ECONOMICS OF PLANNING
Vol. 8, No. 1—2, 1968
Printed in Norway

Digital simulation techniques in power system planning

Vidkunn Hveding
Norwegian Water Resources and Electricity Board

I. PROBLEMS IN POWER SYSTEM PLANNING

Thermal Power Systems

1. Most of the World's power systems are predominantly thermal. The main planning problem is then to provide adequate plant capacity to meet the anticipated maximum load, with due consideration of plant outages for scheduled maintenance and accidental breakdown. The main operational problem is to divide total load at any given moment among the existing units so as to minimize total losses (maximize overall efficiency). If in addition some hydro power is available, planning will also be concerned with the extent to which hydro power can be relied upon to cover peak loads and also rapid load variations at off-peak times, on the basis of recorded runoff and storage possibilities available. But there is usually no problem in covering the energy demand as long as demand at all times stays within the limit of anticipated peak load. It is then simply a question of purchasing sufficient fuel from the market.

Hydro Power Systems: Additional Problem of Water Storage

2. In a predominantly hydro power system it is still necessary for the planning to provide sufficient total plant capacity to meet the peak load, just as in the thermal system. But in addition, the hydro system presents also an energy problem, which tends to be the overriding one, thus making the planning considerably more complex than that of the thermal system. Whereas fuel for thermal plants can be bought as required, water supply to a hydro plant cannot be controlled except by storage, carrying water over from surplus period to deficit periods. Once drawn down, the storage cannot be replenished except by the run-off that may or may not occur. The larger the storage volume is, the better are the

chances that the problem will be mastered at all times, but there is always a certain risk that deficiencies may occur. Hence, the main operational problem is now to operate the given storage facilities so as to minimize the occurence of failures, or rather, to minimize losses incurred by power failure. The main planning problem correspondingly is to provide, in addition to sufficient plant capacity, such storage facilities as will maximize total net benefits from the system, that is, value of power produced less real costs as well as cost of power failures. In highly developed hydro power systems, aiming at near full utilization of available water resources, the risk of occasional power failure can be brought down to zero only at a very high cost, so usually some small risk is accepted, as the best solution, on the reasoning that further reduction could be obtained only at a cost which exceeds the value of loss reduced.

3. The existence of some thermal power within the system does not change the nature of the problem. Thermal power would be used in this case to supplement the storage, thus reducing the storage volume required.

Limitations in Conventional Analysis

4. Since nothing is known about the precipitation and run-off in the future, analysis of a hydro power system is usually based on the assumption that runoff conditions in the future will remain by and large as they have been in the past. This applies also to computer analysis.

5. In addition, conventional analysis has rather severe limits set by the large amount of computation that would be involved. Hydro power developments involving no or very little storage are often analysed on the basis of the runoff duration curve.[1] Power corresponding to the runoff available at all times (100% duration) or, in some cases, at an arbitrary duration near 100% is termed "firm power". Power in excess of that, up to the limit set by the duration curve and plant capacity, is "surplus" or "secondary".

6. The most widely used method of analysing hydro power systems with storage uses the "mass curve" as main computational aid. The mass curve, representing cumulative total of runoff from the beginning of the period of analysis, is simply compared to the cumulative total of discharge which, for a given uniform discharge, is represented by a straight, ascend-

[1] Indicating the length of time (in time units, or in per cent of total time) for which runoff equals or exceeds any given intensity.

ing line. Possible spillover (reservoir overtopping) is accounted for by shifting the discharge curve correspondingly. The maximum gap found between the curves (maximum drawdown) represents the maximum storage that would have been needed to assure operation over the period in question. Repeated calculations with different discharges (the mass curve itself remaining unchanged) establishes the functional relationship discharge to storage required, from which, conversely, the firm discharge (and firm power) obtainable with a given storage volume can be calculated. Some times the criteria for "firm power" are relaxed, in that an arbitrary number of failures is accepted – say, one in every ten years, so that storage needed for a certain amount of "firm power" is the storage which would be sufficient in all but 3 years in a 30 year period of analysis.

7. Some severe limitations will immediately be seen in this. No indication is obtained about the probability (small, but of considerable consequences) of a power failure in the future even though the storage may have been sufficient for all years on record. When a lower than 100% confidence is required (failure allowed to occur in a certain number of years), nothing is said about the duration and severity of the failure. Even more far-reaching, no account is taken of the deviations from normal operation that may be made, to forestall an imminent storage depletion or to utilize temporary surpluses. To a certain extent, these and other deficiencies can be corrected by adjustments to the conventional methods. But the full benefit from the vastly increased calculation possibilities presented by the electronic computer could only be obtained by carefully restating the problem as a whole. That is the subject of the following chapter.

II. RESTATING THE PROBLEM FOR COMPUTER ANALYSIS

8. The problem in planning an expansion to a power system (expansion from zero, in the special case of a new system) is that of selecting, among a number of alternative expansion projects, the one that will give the "best" result. Usually the result is to be measured in terms of benefits versus cost. The benefit-cost criteria can be variously formulated. But if we can assume that all costs related to each project can be defined, including the cost of capital (interest, real or imputed), we can also as-

sume that among a number of projects all yielding the same benefit, the one is best that represents the lowest cost.

9. Considering the low elasticity of demand for electric power, and assuming that the new project will not represent a very marked deviation from the marginal cost on which current selling prices have been based, the *demand* which our expansion project has to meet can be projected with some accuracy, and our planning problem can be regarded as that of selecting the best expansion project to meet this demand. All the alternatives that *do* meet the required demand then yield the same benefit. The best alternative then is the one that has the lowest cost.

10. More often than not, however, one finds that some or all alternatives are capable of yielding additional benefits (saleable surplus power) or of falling short, occasionally, of the required output. Additional benefits can then be offset against costs (reducing the cost). Shortfalls must on the other hand be *added* to the direct costs.

11. Proper planning thus requires, for all alternatives,

a) Cost estimates – initial or investment costs as well as costs of operation,

b) Estimated value of the scheduled firm power production,

c) Estimates of obtainable additional benefits as well as of additional costs due to non-fulfillment of the firm production goal.

Item b), however, is a figure common to all alternatives, and can thus be left out of consideration as far as selection among alternatives is concerned (but it does enter when it comes to deciding about carrying out any development at all).

12. Provided cost estimates for investment and operation are establisned for all alternatives in the conventional way, the remaining problem is that of estimating, for each alternative, the "cost corrections" item c) of para. 11. This can be done only by trying out, i.e. by *simulating* the operation of each alternative over a "sample period" for which runoff records are available.

13. Such a simulation, to be realistic, must make use of no "hindsight". The technique must incorporate, with sufficient approximation, the same type of tactical decisions under uncertainty about all but the nearest future, on which the real operation of a power system has to

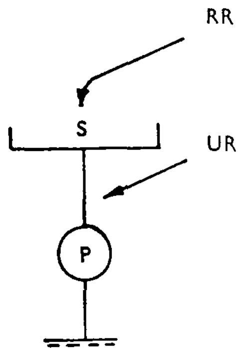

Fig. 1. *Single Plant Model.*

rely. The following chapter will describe the basic concepts of a simulation technique for a "single plant" model of a power system, based on original suggestions by Stage and Larsson.[1] The subsequent chapter (ch. IV) will then describe the further elaboration of the same concepts using a multiple plant model that more closely represents real conditions in a complex system.

III. THE SINGLE PLANT MODEL

The Production Model

14. In the simplest approximation, a hydro power system can be described in terms of a single plant model as diagrammatically shown in Fig. 1.

The total plant capacity of the system is concentrated in one power plant *P*, equipped with a storage reservoir *S*, equal to the combined storage capacity of the system. The runoff received by the system can be divided in two parts, one that flows to storage (regulated runoff, *RR*) and one that is not controlled by storage but flows directly to plant (unregulated runoff, *UR*). The adding up of storage, regulated runoff, and unregulated runoff, throughout the system is done after translating them from quantities of water into *energy* units (storage kWh or GWh, runoff kW, MW or better kWh/day, GWh/month, etc.).

[1] Footnote, page 126.

15. The functioning of this model reflects that of the system as follows: Total system load is met from total system capacity. If total system runoff is insufficient to cover the load, additional energy must be drawn from system storage. It does not matter for the consumer which of the individual storages is drawn on, as long as the required energy is provided.

16. The model would be correct were it not for the following limitations which are imposed in reality, but which the model ignores:

a) Runoff from anywhere in the system cannot be channelled freely to any individual reservoir (some individual reservoirs may well be filled up, unable to receive any more, even though the system as a whole has some unused storage capacity left).

b) Energy cannot be drawn freely from each individual reservoir to cover system deficit (drawdown from a well-filled individual reservoir may be limited by insufficient capacity in the plant downstream of that particular reservoir).

A model respecting these limitations will be described later. The present chapter will describe the use of the single plant model as a first approximation, ignoring the limitations above.

The Demand

17. The desired power output is one that matches exactly the firm power demand. The value of this output, if it could be obtained, would be equal for all alternatives and could thus be left out when making comparisons between them. But if production is for some resaon reduced below the curve, this represents a *loss* which can be different from one alternative to the other, and which must therefore be taken into account. The loss is inflicted partly upon the power company, which loses the corresponding revenue, and partly upon the consumer, since withdrawal of an energy unit for which he has already equipped himself, represents a higher value than the price of the unit itself (such as production cutbacks if the consumer is an industry, discomfort from a cold house if he is a private person, etc.). Our analysis will be based on the total value of the loss – to power company and consumer – but it will be seen that exactly the same methods are valid also in case one wants to base an analysis on the economic considerations of the power company alone.

18. The imputed value or "utility" of a kWh about to be withdrawn must obviously increase as the power deficit increases, as shown in Fig. 2. The difficulty is of course to estimate how steeply the unit value rises. This is a point that needs considerably more analytical and empirical exploration. On the other hand experience shows that rather wide variations – say, varying the marginal value at 45% cutback between 10 and 30 times the marginal selling price – have a rather modest influence on the planning conclusions, i.e. on the ranking of projects, because the mode of operation will in any case adjust itself to the marginal utility assumption that is chosen. This particular problem can therefore be left aside while we carry on with the method analysis, assuming that a plausible utility curve has been established, as shown in Fig. 2. The curve in Fig. 2 has also got a tail end: That is the value per unit (kWh) of surplus power that may be produced. This value usually is markedly lower than the price of firm power, as indicated in the figures; and above a certain limited quantity, surplus power is not saleable at all (price = 0).

19. In this shape, out curve can be regarded as a "cross section" of the seasonal load curve. Furthermore, the curve could be interpreted as a *short term* demand curve for power.

20. The firm power demand (in the figure called Normal Supply) as well as the shape of the marginal utility curve may vary with the time of the year. A full description will thus take the form of a three-dimensional matrix.

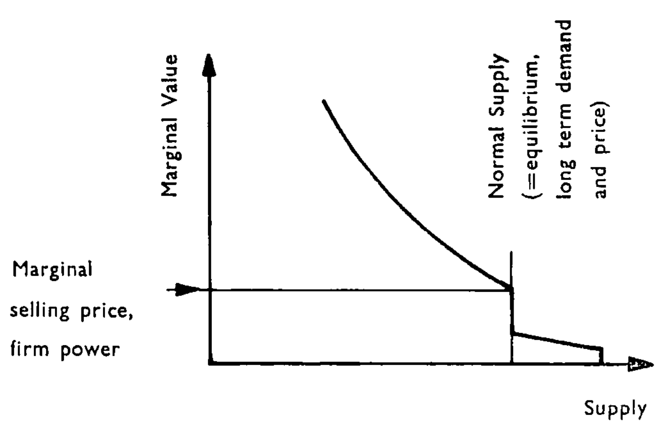

Fig. 2. *Marginal Utility Curve*

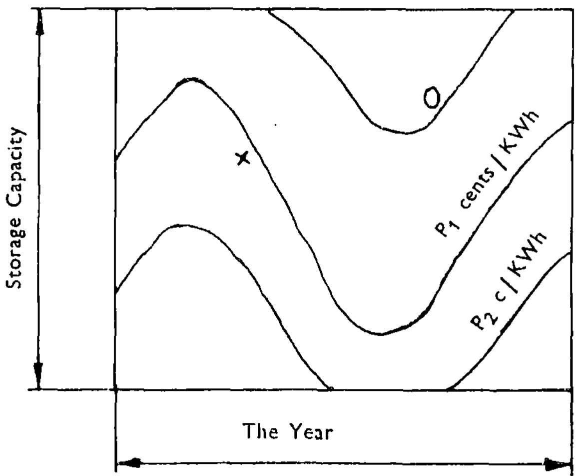

Fig. 3. *Marginal Value of Storage (value of marginal unit stored)*

Simulating the Operation

21. With the demand function thus established, the base is laid for the economic operation of the system: The operation must seek to maximize the economic value of the production. The production is regulated through the control of storage. The (short term) marginal cost of production is equal to whatever *value* a unit of energy drawn from storage may have at the moment; other marginal costs are negligible.

22. The value of a marginal kWh in storage resides in the value it may have if it is *not* used now but withheld and used later: At a later time, its value (realized) is equal to the marginal value of demand at that time. The problem of calculating the marginal value of energy stored may be left out for a moment, while we just suppose that the values are known, and could be demonstrated in a graph like Fig. 3.

If actual energy stored at a certain time of the year is such as marked by **x** in the figure, production at that moment should be limited so as to satisfy only that part of demand which has a value equal to or higher than the marginal storage value at **x**. The correct production level is then found by entering the marginal storage value in the graph Fig. 2. If production is kept at this level for, say, one week, this, together with runoff received, will determine the energy stored at the end of the week. We now again take the marginal storage value from Fig. 3, enter the

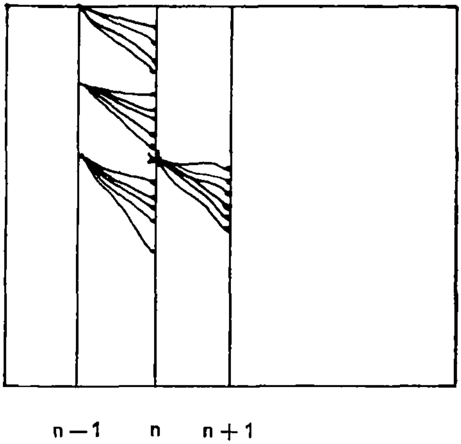

Fig. 4. *Calculating Marginal Value of Storage*

graph Fig. 2 and thus determine production for another week. In this way we could go on simulating production, week by week, throughout our sample period (or 30 years or so).

23. But first we have to establish the diagram of Fig. 3. Stage[1] has shown how this can be done by an iterative process, based on the following reasoning:

A kWh in storage should just be withheld, if its value at the moment is exactly the same as the value it can be *expected* to have at a later point in time. If we are at the point **x** in Fig. 3, we do not know what the runoff will be in the next time interval, and hence which storage situation we would end up with after that interval. But we do have a sample runoff for that time of the year, in previous years. If we regulate the production according to the assumed marginal value at **x**, and calculate storage development for each year in our sample, we get 30 different developments as illustrated in Fig. 4. The marginal kWh stored obtains 30 different values at time $n+1$, the *average* of these values is the value that we *should* have expected at time n. The calculation procedure follows from this (see Fig. 4):

24. Start with an assumed set of marginal storage values, like Fig. 3. Start at time n. For a number of storage amounts (i.e. a number of

[1] Stage, Sven, and Larsson, Yngve: *Incremental Cost of Water Power. Power Apparatus and Systems (AIEE)*, August 1961.

points on the vertical through n), do the following: Read (assumed) marginal value; determine production for the following interval (from the graph Fig. 2); calculate resulting storage amount at time $n+1$ for each of the years on record; Read and average the marginal values at $n+1$. This average represents a *better* approximation of the correct value at time n; the original figures at n are therefore deleted and replaced by the better ones. Then move one step backwards, to time $n-1$. Do the same calculation as above; averaging the marginal values obtained after one step forward (i.e. at time n) now provides a better approximation of the correct value at $n-1$. And so forth. Proceeding step by step backwards, one eventually reaches the same time of the year again (time n), now with better (more accurate) marginal storage values. Calculation is continued for so many rounds (=years) as are needed to obtain stable values.

25. If the storage amount at any time exceeds maximum possible storage, the water is lost, and the value to be ascribed to this case is zero. If on the other hand the storage amount becomes negative, it means that storage is empty and production can be supported only by direct runoff. The value of an extra kWh available in storage in this case would have been equal to the marginal value of the production that can be sustained by runoff alone. This can be read off by entering runoff (in energy units) in the graph Fig. 2. – In fact, these two extremes, storage empty and storage overflowing, represent the only sources of "true" information about storage values. All other values are derived from this information by establishing through successive iterations the internal *relationship* that must exist according to para 23.

26. The basis has now been established for the whole computation process, which falls in two parts:

a) Iterative computations to establish a set of marginal storage values that satisfy the prescribed conditions.
b) Simulation of the operation throughout the time series for which runoff records are available, using the set of marginal storage values established in part a) in conjunction with the short term demand curve of Fig. 2 to determine system production for each interval.

The length of intervals would be chosen at anything between 5 days and 1 month. In most cases, including those reported later in this paper, the author has used intervals of two weeks.

27. The mode of operation thus simulated would maximize net benefits – or minimize net costs, which under our assumptions is the same. We have chosen to use net costs (cost of losses less revenues from surplus power) which are then recorded for each interval during the simulation, and added up for the simulation period. Dividing by the number of years simulated, one gets the average annual net cost of losses for this particular project. – The same procedure, iteration followed by continouous simulation, is then repeated, for each of the alternative projects under consideration.

Comparison of Alternatives

28. Comparison between the various alternatives can then be effected, as to

a) Average annual net cost of losses as obtained from the simulation, *plus*

b) Annual fixed costs (depreciation, interest, maintenance, salaries etc.).

The value of annual firm power production (not corrected for power deficits) need not be included in the comparison since it is the same for all alternatives (all are designed to meet the same firm power demand, cf. para 17).

Combination of Thermal and Hydro Power

29. The effect of adding a termal power plant to the hydro power system can be easily studied by just a small modification to the short term demand curve, on the following reasoning:

30. If a thermal power plant is available in the system, the decision to use it or not would depend solely on its variable cost of operation. As long as the marginal value of power delivered from the system is lower (per kWh) than the variable cost connected with using the thermal station, the station is not used. Such low values prevail when the storage situation is good. Suppose now that the storage situation deteriorates; the marginal value increases, production is adjusted downwards. At a certain stage, marginal value of storage (and correspondingly, of power delivered) reaches equality with the variable cost of thermal power, and eventually exceeds it. But then it becomes more economical to *operate*

the thermal station and save on the water in storage. The thermal station will be kept operating, at *full load*, as long as marginal value of production (and of storage) stays above the variable cost of thermal power.

31. This means that the (short term) demand for *hydro* power alone is now expressed by a curve that has been shifted to the left – above the level of variable cost of thermal power – for as much as the full capacity of the thermal station, see Fig. 6.

Marginal storage values are then determined, and simulation performed, on the basis of this adjusted demand curve. It will be realized that what is now being studied is, strictly speaking, the optimal operation of the hydro part of the system, in the *presence* of thermal power. Furthermore, since thermal power is used as a "substitute" for losses (which would have occurred, had the termal power not been there), the "cost of losses" recorded during the simulation, according to the curve Fig 5, does include the variable costs for thermal power. Only the fixed costs (such as maintenance) need be included under item b) in the comparison described in para 28.

IV. EVALUATING RESULTS FROM THE SINGLE PLANT MODEL

32. The next step in the planning procedure will be the systematic comparison of all the alternative projects, on the basis of the net costs as established by the simulation (cfr. para 11). The techniques to be used for this comparison may vary from case to case.

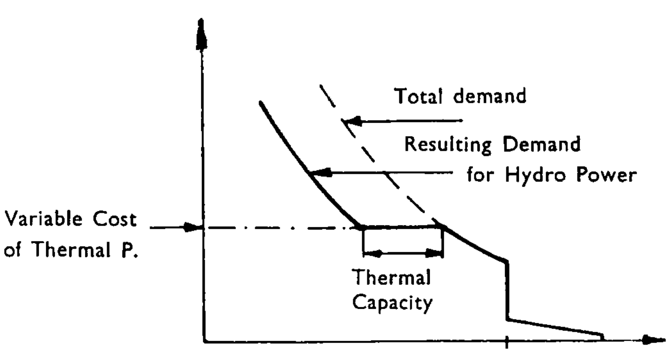

Fig. 5. *Effect of Thermal Back-Up Power*

33. In planning the expansion of a large, established system, the possible expansion projects will often be of rather uniform quality, since they all represent projects that came close to being implemented at the previous expansion step. In the case of a large river basin development, the various projects under consideration may be just variations of one single project – a little more or a little less runoff collected to the intake, a little more or a little less storage provided, etc. In the general exploration of such cases, the textbook method of combining equal-product curves with equal-cost curves can be a useful tool: A contour diagram is drawn up, with contours linking all combinations of annual runoff and storage capacity that will yield the same "product" (i.e., the same net losses according to simulation results). Another set of contours is drawn up, with contours linking the combinations that have equal cost (other than costs already accounted for in the "product", i.e. in the simulation result). If the marginal cost of collecting runoff and that of providing storage volume are constant in the area explored, these latter contours will be straight lines. An example is shown in Fig. 6, from a case in which system parameters had been scaled down uniformly to a "scale model" confronted with a demand of 100 000 units a year. The optimum solution is represented by the combination for which the slope of the equal-product surface is equal to that of the equal-cost surface.

V. A MULTIPLE PLANT MODEL

34. The single plant model ignores some important limitations which are imposed upon the real operation of a hydro power system composed of a number of power stations. The single plant model operates as if all runoff in the "regulated" category can be freely channelled to any individual storage within the system, and as if drawdown requirements can be taken from any individual storage without regard to plant available at each storage, if only total plant capacity is respected.

35. In many instances the observance of these limitations may not make so much of a difference. If all the individual storage reservoirs and power plant are of nearly the same size relative to their runoff, there would be no need for them to call upon each other, and they would act practically as a single plant system. Even if they are more different in characteristics – as long as the combined operation can be carried out

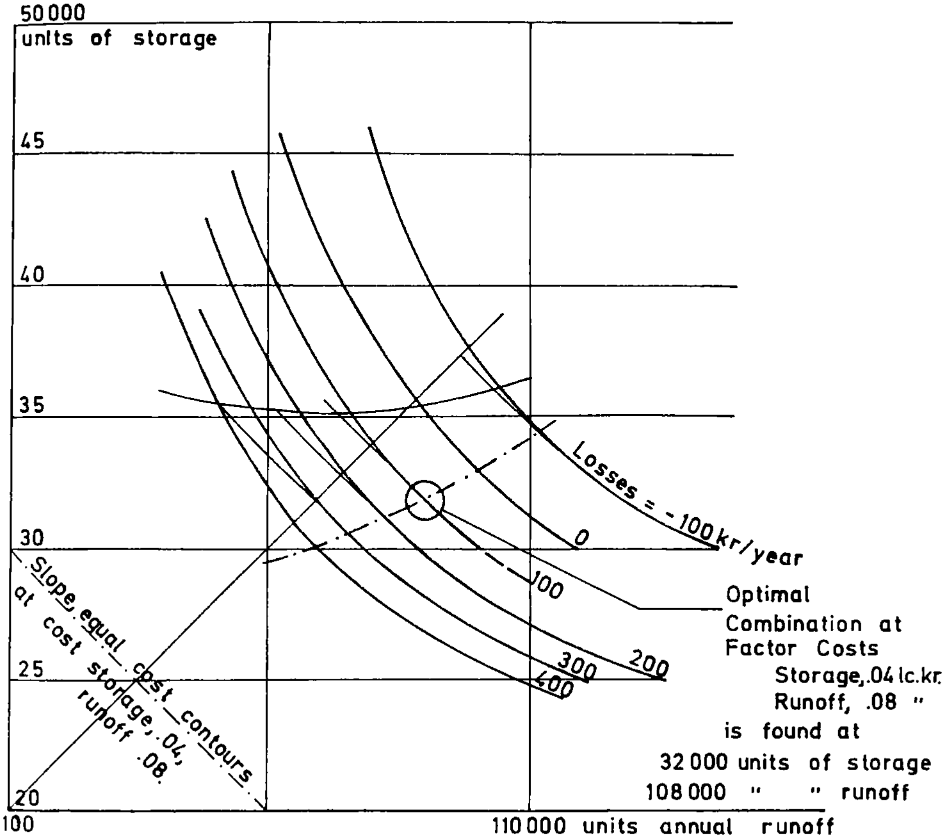

Fig. 6. *Equal Product Curves*

so that no single reservoir is overflowing before all reservoirs are filled up, *and* so that no single reservoir is empty before all are empty, then the result is the same as if all reservoirs were added together.

36. If reservoir and power plant characteristics differ too much, the limitations may cause the real operating result to be inferior to that of the single plant model. Then the real operating result can be found only by simulating a model where individual reservoirs and plants are represented, and dividing (allocating) the combined system load on these individual reservoirs and plants.

37. What is then sought, is a measure of what the system can yield when operated in the best possible way, including the best possible or "optimal" division of the combined load on individual plants. A search for *the* optimal way of dividing up the load would stand small chances of success, however, for formal reasons: In the general case, there would be an infinite number of solutions all as good as "the" best one. Just think of two plants, each with a reservoir, within a large system: A wide variety of load-sharing between them in week no. n would make no difference to the long run result for the system if the differences could be compensated the other way in week no. $n+1$, and so forth.

38. Instead, one can try to find *a* solution that yields a result as *close as possible* to the result of the single plant model (no real solution can yield a better result than the single plant model, for, a system with limitations cannot work better than one where the limitations are lifted – at most, it can work equally good).

Whether there are also other solutions which are *equally* close to the single plant ideal, need not disturb us: What we need for our system planning is just a measure of how good a result can be obtained from the development planned. This is the approach that has been followed in working out the following multiple plant model.

39. The basic concern in the operation is for the storage reservoirs: None of them should be filled up before the others, and none (or at least, not too many of them) should be emptied before the others. Storage movements both up and down will generally be quicker the smaller the reservoir is, relative to the runoff it receives. In the diagram Fig. 7 the reservoirs of a system are shown lined up according to this relationship – volume as percentage of average annual runoff:

Since the width of each column represents average annual runoff, the area of the column represents the volume of the reservoir. Fig. 8 shows a similar line-up for a very large system.

Now, when the reservoirs approach the "full up" situation, it is desirable that the total storage is divided between individual reservoirs more or less as indicated by the line a–a. Only then can premature filling up of one reservoir be avoided, until all reservoir space has been effectively used. Similarly, the reservoirs should approach the "empty" position along a line like b–b.

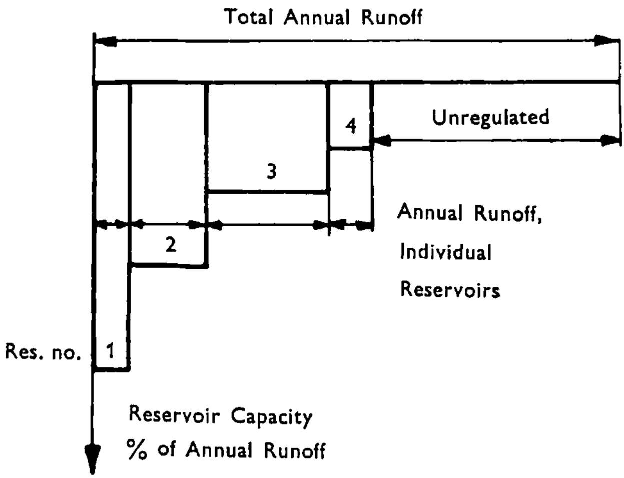

Fig. 7. *Classification of Reservoirs*

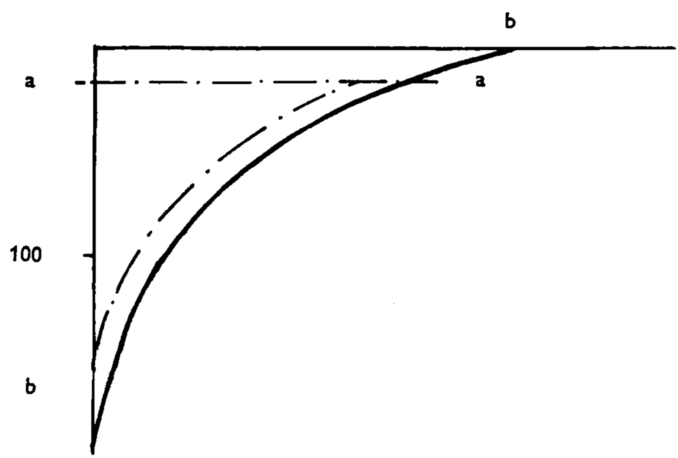

Fig. 8. *"Ideal" Breakdown of Total Storage at Filling (a – a) or Emptying (b – b)*

40. On this basis, a table could be worked out, showing how any given value of total storage should "ideally" be subdivided on the individual storages. Operation should then aim at keeping as close to this ideal as possible.

(Ideal subdivision could also be different for different times of the year, etc. Effects of different "ideals" could in fact easily be studied on the model itself, if deemed necessary.)

When operating the system, aiming at the aim now established, the following limitations will have to be observed: Filling of a reservoir cannot take place quicker than runoff allows, all downstream plants closed down. Emptying cannot be made quicker than by running downstream plants at full gate.

41. Fig. 9 shows diagrammarically the system to which this multiple plant model was first applied, the Hvitá–Thjorsá river system in South Iceland. (Fig. 9.) The components are those typical of any large hydro power system: Storage reservoirs in parallel (reservoirs in series are equally amenable to the model), power plants in series and in parallel, branching river systems, etc.

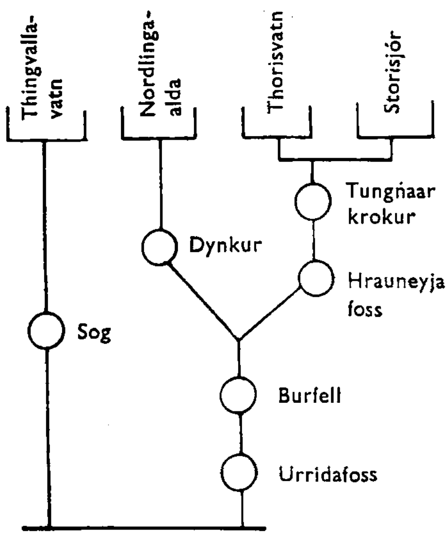

Fig. 9. *Multiple Plant Model*

The computer program for simulation of the operation of this model has 3 main parts:

Part 1 (processing of unregulated runoff): For each power plant in the system, unregulated runoff in the time interval under consideration is checked against plant capacity. Runoff in excess of capacity is registered as waste. If runoff is less than capacity, unused capacity is recorded (available for processing of water from storage, part 3). If all demand for power in the interval concerned is met by processing unregulated runoff, the program stops, records runoff at remaining plants as waste, and proceeds the to next time interval. If not, the program proceeds to part 2.

Part 2 (preparatory steps to processing storage): Total system storage at beginning of interval (=end of previous interval) is entered in the diagram Fig. 3 (or rather, in the corresponding matrix) to find marginal value of storage. This in turn is entered in the diagram Fig. 2 (or matrix) to determine how much power "ought to" be supplied in the time interval. If this has already been satisfied from unregulated runoff (part 1), no storage will be drawn down, and the program proceeds to the next time interval. If not, follows part 3.

Part 3 (processing of storage) starts by comparing "ideal" content of each reservoir (at the prevailing total storage of the system) with actual content, and working out ratios, actual to ideal. The reservoirs are then arranged by order of decreasing ratio. Those first in this line-up should be drawn as much as possible in the next interval; those at the end should be left alone to fill up. This is done, simply by discharging from the first reservoir the quantity needed to achieve full capacity load in the "bottleneck" plant downstream of the reservoir – i.e., that one of the downstream plants that shows the lowest "unused capacity" according to part 1. The same is done with the next reservoir in line, and so on, until the total system load has been covered. The remaining reservoirs are left unused. The procedure may seem crude in the sense that some reservoirs in the middle of the line, which are only slightly off the ideal, will get either full drawdown or full filling, nothing between. But the consequence of this is only that in the next interval, these reservoirs may jump over to a different place in the line, and be treated accordingly, next time.

42. The program contains several additional routines, to observe limitations in plant utilization set by dayly load variation as well as by legal or contractual requirements; to convert from energy to water discharge figures for establishing bottlenecks, etc. etc.

Marginal Storage Value in the Multiple Plant Model

43. The marginal value of storage, which determines how much power should be produced by the system (total system load) must still be established by iterative computation on the single plant model. Similar iterations on the multiple plant model would exhaust a very big computer, and would also be formally difficult. The error in marginal storage values resulting from the single plant idealization is probably small in the general case (the important errors are those in the final simulation); corrections *can* be fed back from the multiple plant model but in the South Iceland case the need for this appeared to be very small, since the main storage reservoirs turned out to be filled and emptied very much in parallel.

Particular Uses of the Multiple Plant Model:

44. The multiple plant model is very versatile, in that the effect on overall system economy caused by a change in *any* of the system components can be stested, by simulating the system *with* and *without* the change. It may be a reservoir, or a plant capacity (particularly when a certain plant is suspected to be a bottleneck in reservoir operation) or even a transmission line out of a particular plant.

45. Thermal plants are introduced indirectly, as in the single plant model, by modifications (leftward shifts) in the short-term demand curve.

VI. A MODIFIED MULTIPLE PLANT MODEL

46. In very large systems such as that of south-east Norway, a multiple plant model may comprise several hundred power plants and a similar number of storage reservoirs, presenting considerable problems in data handling. For many uses a modified model is fully satisfactory, in which power plants with approximately similar reservoir character-

istics (reservoir volume in per cent of annual runoff) are lumped together in groups, each group being represented in the model by one plant. Fig. 10 shows the result of one such classification and aggregation of the South-east Norway system as of Januar 1, 1967, resulting in a model consisting of 3 plants without storage (and with varying ratios capacity to average runoff, Q/R) and 5 plants with storage (reservoir volumes varying from 15 to 300% of average annual runoff).

47. The main feature of this model compared to the complete model of chapter V – apart from size – is that there are no plants or reservoirs connected in series. The corresponding routines can be omitted from the computer program, but for the rest, the simulation procedure remains the same as in the complete model. The working of the modified or "parallel model" has been tested for the South-east Norway system by simulating operation during 1965 and 1966, on a weekly basis, and comparing to actual operation, with extremely satisfactory results. That means that even in this modified form, the multiple plant model overcomes what was the main weakness with the single plant model, namely the aggregation into one reservoir, of individual reservoirs, with widely disparate filling and emptying conditions.

48. The modified model is particularly well suited for general and exploratory studies on the system as a whole, where the main concern is not for single features of individual plants. Thus, the model is being extensively used in Norway to study the relative economy of hydro power versus thermal power expansion. The technique used can be briefly described as follows:

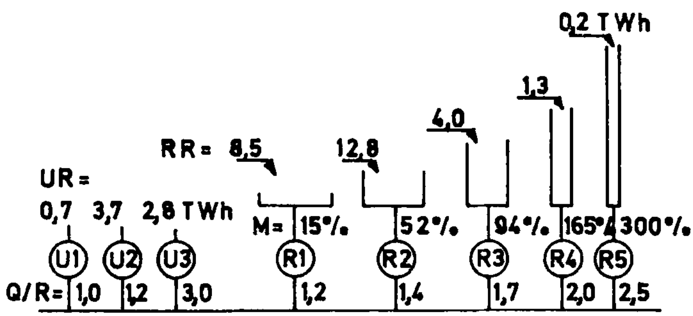

Fig. 10. *Modified Multiple Plant Model, South-East Norway System.*

49. Starting from a system in equilibrium, the demand is assumed to increase by a small amount. The system is then expanded, by expanding just one of it's parameters in turn (runoff in group 1, or group 2, or 3, and so forth, or storage group 1, 2, 3, and so forth, or hydro plant capacities, or steam plant capacity, gas turbine capacity, or nuclear capacity), to find how much each one of these parameters would have to be increased if that alone should bring the system to catch up with the increased demand. Then follows a comparison of costs. Using f.i. the known (estimated) annual fixed cost of conventional steam power as a reference, it can be found how much the development of one additional kWh of runoff to plants in group 1, 2, 3, and so forth, or 1 kWh of storage volume in each of these groups can be allowed to cost, without being less advantageous than steam power. For the system as of Jan. 1, 1967, the "steam equivalent value" of 1 kWh of annual runoff was found to be Norw. kr. 0.008 per year for all groups, and that of 1 kWh storage was also found to be Norw. kr. 0.008 per year for all groups. It will be noted that the figures are annual costs. Corresponding "permissible development costs" are found by multiplying approximately ten times.

50. This establishes a ready scale for preliminary evaluation of any proposed hydro power expansion. If the total estimated cost of the project is higher than the "steam equivalent value" of it's runoff, plus that of it's storage valume, plus a value ascribed to installed capacity (equivalent to what pure capacity can be obtained for, marginally, in other developments), then the project can be expected to be less favorable than an expansion of the system's steam capacity.

51. As already indicated, the relative value of nuclear capacity is also brought out by these studies. For the time being, 1 kW of BWR capacity, with variable costs at Norw. kr. 0.01 per kWh, is found to be worth only about Norw. kr. 95 per year fixed costs, or less than one half of what such capacity would actually cost. The explanation for this seemingly unexpected result is the very favourable interplay that can be counted on between hydro power and conventional steam power, until the steam power component of the system (now practically zero) reaches a certain size. From then on, nuclear power will gain in relative economy.

52. For detailed planning of the system, i.e. for the planning of each individual plant, studies on the complete model will still be required. A "hybrid" model is also being contemplated, in which only the river

system to which the plant in question belongs is represented in full detail, while the remainder of the system is represented in aggregate form as in the modified model. This would give nearly the same accuracy as a complete model, while keeping the data handling problems within bounds.

VI. APPLICATIONS OTHER THAN POWER

53. One obvious application for the same models as described above, is irrigation. Once a marginal utility curve (as per Fig. 2) has been established, for water applied to irrigation, at various times during the season, the optimization procedures would be exactly the same as described above for power. Actually, establishing the marginal utility would probably be easier in agriculture than in power. More generally, the simulation techniques described above would be applicable to the analysis of multiple purpose water resources development, including power, irrigation and other purposes (navigation, water supply etc.).