

Incremental Maximum Flows for Fast Envelope Computation

Nicola Muscettola

NASA Ames Research Center
Moffett Field, CA 94035
mus@email.arc.nasa.gov

Abstract

Resource envelopes provide the tightest exact bounds on the resource consumption and production caused by all possible instantiations of a temporally flexible plan. We present a new algorithm that computes an envelope in $O(\text{Maxflow}(\mathbf{n}, \mathbf{m}, \mathbf{U}))$ where \mathbf{n} , \mathbf{m} and \mathbf{U} measure the size of the flexible plan. This is an $O(\mathbf{n})$ improvement on the best envelope algorithm known so far and makes envelopes more amenable to practical use in scheduling algorithms. The reduction in complexity depends on the fact that when the algorithm computes the constant segment \mathbf{i} of the envelope it makes full reuse of the maximum flow that was computed in order to obtain segment $\mathbf{i}-1$.

Resource Envelopes

The execution of plans greatly benefits from temporal flexibility. Fixed-time plans are brittle and may require extensive replanning due to execution uncertainty. Moreover, when plans must deal with uncontrollable exogenous events (Morris et al., 2001) temporal flexibility cannot be avoided. However, effective algorithms to build temporally flexible plans are rare, especially when activities produce or consume variable amounts of resource capacity. A major obstacle is the difficulty of assessing the resource needs across all possible plan executions. Methods are available to compute resource consumption bounds (Laborie, 2001; Muscettola, 2002). In particular, (Muscettola, 2002) proposes a polynomial algorithm to compute a *resource envelope*, the tightest of these bounds. By being the tightest possible, resource envelopes can potentially save an exponential amount of search (through early backtracking and solution detection) when compared to using looser bounds. Also, methods that compute resource envelopes identify maximally matched sets of resource consumer/producers that balance each other for any plan execution. This and other structural information could be crucial in minimizing the search space and suggesting effective scheduling heuristics, potentially enabling new classes of highly efficient schedulers.

However, preliminary studies on schedulers using envelopes appear not to show advantages with respect to more traditional heuristic methods based on fixed-time resource profiles (Pollicella et al., 2003). When compared to traditional fixed-time profiling methods, it is critical to balance the increased computation cost with the extraction of more structural problem information from the envelope

than backtrack/termination tests and maximum resource contention intervals. Making the trade-off advantageous requires two complementary approaches. The first reduces the cost of computing an envelope; the second devises new envelope analysis methods to extract useful heuristics.

In this paper we address the problem of cost reduction. The fastest known resource envelope algorithm (Muscettola, 2002) computes all piecewise-constant segments of the envelope through as many as $2\mathbf{n}$ stages, where \mathbf{n} is the number of events in the flexible plan. Each stage computes a maximum flow using some maximum flow algorithm. The worst case complexity is $O(\mathbf{n} \text{Maxflow}(\mathbf{n}, \mathbf{m}, \mathbf{U}))$ where \mathbf{m} is the number of temporal constraints between activities in the plan, \mathbf{U} is the maximum level of resource production or consumption at some activity, and $\text{Maxflow}(\mathbf{n}, \mathbf{m}, \mathbf{U})$ is the asymptotic cost of the maximum flow algorithm.

This staged method, however, can be significantly improved since at each stage it recomputes the needed maximum flow completely from scratch. This suggests using an incremental flow method. Starting from the maximum flow at one stage, this method computes the maximum flow at the next stage by minimally reducing flow when deleting nodes and edges, and by minimally increasing flow when adding new nodes and edges (Kumar, 2003). However, without appropriately ordering flow reductions and increases, the asymptotic complexity may not improve (at it appears to be the case in (Kumar, 2003)). In this paper we introduce an incremental method that provably computes an envelope in $O(\text{Maxflow}(\mathbf{n}, \mathbf{m}, \mathbf{U}))$ for a large class of maximum flow algorithms. This reduction of complexity is significant. Experimental analysis has shown that the practical cost of maximum flow is usually as low as $O(\mathbf{n}^{1.5})$ (Ahuja et al., 1993). This compares well with $O(\mathbf{n} \log \mathbf{n})$, the cost of building resource profiles for fixed time schedules. This paper is organized as follows. We first give a succinct introduction to the resource envelope problem and the staged envelope algorithm in (Muscettola, 2002). Next we present the new incremental algorithm identifying all sources of performance improvements. We then prove the complexity result and conclude by discussing future work.

Staged Computation of Envelopes

In this section we introduce the essential information on the envelope problem and the staged algorithm that solves it. For a complete discussion, see (Muscettola, 2002).

Figure 1 shows an activity network with resource allocations. The network has two time variables per activity, a start event and an end event (e.g., e_{1s} and e_{1e} for activity A_1), a non-negative flexible activity duration link (e.g., $[2, 5]$ for activity A_1), and flexible separation links between events (e.g., $[0, 4]$ from e_{3e} to e_{4s}). Two additional events T_s and T_e define a time horizon within which all events occur.

Time origin, events and links constitute a Simple Temporal Network. To describe resource production and consumption each event has also an *allocation value* (e.g., r_{31} for event e_{3s}), a numeric weight that represents the amount of resource allocated when the event occurs. We will assume that all allocations refer to a single, multi-capacity resource. The extension to multiple resources is straightforward. If the allocation is negative an event e^- is a *consumer*, if it is positive e^+ is a *producer*. We assume that the temporal constraints are consistent which means that for any pair of events the shortest path $|e_1e_2|$ from e_1 to e_2 is well defined. Each event e can occur within its time bound, between the earliest time $et(e) = -|eT_s|$ and the latest time $lt(e) = |T_e e|$. The triangular inequality $|e_1e_3| \leq |e_1e_2| + |e_2e_3|$ holds for any three events e_1 , e_2 and e_3 .




Figure 1: An activity network with resource allocations

The *anti-precedence graph*, \mathbf{Aprec} , is a graph containing a path between any two events e_1 and e_2 such that $|e_1e_2| \leq 0$. Figure 2 depicts an anti-precedence graph of the network in Figure 1 with each event labeled with its time bound and resource allocation.

We can now formally define a resource envelope. For any subset of events \mathbf{A} , the *resource level increment* of \mathbf{A} is $\Delta(\mathbf{A}) = 0$ if $\mathbf{A} = \emptyset$, and $\Delta(\mathbf{A}) = \sum_{e \in \mathbf{A}} c(e)$ if $\mathbf{A} \neq \emptyset$. If \mathbf{S} is the set of all possible consistent time instantiations for all events and \mathbf{t} is a time within the time horizon, the resource level at time \mathbf{t} for a specific time instantiation $s \in \mathbf{S}$ is $L_s(\mathbf{t}) = \Delta(\mathbf{E}_s(\mathbf{t}))$. Here $\mathbf{E}_s(\mathbf{t})$ is the set of events e which occur at or before \mathbf{t} in s . The *maximum resource envelope* is $L_{\max}(\mathbf{t}) = \max_{s \in \mathbf{S}} L_s(\mathbf{t})$ and the *minimum resource envelope* is $L_{\min}(\mathbf{t}) = \min_{s \in \mathbf{S}} L_s(\mathbf{t})$. Since L_{\min} can be computed with obvious term substitution on the method that computes L_{\max} , we only focus on L_{\max} .

To compute the resource envelope at time \mathbf{t} we partition all events into three sets depending on the position of their time bound relative to \mathbf{t} : 1) the *closed events* \mathbf{C}_t that must occur before or at \mathbf{t} , i.e., such that that $lt(e) \leq \mathbf{t}$; 2) the




Figure 2: A resource increment flow network

pending events \mathbf{R}_t that can occur before, at or after \mathbf{t} , i.e., such that $(e) \leq \mathbf{t} < lt(e)$; and 3) the *open events* \mathbf{O}_t that must occur strictly after \mathbf{t} , i.e., such that $et(e) > \mathbf{t}$.

Any resource level increment $L_s(\mathbf{t})$ will always include the contribution of all events in \mathbf{C}_t and none of those in \mathbf{O}_t but can include only some subset of events in \mathbf{R}_t , i.e., only those that are scheduled before \mathbf{t} in s . It is possible to show that this subset must be a *predecessor set* $\mathbf{P} \subseteq \mathbf{R}_t$ such that if $e \in \mathbf{P}$ and e' follows e in \mathbf{Aprec} , then $e' \in \mathbf{P}$. We call $\mathbf{P}_{\max}(\mathbf{R}_t)$ the (possibly empty) predecessor set with maximum non-negative resource level increment.

The fundamental result reported in (Muscatella, 2002) is that $L_{\max}(\mathbf{t})$ can be determined from the following equation.

$$\text{Equation 1: } L_{\max}(\mathbf{t}) = \Delta(\mathbf{C}_t) + \Delta(\mathbf{P}_{\max}(\mathbf{R}_t))$$




Figure 3: Anti-precedence graph with time/resource usage

The cost of computing an envelope depends on the cost of computing $\mathbf{P}_{\max}(\mathbf{R}_t)$. We can compute $\mathbf{P}_{\max}(\mathbf{R}_t)$ by solving a maximum flow problem on an auxiliary flow network $\mathbf{F}(\mathbf{R}_t)$, the *resource increment flow network* for \mathbf{R}_t .

The formal definition of a resource increment flow network can be found in (Muscatella, 2002). As an example, Figure 2 gives $\mathbf{F}(\mathbf{R}_t)$ for the activity network in Figure 1. The network has a node for each event in \mathbf{R}_t , an infinite capacity flow edge between two events for each edge in \mathbf{Aprec} (see Figure 2), an edge from the source σ to a producer with capacity equal to the producer's allocation, and an edge from a consumer to the sink τ with capacity equal to the opposite of the consumer's allocation.

A complete discussion of maximum flow algorithms can be found in (Cormen, Leiserson and Rivest, 1990). Here we only highlight a few concepts that we will use in the following. A flow is a function $f(e_1, e_2)$ of pair of events in $\mathbf{F}(\mathbf{R}_t)$ that is skew-symmetric, i.e., $f(e_2, e_1) = -f(e_1, e_2)$, has

a value no greater than the capacity of edge $e_1 \rightarrow e_2$ (assuming capacity zero if the edge is not in $F(R_t)$), and is balanced, i.e., the sum of all flows entering an event must be zero. A pre-flow is a function defined similarly but that relaxes the balance constraint by allowing the sum of preflows entering a node to be positive. The total network flow is defined as $\sum_{e \in R_t} f(\sigma, e) = \sum_{e \in R_t} f(e, \tau)$. The maximum flow of a network is a flow function f_{\max} such that the total network flow is maximum.

A fundamental concept in the theory of flows is the *residual network*, a graph with an edge for each pair of nodes in $F(R_t)$ with positive *residual capacity*, i.e., the difference between edge capacity and flow. Each residual network edge has capacity equal to the residual capacity. An *augmenting path* is a path connecting σ to τ in the residual network. The existence of an augmenting path indicates that additional flow can be pushed from σ to τ . Alternatively, the lack of an augmenting path indicates that a flow is maximum.

We can compute $P_{\max}(R_t)$ according to the next theorem.

Theorem 2: (Muscettola 2002) $P_{\max}(R_t)$ is the (possibly empty) set of events that are reachable from the source σ in the residual network of some f_{\max} of $F(R_t)$.

From Equation 1 and Theorem 2 (Muscettola, 2002) derives a staged envelope algorithm as follows. Consider the $2n$ times t_i corresponding to the earliest and latest times for all events. Since the envelope level can only change at one of these times, the algorithm computes a different level for each of them. At a particular t_i the algorithm determines its closed and pending event sets C_i and R_i , builds $F(R_i)$, solves a maximum flow over it, determines $P_{\max}(R_i)$ according to Theorem 2, and computes $L_{\max}(t_i)$ according to Equation 1. It is easy to see that the worst-case time complexity of this algorithm is $O(n \text{ Maxflow}(n, m, U))$.

Incremental Computation of Envelopes

In the previously described staged envelope algorithm flows are recomputed from scratch for each $F(R_i)$. Assuming that the times t_i are sorted in increasing order, in To reduce cost, we can try and reuse as much as possible of the maximum flow computation performed on $F(R_{i-1})$. At time t_i the set of pending events can undergo two modifications. First, the events $\delta C_i = R_{i-1} - R_i$ move from R_{i-1} to C_i . These are events e such that $t_i = \text{lt}(e)$. Second, the events $\delta R_i = R_i - R_{i-1}$ move from O_{i-1} to R_i . These are the events e such that $t_i = \text{et}(e)$. For example, consider the activity network in Figure 1 and the process through which R_3 is transformed into R_4 . This is described in Figure 4 where the grayed part of the network is deleted and the emphasized part of the network is added at time 4. In particular, we have $\delta C_4 = \{e_{1s}\}$ and $\delta R_4 = \{e_{2s}\}$. For completeness, we note that $F(\delta C_4)$ consists of node e_{1s} and edge $e_{1s} \rightarrow \tau$ while $F(\delta R_4)$ consists of node e_{2s} and edge $\sigma \rightarrow e_{2s}$. All other added and deleted edges are connectives between $F(R_4 - \delta C_4)$ and F

(δC_4) (edges $e_{1e} \rightarrow e_{1s}$ and $e_{3s} \rightarrow e_{1s}$) and between $F(\delta R_4)$ and $F(R_4 - \delta C_4)$ ($e_{2s} \rightarrow e_{1e}$).

The sets δC_i and δR_i satisfy the following fundamental properties.

Lemma 3: δC_i is a predecessor set contained in R_i . δR_i is the complement of predecessor set R_{i-1} in R_i .

Proof: We only give the proof for δC_i since the one for δR_i is analogous. Consider a pair of events $e_1 \in \delta C_i$ and $e_2 \in R_{i-1} - \delta C_i$. From the definition of δC_i we have $\text{lt}(e_1) = t_i$ and $\text{lt}(e_2) \geq t_i + 1$. From the triangular inequality $\text{lt}(e_2) \leq \text{lt}(e_1) + |e_1 e_2|$ we deduce $|e_1 e_2| \geq \text{lt}(e_2) - \text{lt}(e_1) \geq t_i + 1 - t_i = 1 > 0$. \square

Lemma 3 determines what flow edges are eliminated when δC_i is deleted and what are added when δR_i is added. In particular, we can only delete edges that enter events in δC_i or go from δC_i to τ . Similarly, we can only add edges that exit events in δR_i or go from σ to δR_i . Unlike previous proposals for incremental envelope calculation (Kumar, 2003), our methods relies on events and edges exiting and entering the current flow network in a well defined order. This is the primary key to reducing complexity in our




Figure 4: Incremental modification of a resource flow network

incremental envelope algorithm.

Directly related to Lemma 3 is the possibility of computing the maximum flow of $F(R_i)$ by incrementally modifying the flow of $F(R_{i-1})$, reusing both flow values and intermediate data structures across successive invocations of a maximum flow algorithm. We prove that our flow modification operators guarantee the optimality at each intermediate flow. Maintaining maximality of intermediate flows and reusing data structures across flows are keys to reduce complexity for different kinds of maximum flow algorithms.

A final factor is minimizing the size of each intermediate flow network. We will show that as soon as the weight of an intermediate P_{\max} is used in the envelope calculation, $F(P_{\max})$ and all of its connecting edges can be safely eliminated from further consideration. This reduces flow network size and additionally contributes to cost reduction.

Flow Modification Networks

We now define the operators needed for the incremental envelope algorithm. The philosophy of each operator is similar to that used by the flow augmentation method in

maximum flow theory. However, we use this method more generally not only to augment flow but also to shift flow around the network and to reduce flow. The general idea is the following. Given a flow network F and one of its maximum flows f , an operator first defines an auxiliary flow transformation network F_T , then finds one of its maximum flows f_T , and finally produces a flow $f_{\text{new}} = f + f_T$. Each F_T consists of selected edges in the residual network of F for f . Since the properties of flows are preserved when adding a flow of a residual network to the flow that originated the residual network, f_{new} is also a flow for network F .

Consider now the resource increment flow network $F(R_{i-1})$ at stage $i-1$ and assume that the set of new closed events at stage i δC_i is not empty. At stage i all events in δC_i and all of its incoming and outgoing edges will be deleted. This also means that any flow that at the end of stage $i-1$ enters δC_i will necessarily have to be zeroed, i.e., pushed back into $F(R_{i-1})$. The value of this flow is the sum of the residual capacities of all edges $e_1 \rightarrow e_2$ where $e_1 \in \delta C_i$ and $e_2 \in R_{i-1} - \delta C_i$. Once pushed back, this flow can follow two routes. The first reaches τ to fill up some non-saturated exiting edges of $F(R_{i-1} - \delta C_i)$. The second reverses all the way to σ because it cannot find any way to exit $F(R_{i-1} - \delta C_i)$ through its sink. We call this flow push-back operation a *flow contraction*. The first flow route corresponds to a *flow shift* and the second one is a *flow reduction*. For example, consider the network in Figure 4. Assume that at $t=3$ it is $f_{\max}(e_{1s}, \tau) = 4$, $f_{\max}(e_{1e}, \tau) = 1$ and $f_{\max}(e_{3e}, \tau) = 2$. At $t=3$ the elimination of e_{1s} requires pushing back 4 units of flow. However, note that three units can still reach τ by being shifted to $e_{1e} \rightarrow \tau$. Only one unit of flow needs to be pushed back to σ . If we did not shift (as in (Kumar, 2003)), three additional units of flow would have to be pushed again from σ to τ to ensure flow maximality.

Assume now that at stage i there is also a non-empty set δR_i of new pending events. Augmenting $F(R_{i-1} - \delta C_i)$ with the part of the resource increment flow network pertaining to δR_i yields $F(R_i)$. Assume now that $F(R_{i-1} - \delta C_i)$ is traversed by the flow resulting from flow contraction. Even if this flow is maximum for $F(R_{i-1} - \delta C_i)$, it may not be maximum for $F(R_i)$ since additional flow could be pushed through edges $\sigma \rightarrow e$ with $e \in \delta R_i$. We call this flow push-forward operation a *flow expansion*. If at every stage of flow contraction and flow expansion we guarantee flow maximality, we will obtain a maximum flow for $F(R_{(i)})$ that moves a minimal amount of flow.

Flow Contraction

Let us call $f_{\max,i-1}$ the maximum flow for $F(R_{i-1})$. In our discussion we ignore the structure of the flow sub-network for δC_i by using an auxiliary flow network \underline{F}_{i-1} that redirects all flow entering δC_i into the sink τ . Formally, to obtain \underline{F}_{i-1} we first delete from $F(R_i)$ all events in δC_i , together with all their incoming and outgoing flow edges. We then add an *auxiliary edge* $e_1 \rightarrow \tau$ for each set of component edges $e_1 \rightarrow e_2$ in $F(R_{i-1})$ such that $e_1 \in R_{i-1} - \delta C_i$ and $e_2 \in \delta C_i$. The capacity of the auxiliary edge $e_1 \rightarrow \tau$ is the sum of all component edge flows $f_{\max,i-1}(e_1, e_2)$. We call

$f_{\max,i-1}$ a function over the edges of \underline{F}_{i-1} where $f_{\max,i-1}(e_1, e_2)$ is equal to $f_{\max,i-1}(e_1, e_2)$ if $e_1 \rightarrow e_2$ also belongs to $F(R_{i-1})$, and $f_{\max,i-1}(e_1, \tau)$ is equal to the edge's capacity if $e_1 \rightarrow \tau$ is one of the auxiliary flow edges. It is easy to see that $f_{\max,i-1}$ is a maximum flow for \underline{F}_{i-1} . We call $\underline{\text{Res}}_{i-1}$ the residual network of \underline{F}_{i-1} for $f_{\max,i-1}$.

We define a *flow shift network* Shift_i as follows.

Flow shift network: Shift_i is a flow network with the same intermediate events of $\underline{\text{Res}}_{i-1}$. Shift_i has a flow edges $e_1 \rightarrow e_2$ equal to a corresponding one in $\underline{\text{Res}}_{i-1}$ if $e_1 \notin \{\sigma, \tau\}$ and $e_2 \neq \sigma$. Finally, for each edge $\tau \rightarrow e$ in $\underline{\text{Res}}_{i-1}$ such that $e \rightarrow \tau$ is an auxiliary flow edge in \underline{F}_{i-1} , Shift_i has a corresponding edge $\sigma \rightarrow e$ of the same capacity.

Let us now call $\text{Res}(\text{Shift}_i)$ the residual network of \underline{F}_{i-1} for $f' = f_{\max,i} + f_{\max,\text{shift},i}$. We define a *flow reduction network* Reduce_i as follows.

Flow reduction network: Reduce_i is a flow network with the same nodes as $\text{Res}(\text{Shift}_i)$ and edges $e_1 \rightarrow e_2$ identical to $\text{Res}(\text{Shift}_i)$ if one of the following three conditions is satisfied:

- 1) $e_2 \neq \tau$
- 2) $e_1 \neq \sigma$
- 3) $e_1 = \sigma$ and the edge $\sigma \rightarrow e_2$ in Shift_i originates from an auxiliary flow edge for \underline{F}_{i-1} .

Using Shift_i and Reduce_i , we define the **Flow_Contraction** operator needed by the incremental envelope algorithm.

Flow_Contraction($F(R_{i-1})$, $f_{\max,i-1}$, δC_i , Aprec):

- 1) Compute a maximum flow $f_{\max,\text{shift},i}$ for Shift_i ;
- 2) Compute a maximum flow $f_{\max,\text{red},i}$ for Reduce_i ;
- 3) Return $f_{\text{contr},i} = f_{\max,i} + f_{\max,\text{shift},i} + f_{\max,\text{red},i}$

We now prove that the operator keeps the flow maximum.

Lemma 4: The flow $f' = f_{\max,i} + f_{\max,\text{shift},i}$ is maximum for \underline{F}_{i-1} .

Proof: f' is a flow of \underline{F}_{i-1} . It is also maximum since by construction of Shift_i it is $f_{\max,\text{shift},i}(\sigma, e) = 0$. Therefore $f'(\sigma, e) = f_{\max,i-1}(\sigma, e)$ and therefore f' is also maximum for \underline{F}_{i-1} .

Lemma 5: $f_{\text{contr},i}$ is a flow for $F(R_{i-1} - \delta C_i)$.

Proof: $f_{\text{contr},i}$ is a flow for \underline{F}_{i-1} . For it to be a flow for $F(R_{i-1} - \delta C_i)$ it must be $f_{\text{contr},i}(e, \tau) = 0$ if $e \rightarrow \tau$ is an auxiliary edge. If it were $f_{\text{contr},i}(e, \tau) > 0$ for an auxiliary edge, by using the flow conservation constraint we could show that there must be a path from σ to τ , passing through $e \rightarrow \tau$, with all edges having positive flow. Therefore, there must be a flow-reducing path from τ to σ in the corresponding residual network. Such path is an augmenting path in the residual network of Reduce_i for flow $f_{\max,\text{red},i}$, which contradicts the maximality of $f_{\max,\text{red},i}$.

Theorem 6: $f_{\text{contr},i}$ is a maximum flow for $F(R_{i-1} - \delta C_i)$.

Proof: This is clearly true if $f_{\max, \text{red}, i}$ is a null flow since f^* is maximum. If $f_{\max, \text{red}, i}$ is not null, assume that $f_{\text{contr}, i}$ is not maximum. This yields an augmenting path from σ to τ in $F(\mathbf{R}_{i-1} - \delta C_i)$ for $f_{\text{contr}, i}$. Since $f_{\max, i}$ is optimal, such path could only have appeared after the computation of $f_{\max, \text{shift}, i}$. Since f^* is maximum for \mathbf{E}_{i-1} , there must be at least one edge $e_1 \rightarrow e_2$ on the augmenting path that does not belong to the residual network of \mathbf{E}_{i-1} for f^* while the suffix path from e_2 to τ has positive residual capacity in Shift_i for $f_{\max, \text{shift}, i}$. A positive residual for $e_1 \rightarrow e_2$ implies that flow reduction pushed flow in the opposite direction, i.e., $f_{\max, \text{red}, i}(e_2, e_1) > 0$. Consider the last such edge in the augmenting path. By backtracing its flow we find a positive flow path for $f_{\max, \text{red}, i}$ from σ to e_2 . This can only happen if the capacity of the path in Reduce_i is positive, which is equivalent to a prefix path with positive residual capacity in Shift_i for $f_{\max, \text{shift}, i}$. Tying the prefix and postfix at e_2 yields an augmenting path in Shift_i for $f_{\max, \text{shift}, i}$, impossible since $f_{\max, \text{shift}, i}$ is maximum. \square

Flow Expansion

The completion of stage i of the algorithm requires now to incorporate the event set δR_i to yield \mathbf{R}_i and allow the computation of $\mathbf{P}_{\max, i} = \mathbf{P}_{\max}(\mathbf{R}_i)$. Again, we define an incremental operation on an incremental residual flow network, the *flow expansion network*. The network is built on the residual network of $F(\mathbf{R}_{i-1} - \delta C_i)$ for flow $f_{\text{contr}, i}$. We call this residual network $\text{Res}(\text{Contr}_i)$.

Flow expansion network: *Expand_i* is a flow network with the intermediate events \mathbf{R}_i . *Expand_i* all flow edges $e_1 \rightarrow e_2$ in $\text{Res}(\text{Contr}_i)$, all flow edges in $F(\delta R_i)$ and an infinite capacity edge $e_1 \rightarrow e_2$ for each anti-precedence edge between $e_1 \in \delta R_i$ and $e_2 \in \mathbf{R}_{i-1} - \delta C_i$.

Note that by construction Expand_i is the residual network in $F(\mathbf{R}_i)$ for $f_{\text{contr}, i}$. We now define the final operator needed by the incremental envelope algorithm, Flow_Expansion .

Flow_Expansion($F(\mathbf{R}_{i-1} - \delta C_i)$, $f_{\text{contr}, i}$, δR_i , *Apres*):

- 1) Compute a maximum flow $f_{\max, \text{exp}, i}$ for *Expand_i*;
- 2) Return $f_{\max, i} = f_{\text{contr}, i} + f_{\max, \text{exp}, i}$

Theorem 7: $f_{\max, i}$ computed by Flow_Expansion is maximum for $F(\mathbf{R}_i)$.

Proof: $f_{\max, i}$ is clearly a flow for $F(\mathbf{R}_i)$. Moreover, $f_{\max, \text{exp}, i}$ is maximum for *Expand_i* and therefore there is no augmenting path in the corresponding residual network. The maximality of $f_{\max, i}$ follows from the identity between the residual network of *Expand_i* for $f_{\max, \text{exp}, i}$ and the residual network of $F(\mathbf{R}_i)$ for $f_{\max, i}$. \square

Flow Separation for \mathbf{P}_{\max}

We can achieve further performance improvements by minimizing the number of nodes and flow edges that need to be considered at each stage. During stage i two \mathbf{P}_{\max} are computed: $\mathbf{P}_{\max, \text{contr}, i}$ after $\text{Flow_Contraction}_i$ and $\mathbf{P}_{\max, i}$ after

Flow_Expansion. We know that each \mathbf{P}_{\max} is a predecessor set (i.e., it contains all of its successors in the anti-precedence graph), it is flow isolated (i.e., for each pair of events $e_1 \in \mathbf{P}_{\max}$ and $e_2 \in \mathbf{P}_{\max}^c$, $f_{\max}(e_1, e_2) = 0$ and $f_{\max}(e_2, e_1) = 0$) and has all exit edges saturated (i.e., $f_{\max}(e, \tau) = c(e, \tau)$ for all $e \in \mathbf{P}_{\max}$) (Muscuttola, 2002). This will allow us to prove that $F(\mathbf{P}_{\max, i-1})$ can be ignored during the computation of $\text{Flow_Contraction}_i$ and $F(\mathbf{P}_{\max, \text{contr}, i})$ can be ignored during computation of Flow_Expansion_i .

Let us consider each maximum flow operation executed at stage i . The first is flow shifting. Note that by construction, the \mathbf{P}_{\max} of \mathbf{E}_{i-1} , $\mathbf{P}_{\max, i-1}$, contains the events in $\mathbf{P}_{\max, i-1} - \delta C_i$. $\mathbf{P}_{\max, i-1}$ is a predecessor set since δC_i contains events at the bottom of the anti-precedence graph for $F(\mathbf{R}_{i-1})$. However, due to the additional links $e \rightarrow \tau$ the value of the positive residual of $\mathbf{P}_{\max, i-1}$ is equal to $\Delta(\mathbf{P}_{\max, i-1})$. $\mathbf{P}_{\max, i-1}$ is still flow insulated and has all exit edges saturated. Assume that flow shifting pumped flow to reach an event $e' \in \mathbf{P}_{\max, i-1}$. In order for at least part of such flow to reach τ there must be a postfix augmenting path that reaches τ from e' . But this is impossible since $\mathbf{P}_{\max, i-1}$ is a predecessor set, all postfix paths must remain inside $\mathbf{P}_{\max, i-1}$, and all exit edges from $\mathbf{P}_{\max, i-1}$ to τ are saturated. Therefore, any maximum flow algorithm pushing flows that searches for augmenting paths can avoid doing so in $\mathbf{P}_{\max, i-1}$ and any excess flow pumped into events of $\mathbf{P}_{\max, i}$ that can achieve τ will have to be pushed back from $\mathbf{P}_{\max, i-1}$ to $\mathbf{P}_{\max, i-1}^c$. Therefore we can ignore $\mathbf{P}_{\max, i-1}$ during flow shifting.

After flow shifting the maximum predecessor set is still $\mathbf{P}_{\max, i-1}$ since flow shifting simply produces a different maximum flow for \mathbf{E}_{i-1} and $\mathbf{P}_{\max, i-1}$ is independent from the specific flow instance (Muscuttola, 2002).

Considering now flow reduction, $f_{\max, \text{red}, i}$ this can be computed by simply backtracing flow in \mathbf{E}_{i-1} . Because of the flow insulation of $\mathbf{P}_{\max, i-1}$, this backtracing is either performed exclusively over events in $\mathbf{P}_{\max, i-1}^c = \mathbf{P}_{\max, i-1}^c - \delta C_i$ or is confined within the events in $\mathbf{P}_{\max, i-1} = \mathbf{P}_{\max, i-1} - \delta C_i$. Note that since after flow reduction all auxiliary edges must have zero flow, the producers' residual of $\mathbf{P}_{\max, i-1}$ after flow contraction must be equal to $\Delta(\mathbf{P}_{\max, i-1} - \delta C_i)$.

Finally, we can use a similar argument to the one used for flow shifting to show that Flow_Expansion_i can be performed entirely over $F(\mathbf{P}_{\max, \text{contr}, i}^c)$, therefore ignoring $\mathbf{P}_{\max, \text{contr}, i}$.

Incremental Computation of \mathbf{L}_{\max}

We are now ready to derive a recursive equations for the incremental calculation of $\mathbf{L}_{\max}(t)$ by transforming Equation 1 through the application of flow reduction and expansion.

From the discussion on flow separation, we know that, after $\text{Flow_Contraction}_{i-1}$, $\mathbf{P}_{\max, \text{contr}, i} = (\mathbf{P}_{\max}(\mathbf{R}_{i-1}) - \delta C_i) \cup \mathbf{P}_{\max}(\mathbf{P}_{\max}^c(\mathbf{R}_{i-1}) - \delta C_i)$. After Flow_Expansion_i , because of flow separation, we have $\mathbf{P}_{\max, i} = \mathbf{P}_{\max, \text{contr}, i} \cup \mathbf{P}_{\max}(\mathbf{P}_{\max, \text{contr}, i}^c \cup \delta R_i)$.

Theorem 8: $L_{\max}(t)$ satisfies this recursive equation:

if $t = t_1$

$$L_{\max}(t) = \Delta(C_1) + \Delta(P_{\max}(R_1))$$

if $t = t_i$ and $i > 1$

$$L_{\max}(t) = L_{\max}(t_{i-1}) + \begin{array}{l} i \\ \Delta(\delta C_i \cap P_{\max}^C(R_{i-1})) + \\ \Delta(P_{\max}(P_{\max}^C(R_{i-1}) - \delta C_i)) + \\ \Delta(P_{\max}(\delta R_i \cup P_{\max}^C(P_{\max}^C(R_{i-1}) - \delta C_i))); \end{array} \begin{array}{l} ii \\ iii \\ iv \end{array}$$

if $t \neq t_i$, then

$$L_{\max}(t) = L_{\max}(t-1).$$

Proof: $L_{\max}(t)$ only changes when R_t changes, i.e., at a time t_i . Let us consider in turn the application of **Flow_Contraction** _{i} and **Flow_Expansion** _{i} .

- a) **Flow_Contraction** _{i} : the level after flow contraction, $L_{\max, \text{contr}}(t_i)$ is the weight of the closed events after contraction and of $P_{\max, \text{contr}, i}$. Since new events at time t_i are only closed during flow contraction and C_i and $P_{\max, \text{contr}, i}$ are disjoint, $L_{\max, \text{contr}}(t_i) = \Delta(C_{i-1} \cup \delta C_i \cup P_{\max, \text{contr}, i}) = \Delta(C_{i-1} \cup \delta C_i \cup (P_{\max}(R_{i-1}) - \delta C_i) \cup P_{\max}(P_{\max}^C(R_{i-1}) - \delta C_i))$. Since for any two sets A and B it is $A \cup (B - A) = B \cup (A - B)$, with B and $(A - B)$ being disjoint sets, we have $\delta C_i \cup (P_{\max}(R_{i-1}) - \delta C_i) = P_{\max}(R_{i-1}) \cup (\delta C_i - P_{\max}(R_{i-1}))$. Since $\delta C_i \subseteq R_{i-1} = P_{\max}(R_{i-1}) \cup P_{\max}^C(R_{i-1})$, it is easy to see that $\delta C_i - P_{\max}(R_{i-1}) = \delta C_i \cap P_{\max}^C(R_{i-1})$. This yields $L_{\max, \text{contr}}(t_i) = L_{\max}(t_{i-1}) + \Delta(\delta C_i \cap P_{\max}^C(R_{i-1})) + \Delta(P_{\max}(P_{\max}^C(R_{i-1}) - \delta C_i))$, i.e., lines **i**, **ii** and **iii** in the theorem's statement.
- b) **Flow_Expansion** _{i} : the only new increment comes from set $P_{\max}(P_{\max}^C(P_{\max, \text{contr}, i} \cup \delta R_i) = P_{\max}(P_{\max}^C(R_{i-1} - \delta C_i) \cup \delta R_i)$ which yields line **iv** in the theorem's statement.

Algorithm	Time Complexity	Complexity Key
Labeling	$O(nmU)$	Total pushable flow
Capacity scaling	$O(nm \log U)$	Total pushable flow
Successive shortest paths	$O(n^2m)$	Shortest distance to τ
Generic preflow-push	$O(n^2m)$	Distance label
FIFO reflow-push	$O(n^3)$	Distance label

Table 1: Complexity of maximum flow algorithms

Figure 5 shows the pseudocode of the algorithm. The functions **Flow_Contraction** and **Flow_Expansion** receive as arguments the current flow network F_{cur} , which includes the current maximum flow, the incremental set of events that need to be added/deleted $E_{\text{cur}} = \{\text{earliest}, \text{latest}\}$, and the anti-precedence graph **Aprec(N)** for the set of all events in the plan **N**. **Aprec** carries the topological information needed to expand the flow network.

Given the current flow network and its maximum flow both stored in F_{cur} , **Extract_P_max** returns both its maximum increment predecessor set P_{\max} and the restricted network and flow resulting from the elimination of the P_{\max} . Comparing with the formula for $L_{\max}(t_i)$ described by

Incremental_Resource_Envelope (N, Aprec(N))

```

{ 1: E := { Group events in the input set N into entries E_i with three
        members: a time t and two lists earliest and latest. Event
        e ∈ N is included in E_i.earliest if et(e) = t and in E_i.latest if
        lt(e) = t. Sort the E_i in increasing order of t. }
  2: L_max := {<-∞, 0>} /* Maximum resource envelope. */
  3: t_cur := 0; /* Current time */
  4: L_old := 0; /* Envelope level at previous iteration. */
  5: L_new := 0; /* Envelope level at current iteration. */
  6: P_max := ∅; /* Maximum increment predecessors. */
  7: F_cur := ∅; /* Resource increment flow graph with associated
        maximum flow */
  8: E_cur := ∅; /* Entry from E at t_cur. */
  9: while (E is not empty)
  10: { E_cur := pop(E);
  11: t_cur := E_cur.t;
  12: L_new := L_old + weight (intersection (Events(F_cur), E_cur.latest));
  13: F_cur := Flow_Contraction (F_cur, E_cur.latest, Aprec(N));
  14: <P_max, F_cur> := Extract_P_Max (F_cur);
  15: L_new := L_new + weight (P_max);
  16: F_cur := Flow_Expansion (F_cur, E_cur.earliest, Aprec(N));
  17: <P_max, F_cur> := Extract_P_max (F);
  18: L_new := L_new + weight (P_max);
  19: L_max := append (L_max, <t_cur, L_new>);
  20: L_old := L_new;
    }
    return L_max;
}

```

Figure 5: Incremental envelope algorithm

Theorem 8, line 12 in the algorithm computes **i+ii**, line 15 adds **iii** and line 18 adds **iv**. Note that the algorithm is actually more of a method that can be implemented in different ways depending of which maximum flow algorithm is used in **Flow_Contraction** and **Flow_Expansion**. As we shall see the worst-case time complexity of the method is the same as that the maximum flow algorithm used. We will also see that further optimizations are possible when using specific algorithms such as push-preflow methods.

Complexity Analysis

The complexity analysis of the incremental envelope algorithm applies to a large number of maximum flow algorithms used for **Flow_Contraction** and **Flow_Expansion**. Each algorithm has a *complexity key*, i.e., a measurable entity whose static properties or dynamic behavior during its computations determines its time complexity. Table 1 (adapted from (Ahuja, Magnanti and Orlin, 1992)) reports the time complexity and complexity key of several maximum flow algorithms.

The *Labeling* and *Capacity Scaling* algorithms are based on the original Ford-Fulkerson method. The complexity depends on the strict monotonicity of the flow pushed at each iteration of the algorithm and on the fact that the *total pushable flow* is bound by nU where U is the maximum

capacity of an edge $\sigma \rightarrow e$ or $e \rightarrow \tau$. The *successive shortest paths* class of algorithms is based on the original Edmonds-Karp algorithm. The complexity depends on the fact that flow is pushed through augmenting paths of monotonically increasing length. The complexity key for this class of algorithms is the *shortest distance to τ* for each event e . For these algorithms it is crucial to demonstrate that the distance function $d(e)$ increases by at least one unit after each iteration.

Finally, preflow-push algorithms such as *generic preflow-push* and *FIFO preflow-push* (Goldberg and Tarjan, 1988) maintain a *distance labeling* $\underline{d}(e)$. These algorithms use purely local operations that push excess flow available at node e_1 through edges $e_1 \rightarrow e_2$ such that $\underline{d}(e_1) = \underline{d}(e_2) + 1$. When excess flow exists at some node and no such edge exist, the node's distance labeling is increased by the minimum amount that re-establishes a one unit difference in distance label over an edge. This allows more flow to be pushed. The complexity of the algorithms depends on creating a valid labeling at each iteration and on the fact that for each node the distance labeling is monotonically increasing up to $2n-1$.

We now analyze the cumulative cost of computing all flows over $2n$ stages respectively for $f_{\max, \text{shift}, i}$, $f_{\max, \text{red}, i}$ and $f_{\max, \text{exp}, i}$. First note that at each stage $f_{\max, \text{red}, i}$ can be computed by flow backtracing through a backwards depth first search on the resource increment flow network. Since this can cost up to $O(m)$, the total cost of computing flow reduction is $O(nm)$ and is therefore smaller than the cost of applying a regular maximum flow algorithm. Therefore we focus on the cost for the cumulative $f_{\max, \text{shift}, i}$ and $f_{\max, \text{exp}, i}$, respectively $F_{\text{shift}} = \sum_i f_{\max, \text{shift}, i}$ and $F_{\text{exp}} = \sum_i f_{\max, \text{exp}, i}$.

Lemma 9: *Each of F_{shift} and F_{exp} is no greater than nU .*

Proof: Consider F_{shift} (the argument is similar for F_{exp}). The upper bound of the total capacity of the edges $\sigma \rightarrow e$ entering Shift_i is the total capacity of edges $e \rightarrow \tau$ with $e \in \delta(C_i)$. After iteration i all nodes in δC_i are eliminated from further consideration, hence flow can go through each $\sigma \rightarrow e$ only during iteration i . Therefore, the total flow is upper bounded by $\sum_i |\delta C_i| U = nU$. \square

Note that the argument above does not hold for F_{exp} if instead of using flow shifting the flow is simply reduced and then expanded again (Kumar, 2003). In this case the same flow could be pushed up to n times with F_{exp} being $O(n^2U)$. This would not improve on the staged envelope algorithm.

Consider now the distance $d(e)$ for node e and how it changes when computing $f_{\max, \text{shift}, i}$ and $f_{\max, \text{exp}, i}$. Let us call $d_{\text{shift}, i}^0(e)$ and $d_{\text{shift}, i}^f(e)$ the distances at the beginning and at the end of flow shifting for iteration i . We define similarly $d_{\text{exp}, i}^0(e)$ and $d_{\text{exp}, i}^f(e)$.

Lemma 10: $d_{\text{exp}, i-1}^f(e) \leq d_{\text{shift}, i}^0(e)$ and $d_{\text{shift}, i}^f(e) \leq d_{\text{exp}, i}^0(e)$.

Proof: Between the end of flow expansion at iteration $i-1$ and the start of flow shifting at iteration i , the auxiliary

flow network changes through the elimination of nodes and edges in $F(\delta C_i)$. Therefore, the new distances in the remaining residual capacity network can only increase. Since Shift_i only adds edges $\sigma \rightarrow e$, the distances in Shift_i must analogously increase and therefore $d_{\text{exp}, i-1}^f(e) \leq d_{\text{shift}, i}^0(e)$. For Expand_i node distances can further increase because flow reduction can only eliminate residual network edges present in Shift_i for $f_{\max, \text{shift}, i}$. Also, from Lemma 3 the addition of $F(\delta R_i)$ cannot reduce distances since it cannot add any edge from an event in Shift_i to one in $\delta R_{i(i)}$. Therefore, $d_{\text{shift}, i}^f(e) \leq d_{\text{exp}, i}^0(e)$. \square

Note that the argument in Lemma 10 does not hold if events are added in arbitrary order. In this case the addition of edges can reduce the distance function of some node e between a shifting and an expansion phase. In the worst case, this may reduce the distance to one for each application of maximum flow and therefore does not improve on the staged algorithm.

Finally, consider reusing distance labeling across preflow-pushes for shifting and expansion. $\underline{d}_{\text{shift}, i}^0$, $\underline{d}_{\text{shift}, i}^f$, $\underline{d}_{\text{exp}, i}^0$ and $\underline{d}_{\text{exp}, i}^f$ are the distance labelings at the beginning and end of shifting and expansion. Assume also that the distance label of a node that has not yet entered Expand_i or Shift_i is zero.

Lemma 11: $\underline{d}_{\text{shift}, i}^0$ can be made equal to $\underline{d}_{\text{exp}, i-1}^f$ for all nodes in Shift_i . Also, $\underline{d}_{\text{exp}, i}^0$ can be made equal to $\underline{d}_{\text{shift}, i}^f$ for all nodes in Expand_i .

Proof: The distance label of a node remains valid when edges are deleted or new edges are only added to enter it from new nodes. Also, a distance function at node e must be an upper bound of its labeling. From Lemma 10 we know that the distance function can only increase from Expand_{i-1} to Shift_i and from Shift_i to Expand_i . Therefore $\underline{d}_{\text{exp}, i-1}^f$ and $\underline{d}_{\text{shift}, i}^f$ are valid choices respectively for $\underline{d}_{\text{shift}, i}^0$ and $\underline{d}_{\text{exp}, i}^0$. \square

We can now prove the main complexity result.

Theorem 12: *For a large class of maximum flow algorithms, **Incremental_Resource_Envelope** has time complexity $O(\text{Maxflow}(n, m, U))$.*

Proof: Assume we applied one of the maximum flow algorithms in Table 1 to find the full maximum flow on the entire resource increment flow network (e.g., to compute the maximum envelope level over the entire time horizon (Muscettola, 2002)) with cost $O(\text{Maxflow}(n, m, U))$. We use Lemmas 9, 10 and 11 to prove that the cumulative cost of using the same algorithms for flow shifting and flow expansion is $O(\text{Maxflow}(n, m, U))$ for the same algorithm.

1. *Labeling and Capacity scaling:* Lemma 9 shows that the worst case bound for the total flow moved during shifting and expansion is at worst twice that for full flow calculation. Also, at each iteration during shifting and expansion, finding an augmenting path costs at most m as for full flow calculation. Hence shifting and expansion cost at most $O(\text{Maxflow}(n, m, U))$.

2. *Successive shortest paths*: the cost bound for each full flow augmentation is an upper bound for that in shifting and expansion. The algorithm's complexity also depends on the monotonic increase of the distance function up to \mathbf{n} after each elementary operation. Note that until the deletion of a δC_i or a P_{\max} , a node's distance is bound by \mathbf{n} as for the full flow. Monotonic increase is guaranteed by the algorithm within each shifting and expansion phase and by Lemma 10 across phases. Hence, the cost is $O(\text{Maxflow}(\mathbf{n}, \mathbf{m}, \mathbf{U}))$.
3. *Preflow-push methods*: the complexity is found through amortized analysis (Goldberg and Tarjan, 1988), relying on an appropriate potential function Φ and on the determination of its possible variations after the applying a local operation (e.g., a saturating or a non-saturating preflow push). One key observation is the monotonic increase of each node's distance label for each local operation. Both for the incremental and for the full flow this increase is bound by $2\mathbf{n}-1$ and Lemma 11 guarantees monotonic distance label increase across phases. Note that, unlike for the full flow, for shifting and expansion Φ increases also at the beginning of each shifting phase, when nodes are activated by the creation of initial flow excesses. However, a detailed amortized analysis (omitted for space limitations) shows that this increase does not affect the order of complexity of the shifting and expansion phases that remains $O(\text{Maxflow}(\mathbf{n}, \mathbf{m}, \mathbf{U}))$.

The worst case complexity of the other phases of *Incremental_Resource_Envelope* besides shifting and expansion are dominated by $O(\text{Maxflow}(\mathbf{n}, \mathbf{m}, \mathbf{U}))$. Flow reduction is cumulatively $O(\mathbf{nm})$. The total cost of **Extract_P_max** and of incrementally constructing and deleting the flow network, is $2 O(\mathbf{m})$. Finally the sorting of events during initialization is $O(\mathbf{n} \log \mathbf{n})$. \square

Optimized Preflow-Push Implementation

If the incremental algorithm is implemented using a preflow-push method, the previous complexity analysis indicates that, in order to reduce complexity, we need to reuse the distance labeling function from the end of a maximum flow computation to the start of the next.

A further optimization is possible. Consider the maximum flow calculation on **Shift_i**. During initialization, an excess flow is loaded on each event \mathbf{e} for each edge $\sigma \rightarrow \mathbf{e}$ in **Shift_i**. We know that only a fraction of this excess flow may reach τ . The remainder will be pushed back out of **Shift₁** during flow shifting and then pushed again through the flow network during flow reduction. In other words, this flow travels *twice* through the network before being eliminated. We can remove this duplication as flows. Assume that instead of deleting the $\sigma \rightarrow \mathbf{e}$ edges of E_{i-1} when constructing **Shift_i**, we delete instead the $\sigma \rightarrow \mathbf{e}$ edges of **Shift_i** after having performed the appropriate excess loading needed to perform flow shifting. In this case the flow that cannot be shifted will be pushed back to the source in E_{i-1} , i.e., $F(\mathbf{R}_i)$, instead making the additional $O(\mathbf{nm})$ cost of flow reduction

unnecessary. Another possible optimization consists of combining preflow-push through **Shift_i** and **Expand_i** by connecting δR_i before running the shift/reduce preflow-push. In this case the flow excess initializations of contraction and expansion are combined and a single preflow-push is run during phase \mathbf{i} .

These optimizations do not affect asymptotic complexity but may have a significant effect in practice. Empirical studies will be needed to assess their actual usefulness.

Conclusions

We presented a new algorithm that efficiently computes resource envelopes for flexible plans. The method has complexity $O(\text{Maxflow}(\mathbf{n}, \mathbf{m}, \mathbf{U}))$ where \mathbf{n} and \mathbf{m} measure the size of the activity plan and \mathbf{U} measures the maximum resource consumption or production of an activity. We are currently experimenting with various implementations of the methods. While we expect that for large problem sizes the $O(\mathbf{n})$ cost reduction will be evident, practical improvements on smaller problems require careful design of efficient and minimal data structures. We are also addressing the second part of the cost/benefit equation for envelopes by exploiting additional structural information resulting from the method's incrementality and designing scheduling algorithms that use a minimal search space.

References

- R.K. Ahuja, T.L. Magnati, J.N. Orlin, 1993. *Network Flows*, Prentice Hall, NJ.
- R.K. Ahuja, M. Kodialam, A.K. Mishra, J.B. Orlin, 1997. Computational Investigations of Maximum Flow Algorithms, *European Journal of OR*, Vol 97(3).
- T.H. Cormen, C.E. Leiserson, R.L. Rivest, 1990. *Introduction to Algorithms*. Cambridge, MA.
- A.V. Goldberg, R.E. Tarjan, 1988. A New Approach to the Maximum-Flow Problem. *JACM*, Vol. 35(4).
- T.K.S. Kumar, 2003. Incremental Computation of Resource-Envelopes in Producer-Consumer Models, *Procs. of CP2003*, Kinsale, Ireland.
- P. Laborie, 2001. Algorithms for Propagating Resource Constraints in AI Planning and Scheduling: Existing Approaches and New Results, *Procs. ECP 2001*, Spain.
- P. Morris, N. Muscettola, T. Vidal, 2001. Dynamic Control of Plans with Temporal Uncertainty, in *Procs. of IJCAI 2001*, Seattle, WA.
- N. Muscettola, 2002. Computing the Envelope for Stepwise-Constant Resource Allocations, *Procs. of CP2002*, Ithaca, NY.
- N. Pollicella, S.F. Smith, A. Cesta, A. Oddi, 2003. Steps toward Computing Flexible Schedules, *Procs. of Online-2003 Workshop at CP 2003*, Kinsale, Ireland, <http://www.cs.ucc.ie/~kb11/CP2003Online/onlineProceedings.pdf>