6

GENETIC ALGORITHM OPTIMIZATION OF IRRIGATION CHANNEL DELIVERIES

By J.B. Nixon,[1] A.R. Simpson,[2] Member, ASCE, and G.C. Dandy,[3] Member, ASCE

Abstract: This paper identifies water delivery schedules for water orders from an off-farm, open-channel, gravity fed, irrigation network using genetic algorithm (GA) optimization and suitable representations of significant objectives and important constraints identified for this system. In optimizing a schedule of water deliveries, objectives include satisfying customer service agreements and minimizing variations in channel flow rate. Avoiding channel capacity exceedance is one constraint. The GA methodology is applied to and compared against a real-world order scheduling task carried out by a planner at an operating water authority during a past irrigation season. The results show great promise in the ability of GA techniques to identify efficient irrigation order schedules. It is proposed that these optimized schedules will in future be suitable for review and finalization by irrigation planners, and then implementation by field operators.

[1] Senior Research Scientist, Research and Development, United Water International Pty Ltd, PO Box 1875, SA 5001 Australia, T: +61 8 8301 2834, F: +61 8 8357 9728, E: john.nixon@uwi.com.au
[2] Associate Professor, School of Civil and Environmental Engineering, The University of Adelaide, Adelaide, SA 5005, Australia, T: +61 8 8303 5874, F: +61 8 8303 4359, E: angus.simpson@adelaide.edu.au
[3] Professor, School of Civil and Environmental Engineering, The University of Adelaide, Adelaide, SA 5005, Australia, T: +61 8 8303 5472, F: +61 8 8303 4359, E: graeme.dandy@adelaide.edu.au

Key Words: genetic algorithm, irrigation network, off-farm, open-channel, optimization, scheduling

paper is an extension of that work. The reader is directed to that work for background on the problem which was, due to length restrictions, required to be omitted from the present paper.

In this paper the GA techniques are applied to a real-world irrigation system, which consists of a branched open-channel, gravity fed, network supplying over 150 irrigators in an Australian farming community. Each of the irrigators requests an amount of water to be available for their use at the off-take point for their property at certain time. Control structures, chiefly in the form of weirs, exist at the network branch points and other critical locations. It is the during the scheduling activity—carried out at present by water authority planners—that these requested orders are rearranged so that adequate and equitable supply can be delivered without overloading the network. In the future this activity will possibly be carried out with the assistance of the GA methodology presented here.

Advance Notice Ordering for Irrigation Water Deliveries

Irrigators are usually required to request their orders a number of days in advance, with this number specified by the customer service agreement between the irrigation authority and their customer-irrigators. The orders requested to be filled over a certain period of time, usually a single day, are then considered together for scheduling by an irrigation planner. The scheduling of orders for a particular day is usually performed by a planner one day in advance. The advance notice ordering system enables irrigation planners to schedule a set of orders for a given irrigation day such that maximum use is made of the existing channel infrastructure. The scheduling of orders in this manner also allows for the greatest number of orders to be satisfied. At present, the planners have the ability to grant partial orders, i.e. they can reduce the flow rate or flow duration relative to that requested. (In this paper the GA is not formulated so as to allow this. In the case study

SCHEDULING SIMULATION

As part of this research, a simulation was carried out in conjunction with Goulburn–Murray Water (G–MW)—a major rural water authority, and Rubicon Systems Australia Pty Ltd (Rubicon)—a supplier of irrigation control systems. G–MW's region covers 68,000 km^2 of the northern plains of the Australian state of Victoria. This includes the state's major irrigation districts divided into six management areas: Shepparton, Central Goulburn, Rochester–Campaspe, Pyramid–Boort, Murray Valley, and Torrumbarry, as well as three Waterworks Districts and ten river basins. G–MW provides rural water and drainage services to approximately 24,000 properties and manages 6,800 km of irrigation channel and 19 storages, involving 24,500 control structures. The average annual volume delivered to the irrigation area is approximately 7,000 GL per annum (Goulburn–Murray Water, 2000). The case study detailed below used G–MW's database of the Central Goulburn irrigation area, and Rubicon's ordering and management software Irrigation Planning Module (IPM). The authors' GA software was integrated with Rubicon's IPM software by co-developing an appropriate application programming interface (API) between the two otherwise independently developed software layers. Other details of the simulation, and the results thereof, are described in detail, in the case study, later.

OPTIMIZATION GUIDELINES

The *scheduling guidelines* will depend on the priorities of the particular water authority. In this section, those used by G–MW are described. These were obtained by interviewing G–MW personnel. They are typical of irrigation areas where network capacity constraints are often an

Calculating Flow Rate Time-Series at Control Structures

Time-series of flow rates at designated control structures in the irrigation network need to be determined for each particular schedule. These are based on a mass balance, i.e. a summation of the flows downstream of the structure, taking into account the time of travel from the control structure to the relevant off-take point location. At present within the IPM system at G–MW, the channel hydraulics are simulated using such a simple mass balance approach. More sophisticated, and physically realistic, simulations of the unsteady non-uniform flow in irrigation channels (Schuurmans 1991) are beyond the scope of this paper. They would, however, improve the scheduling optimization process proposed, were they to be implemented in the IPM system.

Each time-series represents the flow past a particular control structure in the irrigation network assuming that the orders are delivered as scheduled. At each control structure, it is a constraint that the flow rate should not exceed the channel capacity.

Orders Shifts Fitness Constituent: ϕ_1

The first fitness measure, the *order shift fitness factor*, $\phi_{1a,o}$, is based on the magnitude of the order shifts applied, and is designed to lead the GA to schedules that satisfy optimization guideline A. For example, a proposed schedule may require that an order be shifted 6 hours backwards in time or 12 hours forwards in time, relative to the start requested, and one of these shifts may be deemed by the water authority and its customers to be more desirable than the other.

Another fitness measure, the *order notice fitness factor*, $\phi_{1b,o}$, is based on the notice given by the irrigators for the orders placed, and is designed to encourage ordering schedules that fulfil off. The peaks on the left of a 0 hour shift decay with a half-life of 24 hours, and those to the right decay with a half-life of 12 hours. The troughs of $\phi_{1a,o} = 0$ occur at 12 hour intervals either side of the 0 hour order shift.

Please replace this line with Figure 3, or place Figure 3 near here

The function illustrated in Fig. 3 applies to orders of "long" duration, i.e., those of duration equal to or greater than the cut-off specified by the GA operator. Order shifts of 12 hour multiples are highly rewarded, otherwise those of 6 hour multiples are highly penalised. In general, holding orders off is preferable to bringing orders forward. The peaks on the left of a 0 hour shift decay with a half-life of 12 hours, and those to the right decay with a half-life of 24 hours. The troughs of $\phi_{1a,o} > 0$ occur at 6 hour intervals and decay on the left of a 0 hour order shift with a half-life double that of those to the right.

Note that the functions illustrated in Figs. 2 and 3 over a minus 48 hour to plus 48 hour order shift domain are set to zero for individual orders scheduled with less than the minimum allowable negative shift or greater than or equal to the maximum allowable positive shift. The cut-off between "short" and "long" duration orders is a user-input variable in the GA software but is chosen to be 18 hours in the case study presented later in the paper. This value, and the shapes of the two functions themselves, should be defined following consultations with appropriate irrigation management personnel, as was done here.

$$\phi_2 = \frac{1}{M}\sum_{c=1}^{M} w_{2,c}\phi_{2,c} \text{ such that } \sum_{c=1}^{M} w_{2,c} = 1 \quad (3)$$

where M is the number of control structures for which a channel capacity is defined.

Time-Series Flow Rate Variations Fitness Constituents: ϕ_3 and ϕ_4

Two measures of the "variations" in the flow rate time-series are used to determine the fitness of the associated water delivery schedule. To determine the two *flow rate variations fitnesses*, ϕ_3 **and** ϕ_4, the **same** function is used in two comparisons of two **different** flow rate time-series. The relationship holds for each control structure in the network chosen to be included in the time-series analyses. As per G–MW instructions, a proposed schedule that corresponds to a time-series at a particular control structure that has greater "variation" than that associated with the irrigators' requests is less desirable, and vice versa. The general shape of the variations fitness function illustrated in Fig. 6 (derived in consultation with G–MW personnel) leads the GA to schedules that satisfy optimization guidelines D **and** E.

Please replace this line with Figure 6, or place Figure 6 near here

A weight $w_{3|4,c}$ is specified by the GA operator, for each control structure c, to apply to the two flow rate variations fitnesses $\phi_{3,c}$ and $\phi_{4,c}$ so as to represent the relative importance of the two corresponding measures of time-series flow rate variations at that particular structure.

intervals in time are avoided, as per G–MW operator practice, yet are still taken into account within the flow rate regime. The control structure time-series is derived from the flow rate time-series by the algorithm outlined later, and detailed in Appendix I.

For each control structure, a *control structure regulations fitness* measure is calculated from the ratio of the number of regulations in the control structure time-series for the irrigators' requested schedule to the number of regulations in the control structure time-series for the GA's proposed schedule. This ratio is expressed as a percentage, from which a fitness constituent is calculated, using the function illustrated in Fig. 6, that rewards schedules with order shifts that result in control regulation time-series that are smoother than a schedule with all zero order shifts (i.e. the irrigators' requests), and vice versa. The control structure regulations fitness constituent, ϕ_4, is then defined by Eq. (5).

Two values, each representing a lower-bound associated with the regulations which may required to be made at any particular control structure, are required to be specified by the GA operator: (a) a value representing the smallest time change (in hours) allowed between regulations at any control structure; and (b) a value representing the smallest flow rate change (in % of channel capacity, converted to corresponding ML/day) allowed between regulations at any control structure.

Starting at the beginning of the time-series, if a change in the flow rate time-series (represented by a step in the time-series) equals or exceeds both of these lower-bounds, the step is counted as an "allowed" regulation and that step is recorded in the control structure time-series. If the dimensions of time-series step do not exceed both lower-bounds, the present step is combined with the next step (irrespective of the dimensions of the latter) to form a new step. The dimensions penalized and providing "more" customer service than is desired is **not** penalized or rewarded—thus aiming for efficient water authority practice.

The *delivery targets* as stated in the service agreement between G–MW and their customers are such that, for an individual off-take point and the total irrigation area, the G–MW will deliver:

1. 88% of orders on the day requested; and
2. 95% of orders on the day requested or within one day;

when

$$\text{notice days given} \geq \text{notice days required} \tag{7}$$

holds. In the case study presented below, a minimum of 4 days notice is required to be given.

The definition of the customer service fitness as given above is such that the delivery targets are **not** attempted to be met at individual off-take points. Rather they are attempted to be met ***only*** over the total irrigation area, as represented by an ***average*** over all orders requested from service points included in the plan area specified. Typically, the delivery targets would be calculated by the water authority over a period of time—perhaps even recorded as a running average over and/or across irrigation seasons—and so water authority irrigation planners would be able to not meet these on some plan days, and exceed these on others, while still meeting them over a certain time period. At present the GA, on the other hand, records no periodic or running average, and attempts to meet the delivery targets on ***every*** plan day, individually.

Please replace this line with Figure 7, or place Figure 7 near here

authority personnel operating the GA. Wall (1996) suggested that the fitness constituent weights be scaled so that their total is unity.

The total fitness of any order schedule is hence given by

$$\phi = \tfrac{1}{5}\sum_{f=1}^{5} w_f \phi_f \text{ such that } \sum_{f=1}^{5} w_f = 1 \tag{10}$$

The total is multiplied by a factor of 100 so that

$$\Phi = 100\phi \tag{11}$$

represents the "pseudo-percentage" of the theoretical maximum achievable value.

CASE STUDY

The case study involved scheduling all orders requested to start some time during the 20th of February 1997 for IPM planning Area B of G–MW's Central Goulburn irrigation network in Victoria, Australia. This area, outlined in bold in Fig. 8, consists of the No. 4 and 5 Channels (solid lines), and the Ardmona, Undera, and Rodney Drains (dashed lines), all shown in Fig. 8. Also included is a small part of the mid-section of the No. 6 Channel (solid line in Fig. 8), and Channel 5/A (not shown in Fig. 8) which runs alongside the Tatura–Undera Road (shown in Fig. 8) south of Tatura towards Murchison (these three towns are indicated by circles in Fig. 8). Water stored in Lake Eildon (13,840 ha in area and 3,390,000 ML in volume at full supply level, and located to the south of the area illustrated in Fig. 8) is released via the Goulburn River into a secondary 25,000 ML capacity storage at Goulburn Weir (a 212 m structure, and shown at the bottom of Fig. 8). Water controlled using this weir feeds this part of the channel system, which in general flows south to north, or "upwards" in Fig. 8. The drains are used to collect excess run-off, and irrigation water is also re-drawn from these for irrigator use.

personnel to contribute to the flow rate standard deviation and the control structure regulations fitness constituents ϕ_3 and ϕ_4. The union of these two sets, determined by unique control structure identification numbers in the IPM database, resulted in a total of 39 time-series that were required to be generated by IPM for the GA to analyse.

The Water Authority Parameters

The cut-off between "short" and "long" duration orders associated with the order shift fitness factor $\phi_{1a,o}$ was chosen by G–MW to be 18 hours. The lower-bounds for determination of control structure time-series from flow rate time-series, associated with the control structure regulations fitness $\phi_{4,c}$, were chosen to be 2.0 hours and 5.0% of the channel capacity at each control structure. The fitness constituents for the control structure fitnesses (ϕ_2, ϕ_3, and ϕ_4) were simply averaged, i.e., $w_{2,c} = 1/M$ for $c = 1,2,\ldots,M$ and $w_{3|4,c} = 1/V$ for $c = 1,2,\ldots,V$. Although other weightings are catered for in the software, in this instance $w_{2,c}$ was set to 1.0, for the M control structures at which a capacity (ML/day) is set and to 0.0 otherwise. Similarly, $w_{3,c}$ and $w_{4,c}$ were set to 1.0, for the V control structures which were chosen to contribute, and to 0.0 otherwise. All three sets were then scaled appropriately, viz. the relationships stated above.

The GA Parameters

The population size was kept constant at $P_0 = 150$ strings. Tournament selection was used to determine parent strings, and uniform crossover with probability $p_u = 0.5$ was performed with probability $p_c = 0.9$ (Goldberg 1992). Mutation was performed with per string probability $p_m = 0.05$, wherein a randomly selected order shift value was replaced by a randomly

Also implemented was a procedure by which, at each generation, the members of the population were forced to be unique. This was achieved by replacing each duplicated string with one that was generated randomly. This process was repeated until the proposed replacement was, in fact, different to all other population members. Such an *immigration* operation serves the same purposes as mutation: replacing lost genetic material and resisting premature convergence by increasing population diversity (Bean et al. 1995; Chipperfield 1998).

Please replace this line with Figure 10, or place Figure 10 near here

Fig. 10 illustrates the convergence of total fitness, Φ, towards the maximum found by the GA over all 100 generations, the GA's best solution (83.71). The fitness for the GA's best solution, the irrigators' requests, and the planner's schedule are indicated, respectively, by the upper, middle, and lower horizontal straight lines. Note that the line for the GA's best solution is extended across the entire horizontal axis for illustrative purposes only. The maximum, median, and average fitness of solutions found by the GA as the generations evolved are also indicated in Fig. 10, by the fine-dashed, coarse-dashed, and solid curves, respectively. The maximum GA fitness approached the GA's best solution as the generation number increased. In general, the average and median fitness approached the maximum fitness as the diversity in the population decreased with increasing generations. After approximately 80 generations, however, the average (but not the median) fitness decreased as the replacement of non-unique strings by random generation contaminated the population with relatively unfit schedules. The maximum GA fitness in the initial generation is similar to the fitness of the irrigators' requests as a result of the acceleration bias (towards zero shifts) discussed above. Since this bias was applied to the otherwise randomly generated initial population, each solution in that population had at least half of the orders shifted by zero hours, i.e.

The unequal order shift ranges indicate a possibly unfair test, in that the planners historically tend not to shift orders much more than a day (illustrated here by the –25 to 24 hours range for the planner's schedule) whereas the GA evaluates all schedules in which order shifts keep the starting times within the planning period (illustrated here by the –34 to 41 hours range for the GA's best solution). The mean of the absolute order shift values (2.12 hours for the planner's schedule and 2.85 hours for the GA's best solution) might, at first glance, indicate that, on average, the planner's schedule would result in marginally less disruption to irrigators' requests compared to the GA's best solution. The mean of the moduli of the non-zero order shift values (5.14 hours for the planner's schedule and 17.00 hours for the GA's best solution), however, offer a more indicative measure of mean disruption. The high peaks in frequency at zero shift indicate that most orders were scheduled as requested (98 for the planner's schedule and 139 for the GA's best solution). The GA's best solution clearly shifted fewer orders, but orders that were moved were shifted further, compared to the planner's schedule.

Please replace this line with Table 2, or place Table 2 near here

Table 2 illustrates some clear trends in the fitness constituents and the total fitness of the schedules for the irrigator's requests, the planner's schedule, and the GA's best solution. Shown also are the fitness values for the "optimum" solution: a theoretical (since it is, in practice, impossible to obtain) schedule in which the maximum fitness is scored for all constituents. The weights used for the constituents are also indicated by these "optimum" values.

The irrigators' requests corresponded to no orders being shifted and hence, by definition, scored the maximum practical value for orders shifts fitness constituent ϕ_1. This is less than the "optimum" value only because not all irrigators gave at least the minimum required number of days small number of orders by a small amount and maximized customer service, while keeping control structures time-series mostly under capacity, mainly of low standard deviation, and representing a small number of operations. The irrigators' requests shifted no orders (by definition) and maximized customer service (also by definition), but resulted in many control structures time-series being over capacity, of medium standard deviation, and representing a medium number of operations. The planner's schedule shifted a large number of orders by varying amounts and provided poor customer service, while resulting in some control structures time-series being over capacity, with high standard deviation, and representing a large number of operations.

Please replace this line with Table 3, or place Table 3 near here

Table 3 lists statistics that are direct measures of the optimization guidelines encapsulated in the five fitness constituents:

1. Mean absolute order shifts (in hours) [related to ϕ_1].
2. Mean relative channel capacities exceedance (as percentage of channel capacities, where the percentage is taken first, and all control structures contribute to the average even if they don't exceed capacity) [related to ϕ_2].
3. Mean standard deviations of flow rate time-series for all control structures (in ML/day) [related to ϕ_3].
4. Mean control structure regulations count [related to ϕ_4].
5. Orders shifted (as a percentage of total orders) [related to ϕ_5].

For the orders shifts measure the irrigators' requests schedule results in the best value (zero, by definition) because the schedule shifted no orders. The GA's best solution results in the worst because it shifted more orders than the planner's schedule.

Figs. 12 to 14 illustrate flow rate time-series, over the planning period, at three control structures, whose locations are shown in Fig. 8, of the 39 structures at which time-series were calculated. These are shown for the irrigators' requests (top), the planner's schedule (middle), and the GA's best solution (bottom). The solid curves indicate the flow rate time-series and the finely dashed curves indicate the control structure time-series. The coarse horizontal line indicates channel capacity, and the vertical axis at left (time = 0 hrs) and the solid vertical line at right (time = 47 hours) indicate the "window" within which the three fitness constituents (ϕ_2, ϕ_3, and ϕ_4) associated with the time-series were calculated.

Please replace this line with Figure 12, or place Figure 12 near here

Fig. 12 illustrates the time-series at a control structure, RN.362, for which the three schedules at some stage resulted in capacity exceedance. The exceedance within the first five hours of the planning period cannot be avoided because it corresponds to the completion of orders from the previous planning period. Otherwise the irrigator's requests and the GA's best solution are both under capacity over a greater period than the planner's schedule. The control structure time-series for the GA's best solution requires fewer control structure regulations (5) than does that for the planner's schedule (6), which in turn requires fewer than that for the irrigator's requests (7).

Please replace this line with Figure 13, or place Figure 13 near here

Fig. 13 illustrates the time-series for a control structure, RN.140, that was requested by the irrigators to go over capacity, was not brought under capacity by the planner's schedule, but was brought under capacity by the GA's best solution. The GA's best solution not only minimizes the water volume over capacity, but also minimizes the number of required control structure regulations.

while satisfying a set of constraints. A methodology for applying GA techniques to the optimal scheduling of irrigation orders in real-world open-channel networks was implemented.

Integration of the GA software module with the IPM water ordering, planning, and management software—by developing an appropriate application programming interface (API) layer—enabled the use of an Australian water authority's real-world database to generate irrigation schedules. These have illustrated the ability of GA techniques to identify good irrigation order schedules. The water authority's irrigation planners could be greatly assisted by use of the GA optimization, although they would still have the capacity to override suggestions for shifting an irrigation order backward or forward in time. The intention is that these schedules will be reviewed and finalized by irrigation planners, and then implemented by field operators, if and when the GA approach is fully adopted.

In combination with a suitable GUI, multi-objective GAs such as that described here can become a powerful, and possibly interactive, decision support tool, allowing a decision maker to learn about the problem before committing to a final decision (Fonesca and Fleming 1998). In this particular instance the interface could be built upon that existing in the IPM software, the decisions would be made by the G–MW planners after the GA has presented a collection—most likely derived using different constituent weighting schemes—of alternative "optimal" schedules, and the planners would select between members of the collection on the basis of their expert knowledge of the problem.

The results of this application of GA technology to flow management of open-channel gravity systems have shown that the technology can efficiently schedule irrigation order requests.

APPENDIX I. CONTROL STRUCTURE REGULATIONS FROM FLOW RATE VARIATIONS

With 2.0 hours and 5.0 ML/day as the smallest time change and flow rate change, respectively, allowed between regulations, the determination of the control structure time-series from the flow rate time-series is illustrated in Fig. 15.

Please replace this line with Figure 15, or place Figure 15 near here

The flow rate time-series [Fig. 15(a)] has its first step $[1.0 \text{ hours} \times 5.0 \text{ ML/day}]$, at time $t = 6$ hours, merged with its second $[1.0 \times 12.0]$, at time $t = 8$ hours, to give an intermediate time-series [Fig. 15(b)], because the first step was of dimensions less than the specified lower-bounds $[2.0 \times 5.0]$.

The new step form by this combination, at time $t = 6$ hours, has dimensions $[2.0 \times 11.0]$ at least as great as the lower-bounds and is thus counted as a control structure regulation.

The next step $[5.0 \times 11.0]$, at time $t = 8$ hours, has dimensions greater than the lower-bounds, is thus counted, and takes the control structure regulations total to 2.

The next step $[2.0 \times -12.0]$, at time $t = 13$ hours, has dimensions at least as great as the lower-bounds, is also counted, and gives a control structure regulations total of 3.

The next step $[1.0 \times 14.0]$, at time $t = 15$ hours, has dimensions less than the lower-bounds and is thus merged with the following step $[12.0 \times -5.0]$, at time $t = 16$ hours, to give another intermediate time-series [Fig. 15(c)].

APPENDIX II. REFERENCES

Bean, J. C., Hadj-Alouane, A. B., and Norman, B. A. (1995). "Robust encodings of OR problems for genetic algorithms." *Proc., 1st International Joint Workshop on Artificial Intelligence and Operations Research*, University of Oregon, Timberline, Oregon.

Chipperfield, A. (1998). "Introduction to genetic algorithms." *Genetic Algorithms in Engineering Systems*, A. M. S. Zalzala and P. J. Fleming, eds., number 55 in IEEE Control Engrg. Series, The Institution of Electrical Engineers, chapter 1, 1–45.

Dandy, G. C., Simpson, A. R., and Murphy, L. J. (1996). "An improved genetic algorithm for pipe network optimization." *Water Resources Research*, American Geophysical Union, 32(2), 449–458.

De Jong, K. A. (1985). "Genetic algorithms: A 10 year perspective." *Proc., 1st International Conf. on Genetic Algorithms and Their Applications*, J. J. Grefenstette, ed., Carnegie–Mellon University, Pittsburgh, Pennsylvania, 169–177.

Holland, J. H. (1992). *Adaptation in Natural and Artificial Systems: An Introductory Analysis with Applications to Biology, Control, and Artificial Intelligence.* The MIT Press, Cambridge, Massachusetts, Branford Books edition.

Goldberg, D. E. (1989). *Genetic Algorithms in Search, Optimization, and Machine Learning.* Addison–Wesley, Reading, Massachusetts.

"Goulburn–Murray Water." (2000). <http://www.g-mwater.com.au> (Jan. 31, 2000).

APPENDIX III. NOTATION

The following symbols are used in this paper:

c control structure number;

G number of generations;

M number of structures considering channel capacity exceedance;

O number of irrigation orders;

o irrigation order number;

P_0 (initial) population size;

p_c crossover probability;

p_m mutation probability;

p_u uniform crossover probability;

V number of structures considering time-series flow rate variations;

w fitness constituent weight;

$|\cdots|$ the absolute value of its argument;

Φ overall fitness "percentage"; and

ϕ fitness value.

Subscripts

The following subscripts are used in this paper:

c control structure number;

f fitness constituent;

o irrigation order number;

1 orders shifts;

Order Number (o)	1	2	3	4	5	…	…	166	167
Order Shift (hours)	+2	0	-3	+24	+6	…	…	-24	+12

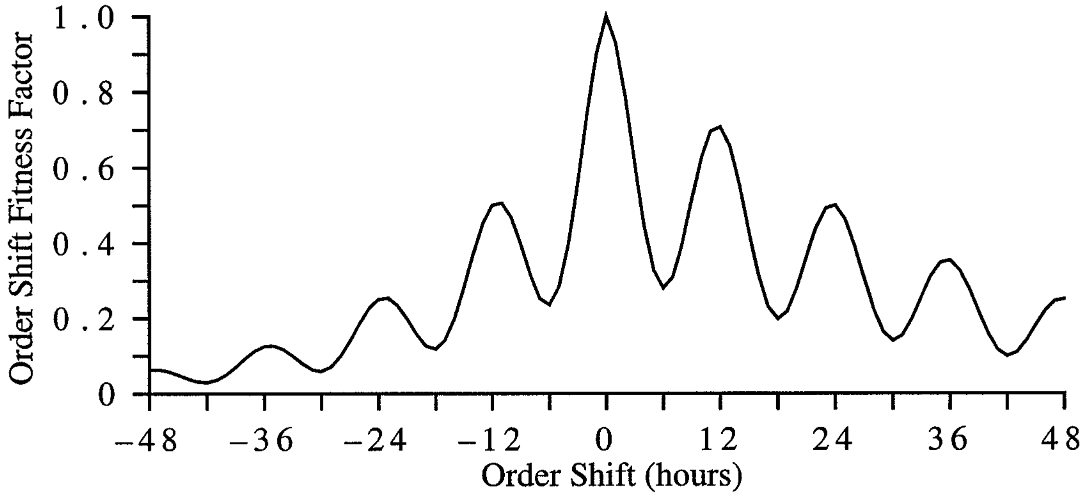

Order Shift Fitness Factor
1.0
0.8
0.6
0.4
0.2
0
−48
−36
−24
−12
0
12
24
36
48
Order Shift (hours)

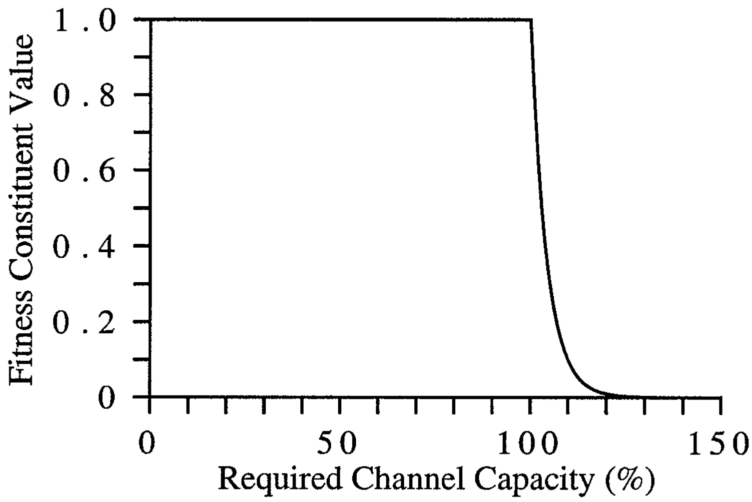
1.0
0.8
0.6
0.4
0.2
0
Fitness Constituent Value
0
50
100
150
Required Channel Capacity (%)

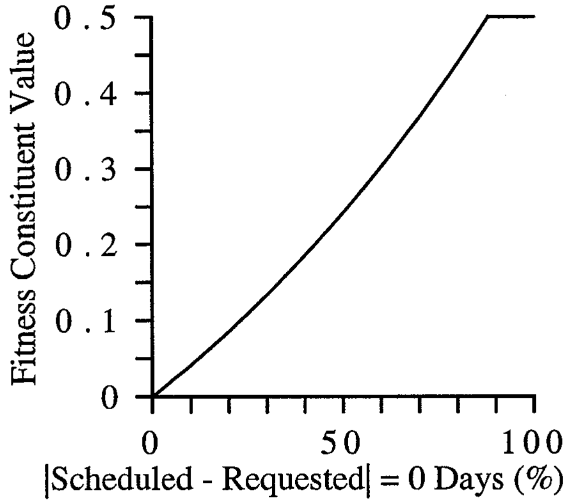
0.5
0.4
0.3
0.2
0.1
0
Fitness Constituent Value
0
50
100
|Scheduled - Requested| = 0 Days (%)

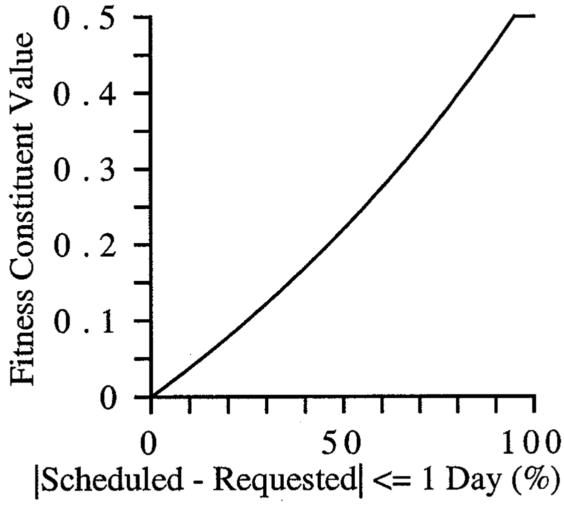
0.5
0.4
0.3
0.2
0.1
0
Fitness Constituent Value
0
50
100
|Scheduled - Requested| <= 1 Day (%)

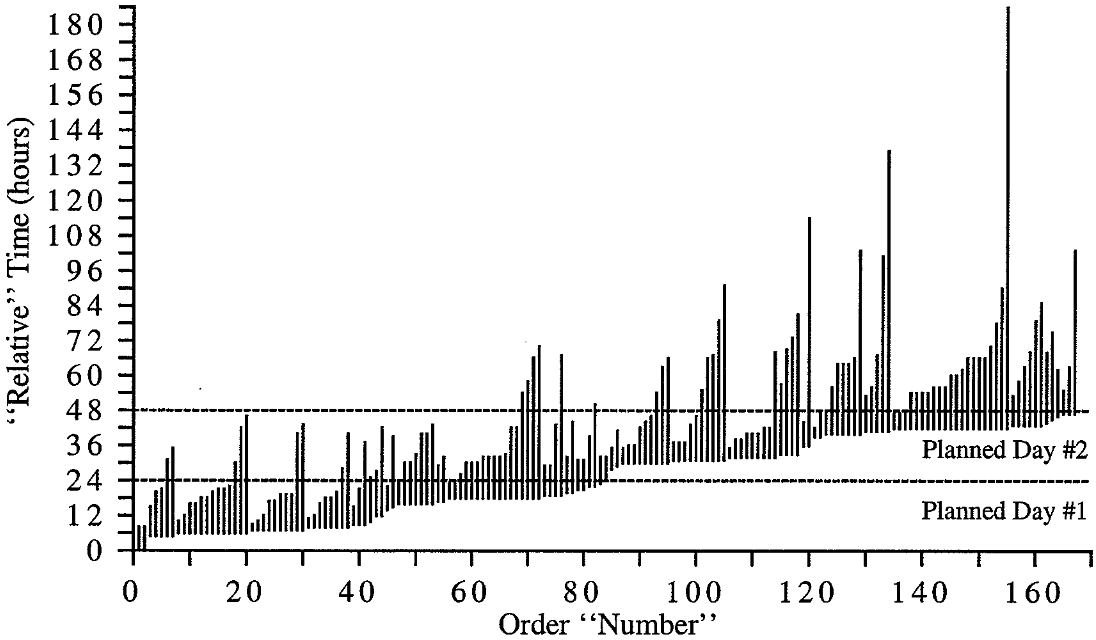

180
168
156
144
132
120
108
96
84
72
60
48
36
24
12
0
''Relative'' Time (hours)
Planned Day #2
Planned Day #1
0
20
40
60
80
100
120
140
160
Order ''Number''

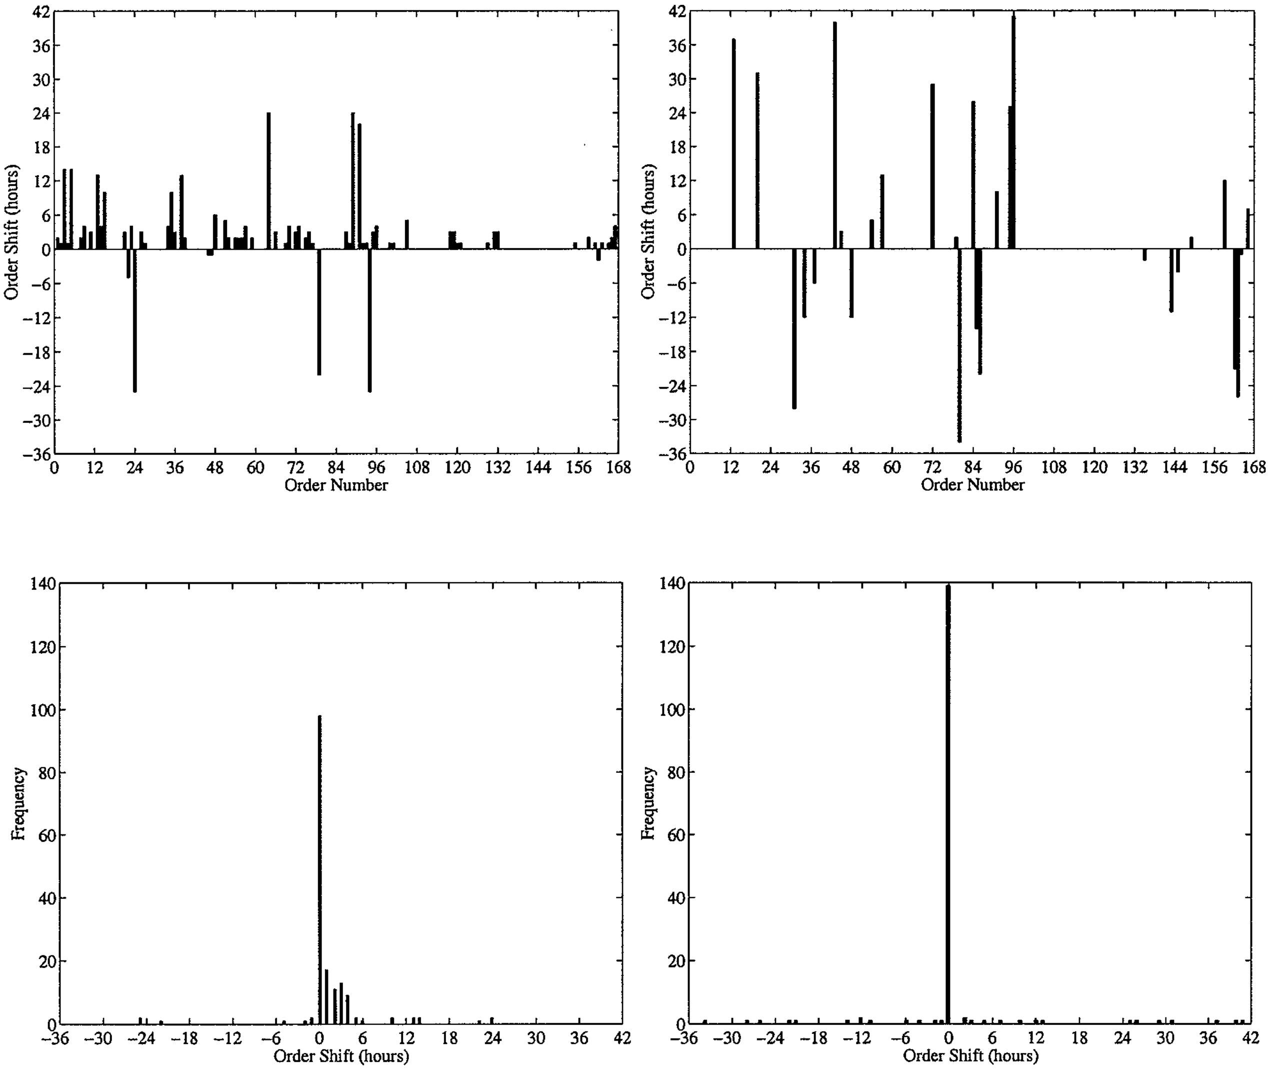
Order Shift (hours)
Order Number
Order Shift (hours)
Order Number
Frequency
Order Shift (hours)
Frequency
Order Shift (hours)

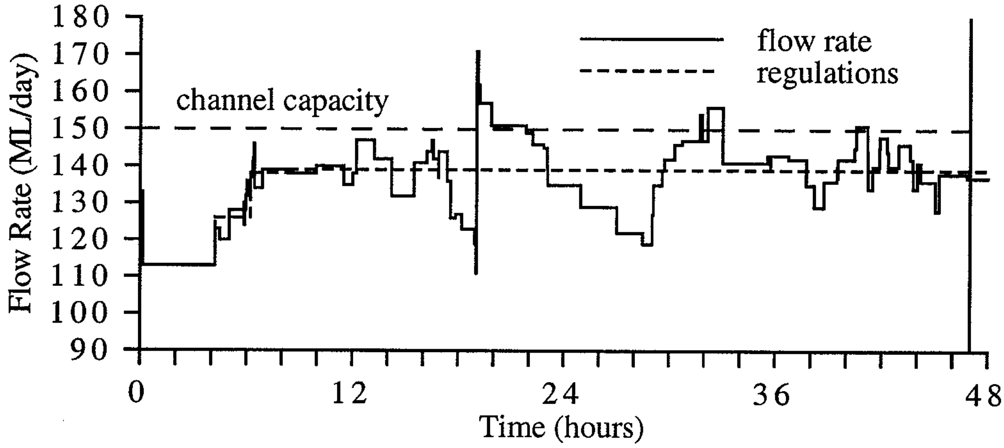

Flow Rate (ML/day)
180
170
160
150
140
130
120
110
100
90
channel capacity
flow rate
regulations
0
12
24
36
48
Time (hours)

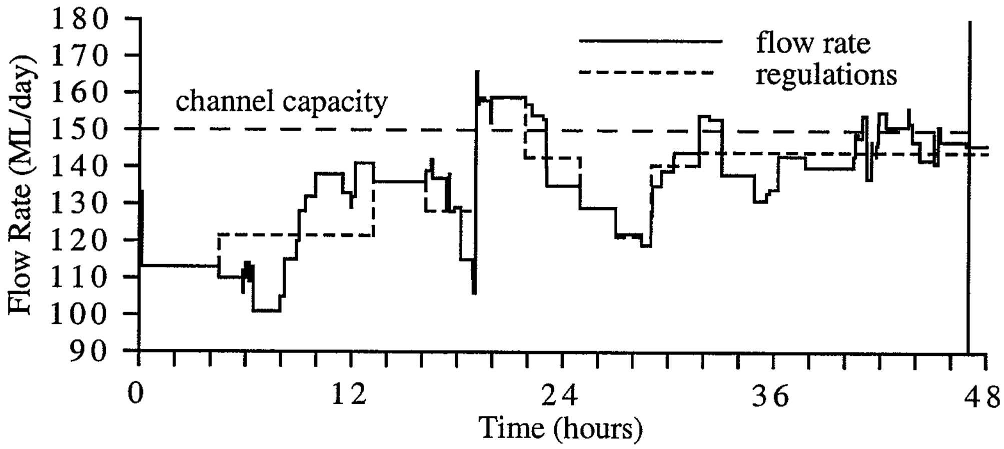

Flow Rate (ML/day)
180
170
160
150
140
130
120
110
100
90
channel capacity
flow rate
regulations
0
12
24
36
48
Time (hours)

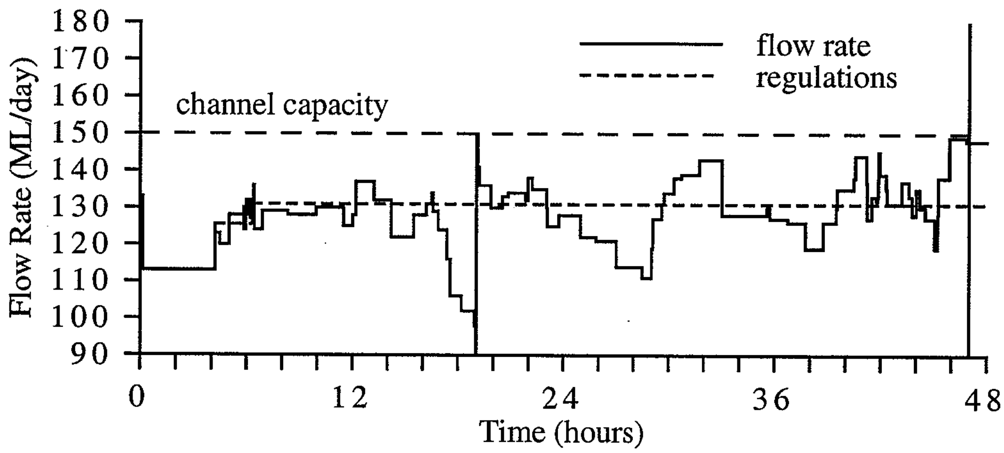

Flow Rate (ML/day)
180
170
160
150
140
130
120
110
100
90
channel capacity
flow rate
regulations
0
12
24
36
48
Time (hours)

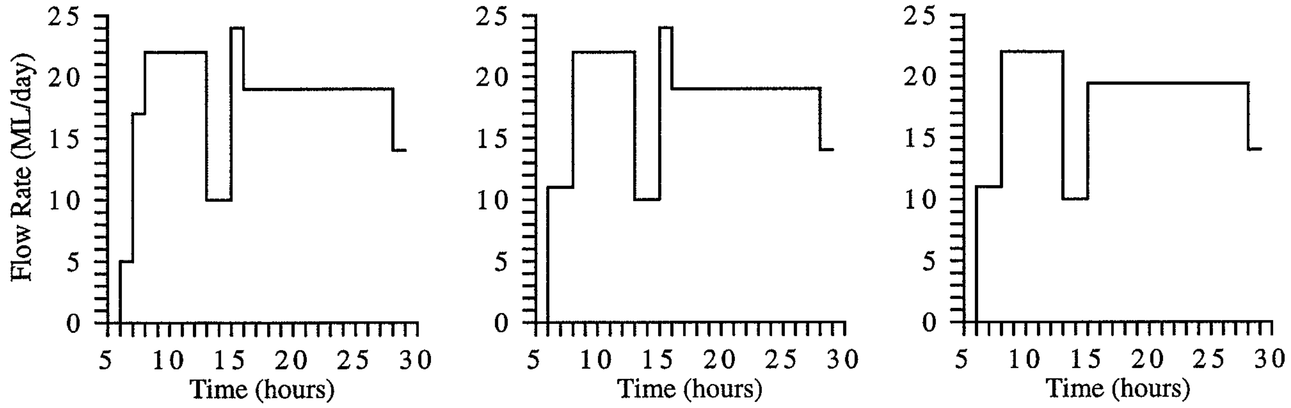

Flow Rate (ML/day)
25
20
15
10
5
0
5 10 15 20 25 30
Time (hours)
25
20
15
10
5
0
5 10 15 20 25 30
Time (hours)
25
20
15
10
5
0
5 10 15 20 25 30
Time (hours)

TABLE 2. Fitness Constituent Values for the Case Study

Irrigation Schedule (which)	Orders Shifts ($100\phi_1$)	Channel Capacities ($100\phi_2$)	Standard Deviations ($100\phi_3$)	Control Structures ($100\phi_4$)	Customer Service ($100\phi_5$)	Total Fitness (Φ)
(1)	(2)	(3)	(4)	(5)	(6)	(7)
"Optimum"	12.50	50.00	12.50	12.50	12.50	100.00
Irrigators'	**10.91**	*40.20*	*6.25*	6.25	12.50	76.11
Planner's	*1.07*	41.71	6.36	*6.09*	*8.15*	*63.38*
GA's Best	9.95	**46.21**	**7.25**	**7.79**	**12.50**	**83.71**

The *minimum*/**maximum** values in each column (excluding the "optimum" values) are in *italic*/**bold** text.

FIG. 1. A Typical GA String for a Scheduling Problem with 167 Orders

FIG. 2. The Order Shift Fitness Factor, $\phi_{1a,o}$, for "Short" Duration Orders

FIG. 3. The Order Shift Fitness Factor, $\phi_{1a,o}$, for "Long" Duration Orders

FIG. 4. The Order Notice Fitness Factor: $\phi_{1b,o}$

FIG. 5. The Channel Capacity Exceedance Fitness Constituent: $\phi_{2,c}$

FIG. 6. The Flow Rate Variations Fitness Constituents: $\phi_{3,c}$ and $\phi_{4,c}$

(a) ϕ_{5a} **(b) ϕ_{5b}**

FIG. 7. The Customer Service Fitness Measures for Orders Scheduled:
(a) on the Same Date as the Requested Day (ϕ_{5a});
(b) between the Nearest Dates of the Requested Day (ϕ_{5b})

FIG. 8. Part of the Central Goulburn Irrigation District

(a) flow rate **(b) "intermediate"** **(c) control structure**

FIG. 15. Control Structure Time-Series from Flow Rate Time-Series

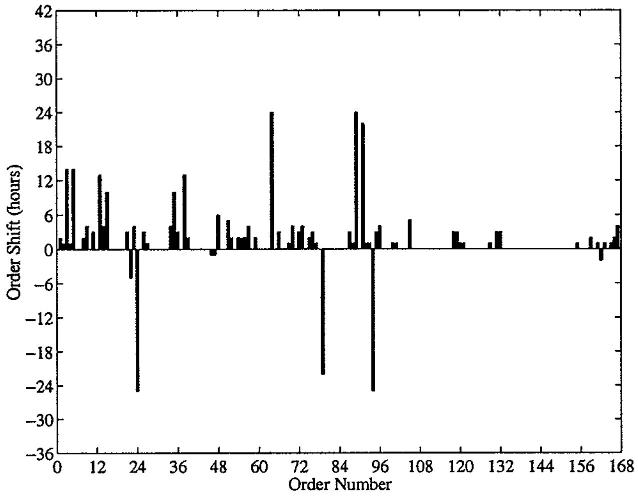

(a) **The planner's schedule**

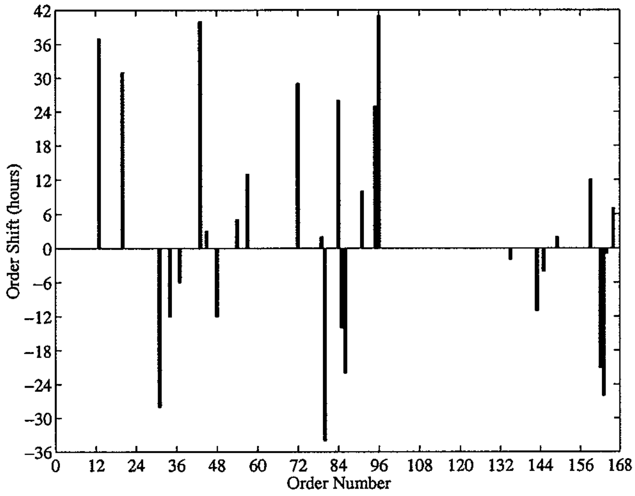

(b) **The GA's best solution**

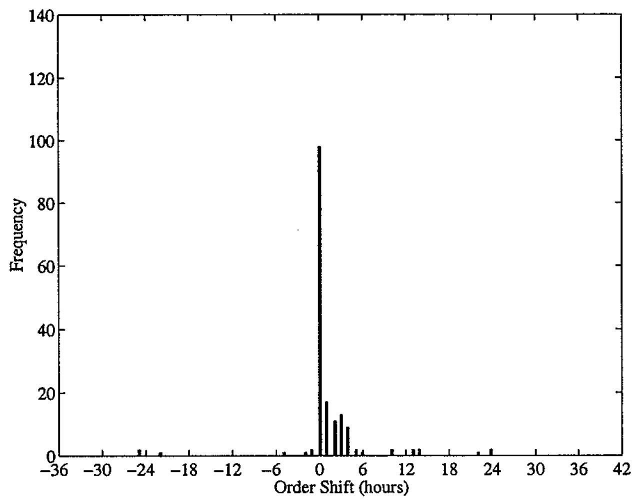

(c) **The planner's schedule**

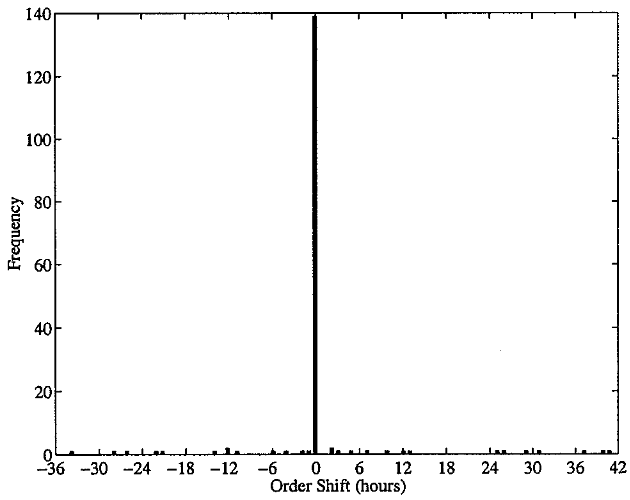

(d) **The GA's best solution**

FIG. 11. Values (top) and Frequencies (bottom) for the Order Shifts, as Scheduled by the Planner (left) and the GA (right), in the Case Study

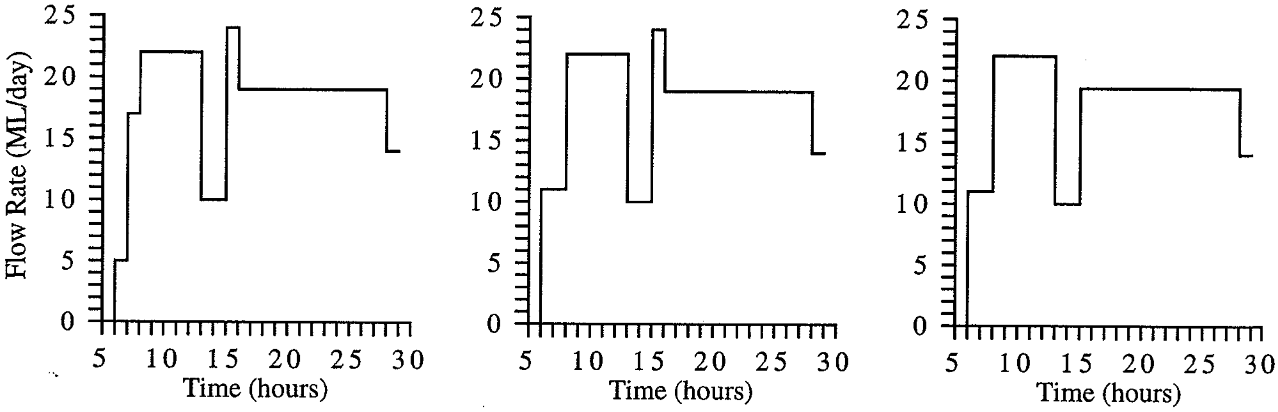

(a) flow rate **(b) "intermediate"** **(c) control structure**

FIG. 15. Control Structure Time-Series from Flow Rate Time-Series