Report Number: WUCS-88-29

1988-09-01

Practical Wide-Sense Nonblocking Generalized Connectors

Jonathan S. Turner

In this note, we show that wide-sense nonblocking networks can be obtained by cascading a pair of Cantor networks or a pair of Clos networks. The only constraint placed on the routing algorithms is that branching be restricted to the second network in the cascade. This result yields practical network for multipoint communication with complexities O(N(logN)2 and O(N1+1/r).

PRACTICAL WIDE-SENSE NONBLOCKING GENERALIZED CONNECTORS

Jonathan S. Turner

WUCS-88-29

Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
Saint Louis, MO 63130-4899

Abstract

In this note, we show that wide-sense nonblocking networks can be obtained by cascading a pair of Cantor networks or a pair of Clos networks. The only constraint placed on the routing algorithm is that branching be restricted to the second network in the cascade. This result yields practical network for multipoint communication with complexities $O(N(logN)^2)$ and $O(N^{1+1/r})$.

This work supported by the National Science Foundation (grant DCI 8600947), Bell Communications Research, Bell Northern Research, Italtel SIT and NEC.

Practical Wide-Sense Nonblocking Generalized Connectors

Jonathan S. Turner

A *network* can be modeled as a directed graph with a designated set of input nodes and a designated set of output nodes. A *connection network* or *connector* is a network used to connect designated pairs of inputs and outputs. The paths used to connect distinct pairs are not permitted to share links. Beneš [1], defines a connector to be strictly nonblocking if given any legal configuration of the network, it is always possible to add a connection between a specified idle input and a specified idle output, without reassigning existing connections to different paths. He also defines a connector to be *wide-sense nonblocking* is there is a routing algorithm which avoids blocking states given any legitimate sequence of connection and disconnection requests. A *generalized connector* is a network used to connect designated inputs to designated output sets; that is a network used to establish multipoint communication channels. The definitions of strict and wide-sense nonblocking extend to generalized connectors in the obvious way. See [1] and [5] for more precise definitions.

In this note, we show that wide-sense nonblocking networks can be obtained by cascading a pair of Cantor networks [2] or a pair of Clos networks [3]. The only constraint placed on the routing algorithm is that branching be restricted to the second network in the cascade. This result yields practical networks for multipoint communication with complexities $O\left(N(\log N)^2\right)$ and $O\left(N^{1+1/r}\right)$. Pippenger [6] describes a wide-sense nonblocking generalized connector with complexity $O\left(N(\log N)^3\right)$. This appears to be the best asymptotic complexity for a network with the potential for practical application. Feldman, Friedman and Pippenger [4] show how to construct wide-sense nonblocking generalized connectors with complexity $O(N \log N)$, but the resulting networks are of purely theoretical interest as the constants involved are quite large and the routing problem is NP-complete.

The Beneš network [1], $B_{N,k}$ can defined by the recursive construction shown in Figure 1. The base of the recursion is a $k \times k$ crossbar. The Cantor network $K_{N,k,m}$ can be defined by the construction shown in Figure 2. The following result is due to Cantor [2]. We include it here for reference.

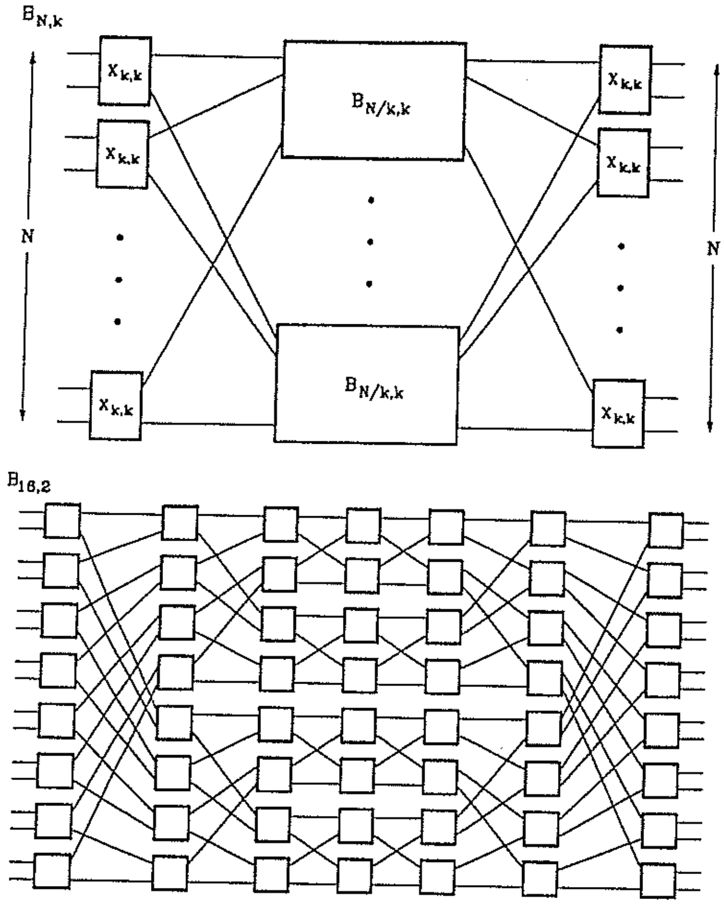

Figure 1: Beneš Network

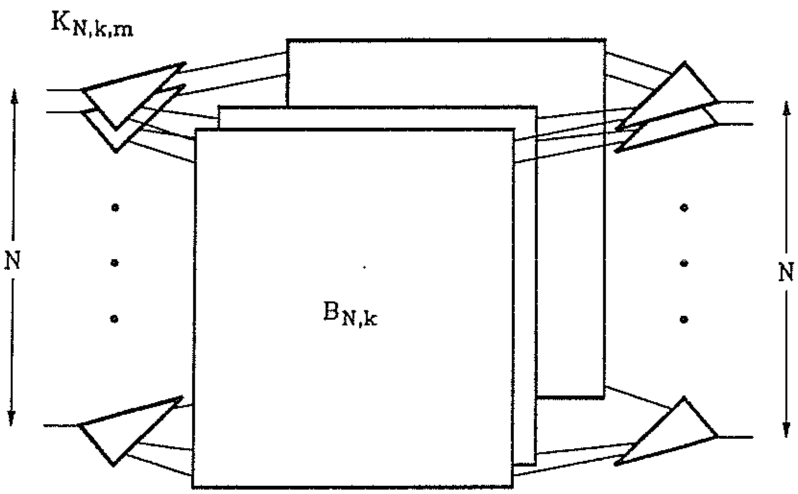

Figure 2: Cantor Network

THEOREM 1. *The Cantor network $K_{N,k,m}$ is a strictly nonblocking connector if $m > 2\frac{k-1}{k}\log_k(N/k)$.*

Proof. Assume the network is in some arbitrary state and we are asked to connect an idle input x to an idle output y. Define α_i to be the number of stage i links accessible from x. Then,

$$\begin{aligned}
\alpha_1 &= m \\
\alpha_2 &\geq mk-(k-1) \\
\alpha_3 &\geq k\alpha_2-(k^2-k) \geq mk^2-2k^2+2k \\
\alpha_4 &\geq k\alpha_3-(k^3-k^2) \geq mk^3-3k^3+3k^2 \\
&\vdots \\
\alpha_i &\geq mk^{i-1}-(i-1)k^{i-1}+(i-1)k^{i-2}
\end{aligned}$$

Substituting $h = \log_k N$, we find that the number of stage h links and hence the number of middle stage crossbars accessible from x is at least

$$\left(m-\frac{k-1}{k}(h-1)\right)(N/k)$$

The number of middle stage crossbars is exactly mN/k. If $\alpha_h > mN/2k$ then more than half the middle stage switches are accessible from x. By a similar argument, more than half the middle stage switches are accessible from y and consequently there is at least one middle stage switch accessible from both x and y allowing us to connect them. Thus it suffices to have

$$\left(m-\frac{k-1}{k}(h-1)\right)(N/k) > mN/2k$$

which holds exactly when $m > 2\frac{k-1}{k}\log_k(N/k)$. □

The next theorem demonstrates that we can obtain a wide-sense nonblocking network by taking two Cantor networks and connecting the outputs of the first network to the inputs of the second. We refer to this as a *cascade connection.*

THEOREM 2. *The network obtained by cascading two Cantor networks, $K_{N,k,m}$ is a wide-sense nonblocking generalized connector if $m > 2\frac{k-1}{k}\log_k(N/k)$.*

Proof. The only constraint placed on the routing algorithm is that connections branch only in the second network. Under this restriction, we show that it is always possible to connect an idle input x simultaneously to any subset of the idle outputs.

First, we need some definitions. For $1 \leq i \leq h = \log_k N$, define $R_i(x)$ to be the set of stage i links that can be reached from input x in an idle Cantor network. We note that $R_i(x) = R_i(y)$ if $\lfloor x/k^{i-1} \rfloor = \lfloor y/k^{i-1} \rfloor$, so

$$R_i(0), R_i\left(k^{i-1}\right), \ldots, R_i\left(jk^{i-1}\right), \ldots, R_i\left(\left(k^{h-i+1} - 1\right)k^{i-1}\right)$$

partitions the links in stage i into k^{h-i+1} groups of mk^{i-1} links each.

Now consider any state of the cascaded Cantor networks with at least one idle output and in which connections branch only in the second network. Suppose we are asked to connect some idle input x to some subset of the idle outputs. Our strategy will be to find some input z of the *second network* from which more than half of the middle stage crossbars in the second network are accessible. We find z by working back from the middle stage of the second network, seeking the most "lightly loaded" portion of the network.

Define $\beta_i(j)$ to be the number of links in $R_i\left(jk^{i-1}\right)$ that are busy and let $\beta_i^* = \min_j \beta_i(j)$. Note that $\beta_i^* \leq (N-1)/k^{h-i+1} = k^{i-1}(1 - 1/N)$, since so long as there is at least one idle output, at most $N-1$ links in any stage can be busy and the number of distinct sets $R_i(\cdot)$ is k^{h-i+1}.

Now suppose that $1 < i \leq h$ and j satisfies $\beta_i(j) \leq (N-1)/k^{h-i+1}$. Then for some $j' \in [jk, \ldots, (j+1)k - 1]$ we have $\beta_i(j') \leq (N-1)/k^{h-i+2}$. Thus, we can work back from the middle stage of the second network to an input z such that for $1 \leq i \leq h$, the number of busy links in $R_i(z) \leq k^{i-1}(1 - 1/N)$. Since the number of busy links must be an integer, this implies $R_i(z) \leq k^{i-1} - 1$. Now, if α_i is the number of stage i links accessible from z then

$$\begin{aligned}
\alpha_1 &= m \\
\alpha_2 &\geq mk - (k-1) \\
\alpha_3 &\geq k\alpha_2 - (k^2 - k) \geq mk^2 - 2k^2 + 2k \\
\alpha_4 &\geq k\alpha_3 - (k^3 - k^2) \geq mk^3 - 3k^3 + 3k^2 \\
&\vdots \\
\alpha_i &\geq mk^{i-1} - (i-1)k^{i-1} + (i-1)k^{i-2}
\end{aligned}$$

Note that this is exactly the same inequality obtained in the previous theorem. Hence, at least half of the second network's middle stage crossbars are accessible from z. Since the branching in the second network does not affect the accessibility of the middle stage crossbars from the output side, each idle output can access more than half the middle stage crossbars implying that every idle output has at least one middle stage crossbar accessible to both it and z. This permits us to connect any subset of the idle outputs to z. Because the first Cantor network is strictly nonblocking for nonbranching connections, it is possible to connect any of its idle inputs to z, establishing the theorem. □

We note without proof that by a similar argument, a pair of cascaded Clos networks forms a wide-sense nonblocking generalized connector.

References

[1] Beneš, V. E. *"Mathematical Theory of Connecting Networks and Telephone Traffic,"* Academic Press, New York, 1965.

[2] Cantor, D. G. "On Non-Blocking Switching Networks", *Networks*, Vol.1, No.4, Winter 1971, pp. 367–377.

[3] Clos, C. "A Study of Non-blocking Switching Networks" *Bell Syst. Tech. J.*, Vol.32, Mar.1953, pp. 406–424.

[4] Feldman, P., J. Friedman and N. Pippenger "Non-Blocking Networks" *Proceedings of STOC 1986* 5/86, pp. 247–254.

[5] Masson, G. M., Gingher, G. C., Nakamura, S. "A Sampler of Circuit Switching Networks", *Computer*, June 1979, pp.145–161.

[6] Pippenger, N. "The Complexity of Switching Networks," MIT, PhD thesis, 1973.

[7] Richards, G. and F. K. Hwang "A Two Stage Rearrangeable Broadcast Switching Network," *IEEE Transactions on Communications*, 10/85.

PRACTICAL WIDE-SENSE NONBLOCKING GENERALIZED CONNECTORS

Jonathan S. Turner

WUCS-88-29

Department of Computer Science
Washington University
Campus Box 1045
One Brookings Drive
Saint Louis, MO 63130-4899

Abstract

In this note, we show that wide-sense nonblocking networks can be obtained by cascading a pair of Cantor networks or a pair of Clos networks. The only constraint placed on the routing algorithm is that branching be restricted to the second network in the cascade. This result yields practical network for multipoint communication with complexities $O(N(logN)^2)$ and $O(N^{1+1/r})$.

This work supported by the National Science Foundation (grant DCI 8600947), Bell Communications Research, Bell Northern Research, Italtel SIT and NEC.

Practical Wide-Sense Nonblocking Generalized Connectors

Jonathan S. Turner

A *network* can be modeled as a directed graph with a designated set of input nodes and a designated set of output nodes. A *connection network* or *connector* is a network used to connect designated pairs of inputs and outputs. The paths used to connect distinct pairs are not permitted to share links. Beneš [1], defines a connector to be strictly nonblocking if given any legal configuration of the network, it is always possible to add a connection between a specified idle input and a specified idle output, without reassigning existing connections to different paths. He also defines a connector to be *wide-sense nonblocking* is there is a routing algorithm which avoids blocking states given any legitimate sequence of connection and disconnection requests. A *generalized connector* is a network used to connect designated inputs to designated output sets; that is a network used to establish multipoint communication channels. The definitions of strict and wide-sense nonblocking extend to generalized connectors in the obvious way. See [1] and [5] for more precise definitions.

In this note, we show that wide-sense nonblocking networks can be obtained by cascading a pair of Cantor networks [2] or a pair of Clos networks [3]. The only constraint placed on the routing algorithm is that branching be restricted to the second network in the cascade. This result yields practical networks for multipoint communication with complexities $O\left(N(\log N)^2\right)$ and $O\left(N^{1+1/r}\right)$. Pippenger [6] describes a wide-sense nonblocking generalized connector with complexity $O\left(N(\log N)^3\right)$. This appears to be the best asymptotic complexity for a network with the potential for practical application. Feldman, Friedman and Pippenger [4] show how to construct wide-sense nonblocking generalized connectors with complexity $O(N \log N)$, but the resulting networks are of purely theoretical interest as the constants involved are quite large and the routing problem is NP-complete.

The Beneš network [1], $B_{N,k}$ can defined by the recursive construction shown in Figure 1. The base of the recursion is a $k \times k$ crossbar. The Cantor network $K_{N,k,m}$ can be defined by the construction shown in Figure 2. The following result is due to Cantor [2]. We include it here for reference.

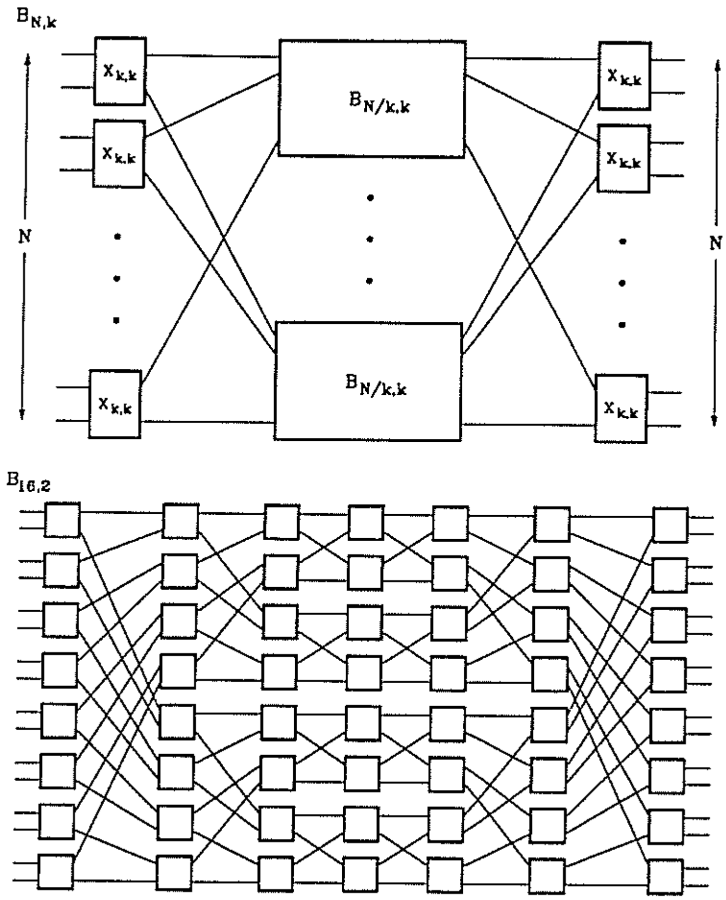

Figure 1: Beneš Network

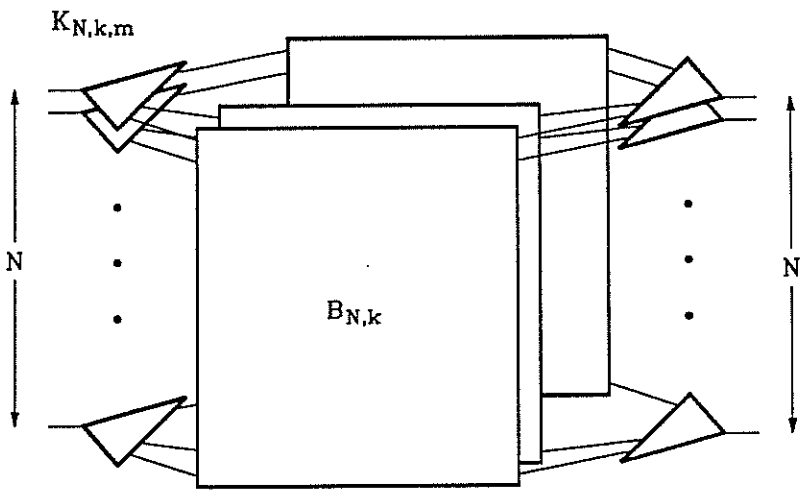

Figure 2: Cantor Network

THEOREM 1. *The Cantor network $K_{N,k,m}$ is a strictly nonblocking connector if $m > 2\frac{k-1}{k}\log_k(N/k)$.*

Proof. Assume the network is in some arbitrary state and we are asked to connect an idle input x to an idle output y. Define α_i to be the number of stage i links accessible from x. Then,

$$\begin{array}{rcl}\alpha_1 & = & m \\ \alpha_2 & \geq & mk-(k-1) \\ \alpha_3 & \geq & k\alpha_2-(k^2-k) \geq mk^2-2k^2+2k \\ \alpha_4 & \geq & k\alpha_3-(k^3-k^2) \geq mk^3-3k^3+3k^2 \\ & \vdots & \\ \alpha_i & \geq & mk^{i-1}-(i-1)k^{i-1}+(i-1)k^{i-2}\end{array}$$

Substituting $h = \log_k N$, we find that the number of stage h links and hence the number of middle stage crossbars accessible from x is at least

$$\left(m - \frac{k-1}{k}(h-1)\right)(N/k)$$

The number of middle stage crossbars is exactly mN/k. If $\alpha_h > mN/2k$ then more than half the middle stage switches are accessible from x. By a similar argument, more than half the middle stage switches are accessible from y and consequently there is at least one middle stage switch accessible from both x and y allowing us to connect them. Thus it suffices to have

$$\left(m - \frac{k-1}{k}(h-1)\right)(N/k) > mN/2k$$

which holds exactly when $m > 2\frac{k-1}{k}\log_k(N/k)$. □

The next theorem demonstrates that we can obtain a wide-sense nonblocking network by taking two Cantor networks and connecting the outputs of the first network to the inputs of the second. We refer to this as a *cascade connection.*

THEOREM 2. *The network obtained by cascading two Cantor networks, $K_{N,k,m}$ is a wide-sense nonblocking generalized connector if $m > 2\frac{k-1}{k}\log_k(N/k)$.*

Proof. The only constraint placed on the routing algorithm is that connections branch only in the second network. Under this restriction, we show that it is always possible to connect an idle input x simultaneously to any subset of the idle outputs.

First, we need some definitions. For $1 \le i \le h = \log_k N$, define $R_i(x)$ to be the set of stage i links that can be reached from input x in an idle Cantor network. We note that $R_i(x) = R_i(y)$ if $\lfloor x/k^{i-1} \rfloor = \lfloor y/k^{i-1} \rfloor$, so

$$R_i(0), R_i\left(k^{i-1}\right), \ldots, R_i\left(jk^{i-1}\right), \ldots, R_i\left(\left(k^{h-i+1}-1\right)k^{i-1}\right)$$

partitions the links in stage i into k^{h-i+1} groups of mk^{i-1} links each.

Now consider any state of the cascaded Cantor networks with at least one idle output and in which connections branch only in the second network. Suppose we are asked to connect some idle input x to some subset of the idle outputs. Our strategy will be to find some input z of the *second network* from which more than half of the middle stage crossbars in the second network are accessible. We find z by working back from the middle stage of the second network, seeking the most "lightly loaded" portion of the network.

Define $\beta_i(j)$ to be the number of links in $R_i\left(jk^{i-1}\right)$ that are busy and let $\beta_i^* = \min_j \beta_i(j)$. Note that $\beta_i^* \le (N-1)/k^{h-i+1} = k^{i-1}(1-1/N)$, since so long as there is at least one idle output, at most $N-1$ links in any stage can be busy and the number of distinct sets $R_i(\cdot)$ is k^{h-i+1}.

Now suppose that $1 < i \le h$ and j satisfies $\beta_i(j) \le (N-1)/k^{h-i+1}$. Then for some $j' \in [jk, \ldots, (j+1)k-1]$ we have $\beta_i(j') \le (N-1)/k^{h-i+2}$. Thus, we can work back from the middle stage of the second network to an input z such that for $1 \le i \le h$, the number of busy links in $R_i(z) \le k^{i-1}(1-1/N)$. Since the number of busy links must be an integer, this implies $R_i(z) \le k^{i-1} - 1$. Now, if α_i is the number of stage i links accessible from z then

$$\begin{aligned}
\alpha_1 &= m \\
\alpha_2 &\ge mk - (k-1) \\
\alpha_3 &\ge k\alpha_2 - (k^2 - k) \ge mk^2 - 2k^2 + 2k \\
\alpha_4 &\ge k\alpha_3 - (k^3 - k^2) \ge mk^3 - 3k^3 + 3k^2 \\
&\vdots \\
\alpha_i &\ge mk^{i-1} - (i-1)k^{i-1} + (i-1)k^{i-2}
\end{aligned}$$

Note that this is exactly the same inequality obtained in the previous theorem. Hence, at least half of the second network's middle stage crossbars are accessible from z. Since the branching in the second network does not affect the accessibility of the middle stage crossbars from the output side, each idle output can access more than half the middle stage crossbars implying that every idle output has at least one middle stage crossbar accessible to both it and z. This permits us to connect any subset of the idle outputs to z. Because the first Cantor network is strictly nonblocking for nonbranching connections, it is possible to connect any of its idle inputs to z, establishing the theorem. □

We note without proof that by a similar argument, a pair of cascaded Clos networks forms a wide-sense nonblocking generalized connector.

References

[1] Beneš, V. E. *"Mathematical Theory of Connecting Networks and Telephone Traffic,"* Academic Press, New York, 1965.

[2] Cantor, D. G. "On Non-Blocking Switching Networks", *Networks*, Vol.1, No.4, Winter 1971, pp. 367–377.

[3] Clos, C. "A Study of Non-blocking Switching Networks" *Bell Syst. Tech. J.*, Vol.32, Mar.1953, pp. 406–424.

[4] Feldman, P., J. Friedman and N. Pippenger "Non-Blocking Networks" *Proceedings of STOC 1986* 5/86, pp. 247–254.

[5] Masson, G. M., Gingher, G. C., Nakamura, S. "A Sampler of Circuit Switching Networks", *Computer*, June 1979, pp.145–161.

[6] Pippenger, N. "The Complexity of Switching Networks," MIT, PhD thesis, 1973.

[7] Richards, G. and F. K. Hwang "A Two Stage Rearrangeable Broadcast Switching Network," *IEEE Transactions on Communications*, 10/85.