CLUSTERING FORMALISM FOR SYNCHRONOUS DATAFLOW

by

Shuvra S. Bhattacharyya

Memorandum No. UCB/ERL M92/30

6 April 1992

CLUSTERING FORMALISM FOR SYNCHRONOUS DATAFLOW

by

Shuvra S. Bhattacharyya

Memorandum No. UCB/ERL M92/30

6 April 1992

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

TITLE PAGE

CLUSTERING FORMALISM FOR SYNCHRONOUS DATAFLOW

by

Shuvra S. Bhattacharyya

Memorandum No. UCB/ERL M92/30

6 April 1992

ELECTRONICS RESEARCH LABORATORY

College of Engineering
University of California, Berkeley
94720

CLUSTERING FORMALISM FOR SYNCHRONOUS DATAFLOW

Shuvra S. Bhattacharyya

ABSTRACT

This document presents a formalization of "clustering" nodes during the scheduling of a synchronous dataflow (SDF) graph. The interpretation of clustering in this document is specific to our work in *scheduling for efficient iteration*, and it is possible that this formalism may not apply to clustering in other contexts. The goal of this document is to define our interpretation of a cluster; to define the process by which a cluster in a graph derives a new graph; and then to show that scheduling this new graph is equivalent to scheduling the original graph with the desired clustering effect.

The definitions and theorems presented here apply to the single-processor scheduling problem. The intent presently is to gain intuition about the iteration problem alone, before attempting to combine these considerations with parallel-processor scheduling. It is expected that this formalism will eventually be extended to the multiprocessor case.

This document expands on the discussion in chapter 3 of [1] on "consolidating a subgraph" of an SDF graph.

The following notation will be used throughout this document:

Notation

(N1) An SDF graph is denoted by an ordered pair {X,Y}, where X is the set of nodes and Y is the set of arcs.

(N2) For an SDF arc α, source(α) and sink(α) denote respectively the source node and sink node of α. p(α) denotes the number of samples produced on α for every invocation of source(α). Similarly, c(α) denotes the number of samples consumed from α by every invocation of sink(α). Finally, delay(α) denotes the delay on α.

(N3) We use the term MPASS (*minimum*-PASS) to denote a PASS (periodic admissible sequential schedule) of blocking factor one.

(N4) For set operands, we use "+" and "-" to denote the union and exclusion operations respectively.

Definition 1 Let G = {N, A} be an SDF graph; let S be a PASS for G; and let V be a connected subset of N. Then we say that V *is a cluster in* G *for* S if and only if S can be expressed as

$$(1)\quad \Xi_0\Pi_1\Xi_1\Pi_2\cdots\Xi_{n-1}\Pi_n\Xi_n,$$

where each Ξ_i is a (possibly empty) sequence of invocations of nodes in N – V, and each Π_i is an MPASS for the subgraph associated with V.

Observe that if the decomposition (1) exists for a schedule S and subset V, then it must be unique for (S,V). Thus we can define

$$\Xi(S,V,G) = \{\Xi_0,\Xi_1,\cdots,\Xi_n\},\text{ and}$$
$$\Pi(S,V,G) = \{\Pi_1,\Pi_1,\cdots,\Pi_n\}$$

whenever V is a cluster in G for S.

* * * * *

The following assumptions will be used throughout the remainder of this document.

Assumptions

(A1) Let $G = \{N, \Lambda\}$ be an SDF graph for which a PASS exists.

(A2) Let V be a connected subset of N and let $G_v = \{V, \Lambda_v\}$ be the subgraph associated with V.

(A3) Define a node n_v (to represent an instance of the cluster V).

(A4) Let Γ_v be the topology matrix for G_v and let q_v be the smallest positive integer vector in the null space of Γ_v. The existence of q_v is guaranteed by lemma 1 below and theorem 4 in chapter 2 of [1].

(A5) For each arc $\alpha \in \Lambda$ which is directed *from* a node $n_1 \in N-V$ *to* a node $n_2 \in V$, define the arc $\hat{\alpha}$ to be directed *from* n_1 *to* n_v, with $p(\hat{\alpha}) = p(\alpha)$, $c(\hat{\alpha}) = q_v[n_2]c(\alpha)$, and $\mathrm{delay}(\hat{\alpha}) = \mathrm{delay}(\alpha)$.

(A6) Similarly, for each arc $\beta \in \Lambda$ which is directed *from* a node $n_1 \in V$ *to* a node $n_2 \in N-V$, define the arc $\hat{\beta}$ to be directed *from* n_v *to* n_2, with $p(\hat{\beta}) = q_v[n_1]p(\beta)$, $c(\hat{\beta}) = c(\beta)$, and $\mathrm{delay}(\hat{\alpha}) = \mathrm{delay}(\alpha)$.

(A7) Let $\Lambda_1 = \{\gamma \in \Lambda \mid \mathrm{source}(\gamma) \in N-V, \text{ and } \mathrm{sink}(\gamma) \in V\}$;

$\Lambda_2 = \{\gamma \in \Lambda \mid, \mathrm{source}(\gamma) \in V, \text{ and } \mathrm{sink}(\gamma) \in N-V\}$;

$\hat{\Lambda}_1 = \{\hat{\gamma} \mid \gamma \in \Lambda_1\}$;

$\hat{\Lambda}_2 = \{\hat{\gamma} \mid \gamma \in \Lambda_2\}$; and

$\hat{\Lambda} = \hat{\Lambda}_1 + \hat{\Lambda}_2$.

Definition 2 We define $F_v(G)$, called *the graph obtained by clustering V in G*, by

$$F_v(G) = \{N - V + \{n_v\}, \Lambda - \Lambda_v + \hat{\Lambda} - \Lambda_1 - \Lambda_2\}$$

Lemma 1 Assume A1-A2 and suppose S is a PASS for G. Then the schedule S', obtained by eliminating from S the firings of nodes not in V, is a PASS for the subgraph associated with V.

Proof The subset of firings in S which produce or consume data on arcs in G_v is precisely S'. Let x be a firing in S' and suppose x does not have enough data on input arc $\alpha \in \Lambda_v$ to fire. Then there must be some $n \in V$, which is fired before x in S — but not fired in S' — to produce one or more of the samples which S' is "missing" on α for x. But this contradicts our definition of S', and it follows that S' is admissible. Furthermore, since S leaves no net change in the number of samples on any arc in G, S' leaves no net change on any arc in G_v. We conclude that S' is a PASS for G_v.

Definition 3 Let G be an SDF graph and let S be a schedule for G. Given an arc α in G and a firing f in S, we define

$b_S(\alpha, f) =$
{the number of times source(α) is invoked in S before f } * $p(\alpha)$
- {the number of times sink(α) is invoked in S before f } * $c(\alpha)$
+ delay(α)

If S_1 is a nonempty subschedule (a subsequence of firings) of S then we define $b_S(\alpha, S_1) = b_S(\alpha, f_1)$, where f_1 is the first firing in S_1.

The following fact is immediately apparent from definition 3:

Fact 1: If S is an admissible schedule, then $b_S(\alpha, f)$ equals the number of samples on α immediately before f is fired in S.

Lemma 2 Assume the hypotheses and notation developed in A1-A7 above. Let S be a schedule for $F_v(\mathrm{G})$ and let S' be the schedule for G which results from replacing in S each firing of n_v with an MPASS for G_v. Suppose S_i is the ith MPASS for G_v in S'. Then:

(A) $\alpha \in \Lambda_1 \Rightarrow b_S(\hat{\alpha}, \text{the } i\text{th invocation of } n_v) = b_{S'}(\alpha, S_i)$; and

(B) $\alpha \in \Lambda_2 \Rightarrow b_S(\hat{\alpha}, \text{the } i\text{th invocation of sink}(\hat{\alpha})) = b_{S'}(\alpha, \text{the } i\text{th invocation of sink}(\hat{\alpha}))$.

Proof of A

From our assumptions, S and S' can be decomposed as :

(2A) $S = \Xi_0 n_v \Xi_1 n_v \cdots \Xi_{M-1} n_v \Xi_M$

(2B) $S' = \Xi_0 \Pi_1 \Xi_1 \Pi_2 \cdots \Xi_{M-1} \Pi_M \Xi_M$,

where each Ξ_i is a subschedule involving nodes in N−V and each Π_i is an MPASS for G_v. We wish to show that for $1 \le i \le M$, $b_{S'}(\alpha, \Pi_i) = b_S(\alpha, \text{the } i\text{th invocation of } n_v)$. Now,

(3) $b_{S'}(\alpha, \Pi_i) =$
{the number of times source(α) is invoked in $\Xi_0 \cdots \Xi_{i-1}$} * p(α)
- $(i-1)$ * {the number of times sink(α) is invoked in an MPASS for G_v} * c(α)
+ delay(α)

The second term in (3) is equal to

$(i-1) * q_v[\text{sink}(\alpha)] * \text{c}(\alpha)$
$= (i-1) * \text{c}(\hat{\alpha})$. (from A5)

Also, from A5, $\text{p}(\hat{\alpha}) = \text{p}(\alpha)$, $\text{delay}(\hat{\alpha}) = \text{delay}(\alpha)$, and $\text{source}(\hat{\alpha}) = \text{source}(\alpha)$, so (3) becomes

{the number of times source($\hat{\alpha}$) is invoked in $\Xi_0 \cdots \Xi_{i-1}$} * p($\hat{\alpha}$)
- $(i-1)$ * c($\hat{\alpha}$)
+ delay($\hat{\alpha}$),

which, from definition 3, is clearly equal to $b_S(\hat{\alpha}, \text{the } i\text{th invocation of } n_v)$.

Proof of B

Suppose the ith invocation of sink(α) in S' is in Ξ_k of decomposition 2B. Then, since sink(α) is not in V, it is clear that the ith invocation of sink($\hat{\alpha}$) — which is the same as sink(α) — in S is in Ξ_k of decomposition 2A. Now since subschedules $\Pi_1, \cdots, \Pi_k$ precede Ξ_k in S',

$$b_{S'}(\alpha, \text{the } i\text{th invocation of sink}(\alpha)) = k q_v[\text{source}(\alpha)] * \text{p}(\alpha) - (i-1) * \text{c}(\alpha) + \text{delay}(\alpha),$$

which from A5-A6 equals :

$$k * \text{p}(\hat{\alpha}) - (i-1) * \text{c}(\hat{\alpha}) + \text{delay}(\hat{\alpha})$$

$$= b_S(\hat{\alpha}, \text{the } i\text{th invocation of sink}(\hat{\alpha})).$$

QED.

Lemma 3 Let G = {N, Λ} be an SDF graph. Suppose V is a connected subset of N and let G_v = {V, Λ_v} denote the subgraph associated with V. Let Γ, Γ_v, and Γ_F respectively denote topology matrices for G, G_v and F_v(G). Finally, suppose that PASSes for G and G_v exist, and let q, q_v and q_F be the smallest positive-integer vectors in the left-side-null-spaces of Γ, Γ_v and Γ_F, respectively.

Then, $x \in V \Rightarrow q[x] = q_F[n_v] * q_v[x]$.

Proof

Γ can be expressed as

$$\begin{bmatrix} \Gamma_1 & 0 \\ \Gamma_2 & \Gamma_v \end{bmatrix},$$

where Γ_1 denotes the submatrix associated with the interaction between nodes in N−V and arcs in $\Lambda - \Lambda_v$, and Γ_2 denotes the submatrix associated with V and $\Lambda - \Lambda_v$.

Also, q can be decomposed as

$$\begin{bmatrix} q_1 \, q_2 \end{bmatrix},$$

where q_1 denotes the subvector associated with nodes in N−V and q_2 denotes the entries associated with the elements of V.

With these decompositions, we have that

$$\begin{bmatrix} q_1 \, q_2 \end{bmatrix} \begin{bmatrix} \Gamma_1 & 0 \\ \Gamma_2 & \Gamma_v \end{bmatrix} = 0,$$

which implies that

$$q_1\Gamma_1 + q_2\Gamma_2 = 0$$

and

$$q_2\Gamma_v = 0.$$

Thus q_2 must be some positive integer multiple k of q_v. I.e. $q_2 = kq_v$ for some $k \in \{1, 2, 3, \cdots\}$

Now from A5-A6, it can be deduced that we may express Γ_F as

$$\begin{bmatrix} \Gamma_1 \\ q_v\Gamma_2 \end{bmatrix},$$

where $q_v\Gamma_2$ determines the row corresponding to the node n_v. It follows then, that

$$q' \equiv \begin{bmatrix} q_1 & k \end{bmatrix}$$

is in the null space of Γ_F, since

$$q_1\Gamma_1 + kq_v\Gamma_2 = q_1\Gamma_1 + q_2\Gamma_2 = 0$$

Now suppose that there was an integer $r>1$ such that

$$\frac{q'}{r}$$

was also a positive integer vector in $\eta(\Gamma_F)$. Then

$$\frac{q_1}{r}\Gamma_1 + \frac{kq_v}{r}\Gamma_2 = \frac{q_1}{r}\Gamma_1 + \frac{q_2}{r}\Gamma_2 = 0$$

and thus $q\,/\,r$ is a positive integer vector in the null space of Γ. But this contradicts our assumption that q is the smallest such vector.

We conclude that

$$q_F = \left[\begin{array}{cc} q_1 & k \end{array}\right].$$

Thus, if $x \in V$ then $q[x] = q_2[x] = kq_v[x] = q_F[n_v]q_v[x]$.

QED.

Theorem 1 Assume A1-A7. Let S_v be an PASS of blocking-factor k for F_v(G). Then the schedule S', obtained by substituting each occurrence of n_v in S_v with an MPASS for G_v, is a PASS of blocking-factor k for G, and V is a cluster in G for S'.

Proof

From these assumptions, S_v and S' can be decomposed as S and S' are (respectively) decomposed in 2A-2B. Thus

(4A) $S_v = \Xi_0 n_v \Xi_1 n_v \cdots \Xi_{M-1} n_v \Xi_M$
(4B) $S' = \Xi_0 \Pi_1 \Xi_1 \Pi_2 \cdots \Xi_{M-1} \Pi_M \Xi_M$,

and each Ξ_i is a subschedule involving nodes in N−V, and each Π_i is an MPASS for G_v. It is obvious from (4B) that V is a cluster in G for S'.

Now we show by contradiction that S' is admissible. Suppose that x is the first firing in S' which does not have enough data to fire, and suppose x is "missing" samples from arc α.

(A) Suppose $\alpha \in \Lambda - \Lambda_v - \Lambda_1 - \Lambda_2$. Then α is in both G and F_v(G). Since neither n_v nor any Π_i interacts with $(\Lambda - \Lambda_v - \Lambda_1 - \Lambda_2)$, it follows that $x \in$ some Ξ_i, and that x does not have enough samples from α in S_v as well. This contradicts the assumption that S_v is admissible.

(B) Suppose $\alpha \in \Lambda_v$. Then $x \in$ some Π_i. Since the Ξ_i's don't interact with Λ_v, it follows that Π_i is not admissible, which is another contradiction.

(C) Suppose $\alpha \in \Lambda_1$. Then $x \in$ some Π_i. Since source(α) is outside G_v, $b_S(\alpha, \Pi_i) < q_v[\text{sink}(\alpha)]$ * c(α). From Lemma 2A and assumption A5, this inequality implies that $b_{S_v}(\hat{\alpha}$, the ith invocation of $n_v) <$ c($\hat{\alpha}$). From fact 1, we see that this last inequality contradicts our assumption that S_v is a admissible.

(D) Finally, suppose that $\alpha \in \Lambda_2$. Then then x denotes some firing, say the ith one, of sink(α) and $b_{S'}(\alpha$, the ith invocation of sink(α)) < c(α).

From lemma 2B and assumption A5, this inequality implies that $b_{S_v}(\hat{\alpha}$, the ith invocation of sink($\hat{\alpha}$)) < c($\hat{\alpha}$). From fact 1, this contradicts our assumption that S_v is admissible.

(A)-(D) together prove that the assumption that S' is not admissible cannot hold.

To prove that S' is periodic, we must show that S' produces no net change in the number of samples on any arc α in G.

(E) Suppose $\alpha \in \Lambda - \Lambda_1 - \Lambda_2 - \Lambda_v$. Then $\alpha \in F_v$(G) and source(α) and sink(α) are the same in G and F_v(G). Furthermore, for both schedules, source(α) and sink(α) are fired only within the Ξ_i subschedules of decompositions 4A-4B. Thus the number of firings of source(α) and sink(α) in S are the same, respectively, as the number of firings of source(α) and sink(α) in S_v. Since S_v is periodic it follows that S' can produce no net change in the number of tokens on α.

(F) Suppose $\alpha \in \Lambda_v$. Then in S', source(α) and sink(α) are fired only within $\Pi_1, \Pi_2, \cdots, \Pi_M$. Since each Π_i is assumed to be an MPASS, it follows that there can be no net change on α.

(G) Suppose $\alpha \in \Lambda_1$. Then from A5, each invocation of n_v in S_v consumes the same number of tokens from $\hat{\alpha}$ as each Π_i, and thus the total number of samples consumed on α through S' is the same as the number consumed from $\hat{\alpha}$ through S_v. Also source(α)=source($\hat{\alpha}$) is not in V, so it is fired the same number of times in both schedules. Since p(α) = p($\hat{\alpha}$), it follows that the net change on α is equal to the net change on $\hat{\alpha}$, which is 0 since S_v is a pass.

(H) Suppose $\alpha \in \Lambda_2$, then sink(α)=sink($\hat{\alpha}$) is outside V, so the number of samples consumed from α during S' is the same as the number consumed from $\hat{\alpha}$ during S_v. Furthermore, each invocation of n_v produces the same number of samples on $\hat{\alpha}$ as each Π_i produces on α. Thus the net change on α equals the net change on $\hat{\alpha}$, which is 0.

(E)-(H) show that S' produces no net change in sample-population on any arc in G, and we conclude that S' is periodic.

Finally, we show that S' has blocking factor k. This is equivalent to showing that for some node x in G, S' fires x ($k * q[x]$) times, where q is the smallest integer vector in the null space of Γ. Let our x be chosen from within V. Also let q_F and q_v denote the smallest positive - integer - null - space - vectors for the topology matrices of F_v(G) and G_v respectively. Then S' fires k sets of $q_F[n_v]$ invocations of MPASSes for G_v, each containing $q_v[x]$ invocations of x, for a total of ($k * q_F[n_v] * q_v[x]$) firings of x. From lemma 3, this is equal to ($k * q[x]$) firings.

QED.

Theorem 2 Assume A1-A7. Let S be a PASS of blocking-factor k for G, and suppose that V is a cluster in G for S. Then the schedule S', obtained by substituting in S each occurrence of an MPASS for G_v with a firing of n_v, is a PASS of blocking-factor k for F_v(G).

The proof is analogous to that of Theorem 1, and we omit it here for brevity. Theorems 1 and 2 together indicate that scheduling the "clustered" graph F_v(G) is equivalent to producing schedules in G for which V is cluster — every PASS for G in which V is a cluster can be obtained by scheduling F_v(G), and from every PASS for F_v(G), we can derive a schedule for G in which V is a cluster.

We conclude this document with a fact that is useful when considering only two nodes as a cluster. We have found that such *pairwise clustering* is useful for organizing looping in a schedule.

Fact: If V = {A,B} is a two-element subset of N then

$$q_v[\mathrm{A}] = \frac{q[\mathrm{A}]}{\gcd(q[\mathrm{A}], q[\mathrm{B}])} \; ;$$

$$q_v[\mathrm{B}] = \frac{q[\mathrm{B}]}{\gcd(q[\mathrm{A}], q[\mathrm{B}])} \; ; \text{ and}$$

$$q_F[n_v] = \gcd(q[\mathrm{A}], q[\mathrm{B}]) \; .$$

(''gcd'' denotes the *greatest common divisor*)

Proof:

From lemma 3, each of the equalities above implies the other two. Here we will directly prove the first two.

From lemma 1, we know that G_v has a PASS, and thus that Γ_v has rank 1. It follows that there are integers a and b such that

(5) $[q_1 \; q_2] \in \eta(\Gamma_v) \Leftrightarrow q_1 a + q_2 b = 0.$

From lemma 3, we know that $\forall\, x \in \mathrm{V}$, $q[x] = q_F[n_v] * q_v[x]$, and thus $[q[\mathrm{A}] \; q[\mathrm{B}]] \in \eta(\Gamma_v)$. Hence:

$q[\mathrm{A}]*a + q[\mathrm{B}]*b = 0.$

=> the smallest positive integer vector that satisfies the RHS of (5) is

$$\left[\frac{q[\mathrm{A}]}{\gcd(q[\mathrm{A}], q[\mathrm{B}])} \quad \frac{q[\mathrm{B}]}{\gcd(q[\mathrm{A}], q[\mathrm{B}])} \right]$$

QED.

References

1. E.A. Lee, Ph.D. Thesis, U.C. Berkeley, June 1986.