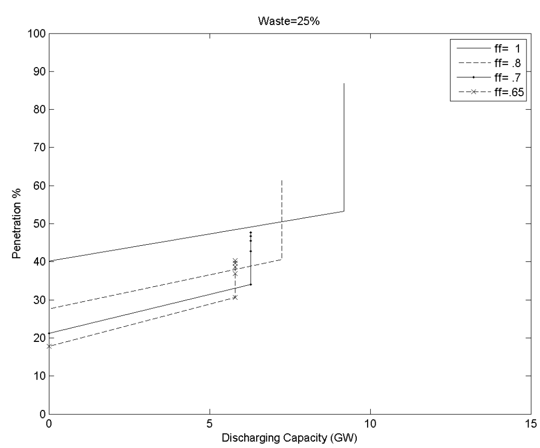

Note that the only points which require no charging or discharging are the points where the energy storage system is 0GWh. Beyond that, there is an almost *immediate* jump in both charging capacity and discharging capacity to bring these graphs toward matching the earlier plots of these variables against plant size. For discharging capacity, that means flat-lining around 9GWp. For charging capacity, (since plant size is linear in penetration for this range, and charging capacity is linear in plant size, but offset from the origin) this is an up-sloping line. So, to finish our sample interpretation of the previous set of graphs, the system discussed would require a power capacity of about 9GW.

Conclusions

The original goal of this paper was to demonstrate that the results of Solomon et al. hold across different time scales. Scaling down, energy capacity remains effectively constant, while power capacities increase by several percent. Thus, when planning for a solar system, working with hourly data for the site and for the grid demand should be acceptable, if enough margin is allotted in power requirements. We have also, re-demonstrated the results for concentrator photovoltaic systems.

Furthermore, we have examined the data in absolute numbers rather than ND-multiples; should the reader wish to make any conversions or generalization, a table of ND sizes has been provided below in the appendix.

And finally, we have shown how, with different tolerances to wasted energy, plant size or storage size can determine the amount of penetration possible for the grid at a wide range of flexibility. The methods employed here, and the shape of the plots, if not the raw numbers, should generalize to other grids.

Appendix: ND Sizes* in GWp (for CPV systems)

	2006 Hourly	2009 Hourly	2009 10-minute	2009 Minutely
ff=1	4.59	4.56	4.55	4.54
ff=.8	2.43	2.37	2.3	2.28
ff=.7	1.23	1.22	1.13	1.06
ff=.65	0.482	0.613	0.402	0.388

*One of the reasons for the use of absolute sizes in this paper was the often significant difference between 2009 Hourly and 2009 Minutely ND sizes, which would have shrunk the horizontal scale of plots for different ff values, within a given graph, as the time scale decreased.