Lemma 4-6: Let p be nonfaulty. Then $|ADJ^i_p| \leq (1 + \rho)(\beta + \varepsilon) + \rho\delta$.

Proof: $ADJ^i_p = T^i + \delta - AV^i_p$.

Thus, for some nonfaulty q and r, Lemma 4-5 implies that

$T^i + \delta - ARR^i_p(q) \leq ADJ^i_p \leq T^i + \delta - ARR^i_p(r)$.

Then Lemma 4-4 implies that:

(a) $ADJ^i_p \geq T^i + \delta - (T^i + (1 + \rho)(\beta + \delta + \varepsilon)) = -(1 + \rho)(\beta + \varepsilon) - \rho\delta$.

(b) If $\delta - \varepsilon \geq \beta$, then $ADJ^i_p \leq T^i + \delta - (T^i + (1 - \rho)(\delta - \varepsilon - \beta)) = (1 - \rho)(\beta + \varepsilon) + \rho\delta$.

(c) If $\delta - \varepsilon \leq \beta$, then $ADJ^i_p \leq T^i + \delta - (T^i - (1 + \rho)(\beta - \delta + \varepsilon)) = (1 + \rho)(\beta + \varepsilon) - \rho\delta$.

The conclusion is immediate. ∎

4.5.4 Timers Are Set in the Future

Earlier, we gave a lower bound on P and described two conditions which that bound was supposed to guarantee (that timers are set in the future and that messages arrive after the appropriate clocks have been set). In this subsection, we show that the given bound on P is sufficient to guarantee that the first of these two conditions holds.

Lemma 4-7: Let p be nonfaulty. Then $U^i + ADJ^i_p < T^{i+1}$.

Proof: $U^i + ADJ^i_p \leq U^i + (1 + \rho)(\beta + \varepsilon) + \rho\delta$, by Lemma 4-6

$= U^i + (2(1 + \rho)(\beta + \varepsilon) + (1 + \rho)\delta + \rho\delta) - (1 + \rho)(\beta + \delta + \varepsilon)$

$< U^i + P - (1 + \rho)(\beta + \delta + \varepsilon)$, by the assumed lower bound on P

$= T^{i+1}$. ∎

This lemma implies that timers are set in the future and that t^{i+1}_p is defined, the first of the three inductive properties which we must verify.

4.5.5 Bounding the Separation of Clocks

Next, we prove several lemmas which lead to bounds on the distance between the new clocks of nonfaulty processes. The first lemma gives an upper bound on the error in a process' estimate of the difference in real time between its own clock and another nonfaulty process' clock reaching T^i.

Lemma 4-8: Let p, q and r be nonfaulty. Then

$|(ARR^i_p(q) - (T^i + \delta)) - (c^i_q(T^i) - c^i_p(T^i))| \leq \varepsilon + \rho(\beta + \delta + \varepsilon)$.

Proof: Let a be the real time of arrival of q's message at process p. Then a is at most $c^i_q(T^i) + \delta + \varepsilon$. Define a new auxiliary clock, D, with rate exactly equal to 1, and such that $D(a) = C^i_p(a)$. Thus, $ARR^i_p(q) = D(a)$. So the expression we want to bound is at most equal to:

$$|(D(a) - (T^i + \delta)) - (c^i_q(T^i) - d(T^i))| + |c^i_p(T^i) - d(T^i)|.$$

First we demonstrate that the first of these two terms is at most ε.

$$|D(a) - (T^i + \delta) - c^i_q(T^i) + d(T^i)|$$

$$= |a - d(T^i + \delta) - c^i_q(T^i) + d(T^i)|, \text{ since D has rate 1}$$

$$= |a - c^i_q(T^i) + T^i - (T^i + \delta)|$$

$$\leq |c^i_q(T^i) + \delta + \varepsilon - c^i_q(T^i) - \delta|$$

$$= \varepsilon.$$

Next we show that the second term, $|c^i_p(T^i) - d(T^i)|$, is at most $\rho(\beta + \delta + \varepsilon)$.

Case 1: $c^i_p(T^i) \leq a$. So p reaches T^i before q's message arrives.

Let $\gamma = a - c^i_p(T^i)$. Then $\gamma \leq \beta + \delta + \varepsilon$.

Subcase 1a: $d(T^i) \geq c^i_p(T^i)$. So C_p has rate slower than real time.

Then $d(T^i) - c^i_p(T^i)$ is largest when C_p goes at the slowest possible rate, $1/(1 + \rho)$. In this case, $d(T^i) - c^i_p(T^i) = \gamma - (a - d(T^i))$, where $a - d(T^i) = \gamma/(1 + \rho)$. Thus, $d(T^i) - c^i_p(T^i) = \gamma(1 - 1/(1 + \rho)) = \gamma\rho/(1 + \rho) \leq \gamma\rho \leq \rho(\beta + \delta + \varepsilon)$.

Subcase 1b: $d(T^i) \leq c^i_p(T^i)$. So C_p has rate faster than real time.

Then $c^i_p(T^i) - d(T^i)$ is largest when C_p goes at the fastest possible rate, $1 + \rho$. Then $c^i_p(T^i) - d(T^i) = \gamma(1 + \rho) - \gamma = \gamma\rho \leq \rho(\beta + \delta + \varepsilon)$.

Case 2: $c^i_p(T^i) \geq a$. So p reaches T^i after q's message arrives.

Let $\gamma = c^i_p(T^i) - a$. Then $\gamma \leq \beta - \delta + \varepsilon$.

Subcase 2a: $d(T^i) \geq c^i_p(T^i)$. So C_p has rate faster than real time.

An argument similar to that for case 1b shows that $d(T^i) - c^i_p(T^i) \leq \gamma\rho \leq \rho(\beta - \delta + \varepsilon)$, which suffices.

Subcase 2b: $d(T^i) \leq c^i_p(T^i)$. So C_p has rate slower than real time.

An argument similar to that for case 1a shows that $c^i_p(T^i) - d(T^i) \leq \gamma\rho \leq \rho(\beta - \delta + \varepsilon)$,

which suffices. ∎

In order to prove the next lemma, we use some results about multisets, which are presented in the Appendix. This is a key lemma because the distance between the clocks is reduced from β to $\beta/2$, roughly. The halving is due to the properties of the fault-tolerant averaging function used in the algorithm. Consequently, the averaging function can be considered the heart of the algorithm.

Lemma 4-9: Let p and q be nonfaulty. Then

$|(c^i_p(T^i) - c^i_q(T^i)) - (ADJ^i_p - ADJ^i_q)| \leq \beta/2 + 2\varepsilon + 2\rho(\beta + \delta + \varepsilon)$.

Proof: We define multisets U, V, and W, and show they satisfy the hypotheses of Lemma A-4. Let

$U = c^i_p(T^i) - (T^i + \delta) + ARR^i_p$,

$V = c^i_q(T^i) - (T^i + \delta) + ARR^i_q$, and

$W = \{c^i_r(T^i)$: r is nonfaulty$\}$.

U and V have size n and W has size n – f.

Let $x = \varepsilon + \rho(\beta + \delta + \varepsilon)$.

Define an injection from W to U as follows. Map each element $c^i_r(T^i)$ in W to $c^i_p(T^i) - (T^i + \delta) + ARR^i_p(r)$ in U. Since Lemma 4-8 implies that $|(ARR^i_p(r) - (T^i + \delta)) - (c^i_r(T^i) - c^i_p(T^i))| \leq \varepsilon + \rho(\beta + \delta + \varepsilon)$ for all the elements of W, $d_x(W,U) = 0$. Similarly, $d_x(W,V) = 0$.

Since any two nonfaulty processes reach T^i within β real time of each other, diam(W) = β.

By Lemma A-4, $|mid(reduce(U)) - mid(reduce(V))| \leq \beta/2 + 2\varepsilon + 2\rho(\beta + \delta + \varepsilon)$.

Since $mid(reduce(U)) = mid(reduce(c^i_p(T^i) - (T^i + \delta) + ARR^i_p)) = c^i_p(T^i) - ADJ^i_p$, and similarly $mid(reduce(V)) = c^i_q(T^i) - ADJ^i_q$, the result follows. ∎

The next lemma is analogous to the previous one, except that it involves U^i instead of T^i.

Lemma 4-10: Let p and q be nonfaulty. Then

$|(c^i_p(U^i) - c^i_q(U^i)) - (ADJ^i_p - ADJ^i_q)| \leq \beta/2 + 2\varepsilon + 2\rho(2 + \rho)(\beta + \delta + \varepsilon)$.

Proof: The given expression is

$\leq |(c^i_p(T^i) - c^i_q(T^i)) - (ADJ^i_p - ADJ^i_q)| + |(c^i_p(U^i) - c^i_q(U^i)) - (c^i_p(T^i) - c^i_q(T^i))|$

$\leq \beta/2 + 2\varepsilon + 2\rho(\beta + \delta + \varepsilon) + 2\rho(1 + \rho)(\beta + \delta + \varepsilon)$, by Lemmas 4-9 and 4-2.

This reduces to the claimed expression. ∎

Next we bound the distance in real time between two nonfaulty processes switching to their new clocks. It is crucial that the distance between the new clocks reaching U^i be less than β in order to accommodate their relative drift during the interval between U^i and T^{i+1}.

Lemma 4-11: Let p, q be nonfaulty. Then

$|c^{i+1}_p(U^i) - c^{i+1}_q(U^i)| \leq \beta/2 + 2\varepsilon + 2\rho(3\beta + 2\delta + 3\varepsilon) + 4\rho^2(\beta + \delta + \varepsilon)$.

Proof: We define idealized clocks, D_p and D_q, as follows. Both have rate exactly 1. Also, $D_p(u^i_p) = C^{i+1}_p(u^i_p) = U^i + ADJ^i_p$, and similarly for q. Then

$|c^{i+1}_p(U^i) - c^{i+1}_q(U^i)| \leq |c^{i+1}_p(U^i) - d_p(U^i)| + |d_p(U^i) - d_q(U^i)| + |d_q(U^i) - c^{i+1}_q(U^i)|$.

We bound each of these three terms separately.

First, consider $|c^{i+1}_p(U^i) - d_p(U^i)|$. Now, $U^i + ADJ^i_p = D_p(u^i_p) = C^{i+1}_p(u^i_p)$. So

$|c^{i+1}_p(U^i) - d_p(U^i)| \leq |(c^{i+1}_p(U^i) - d_p(U^i)) - (c^{i+1}_p(U^i + ADJ^i_p) - d_p(U^i + ADJ^i_p))|$

$\leq \rho|ADJ^i_p|$, by Lemma 4-2

$\leq \rho((1 + \rho)(\beta + \varepsilon) + \rho\delta)$, by Lemma 4-6.

The same bound holds for the third term.

Finally, consider the middle term, $|d_p(U^i) - d_q(U^i)|$. We know that $d_p(U^i) = d_p(U^i + ADJ^i_p) - ADJ^i_p = u^i_p - ADJ^i_p$, and similarly for q.

$|d_p(U^i) - d_q(U^i)| = |(u^i_p - u^i_q) - (ADJ^i_p - ADJ^i_q)|$

$\leq \beta/2 + 2\varepsilon + 2\rho(2 + \rho)(\beta + \delta + \varepsilon)$, by Lemma 4-10.

Combining these three bounds, we get the required bound. ∎

Finally, we can show the second of our inductive properties, bounding the distance between times when clocks reach T^{i+1}.

Lemma 4-12: Let p, q be nonfaulty. Then $|t^{i+1}_p - t^{i+1}_q| \leq \beta$.

Proof: $|t^{i+1}_p - t^{i+1}_q|$

$= |c^{i+1}_p(T^{i+1}) - c^{i+1}_q(T^{i+1})|$

$\leq |(c^{i+1}_p(T^{i+1}) - c^{i+1}_q(T^{i+1})) - (c^{i+1}_p(U^i) - c^{i+1}_q(U^i))| + |c^{i+1}_p(U^i) - c^{i+1}_q(U^i)|$

$\leq 2\rho(P - (1 + \rho)(\beta + \delta + \varepsilon)) + \beta/2 + 2\varepsilon + 2\rho(3\beta + 2\delta + 3\varepsilon) + 4\rho^2(\beta + \delta + \varepsilon)$, by Lemmas 4-2 and 4-11.

The assumed upper bound on P implies that this expression is at most β. ∎

4.5.6 Bound on Message Arrival Time

In this subsection, we show that the third and final inductive assumption holds. That is, we show that messages arrive after the appropriate clocks have been set.

Lemma 4-13: Let p and q be nonfaulty. Then $t^{i+1}_q + \delta - \varepsilon > u^i_p$.

Proof: Since $t^{i+1}_q + \delta - \varepsilon \geq t^{i+1}_p - \beta + \delta - \varepsilon$, it suffices to show that

$t^{i+1}_p - u^i_p > \beta - \delta + \varepsilon$.

Now, $t^{i+1}_p - u^i_p \geq (P - (1 + \rho)(\beta + \delta + \varepsilon) - ADJ^i_p)/(1 + \rho)$ since the numerator represents the smallest possible difference in the values of the clock C^{i+1}_p at the two given real times.

But the lower bound on P implies that $P > 3(1 + \rho)(\beta + \varepsilon) + \rho\delta$. Also, the bound on the adjustment shows that $ADJ^i_p \leq (1 + \rho)(\beta + \varepsilon) + \rho\delta$. Therefore,

$t^{i+1}_p - u^i_p > (3(1 + \rho)(\beta + \varepsilon) + \rho\delta - (1 + \rho)(\beta + \delta + \varepsilon) - (1 + \rho)(\beta + \varepsilon) - \rho\delta) / (1 + \rho)$

$= \beta - \delta + \varepsilon$, as needed. ∎

Thus, we have shown that the three inductive hypotheses hold. Therefore, the claims made in this section for a particular i, in fact hold for all i.

4.6 Some General Properties

In this section, we state several consequences of the results proved in the preceding section.

First, we state a bound on the closeness with which the various clocks reach corresponding values.

Lemma 4-14: Let p, q be nonfaulty, $i \geq 0$. Assume that T is chosen so that $U^{i-1} \leq T \leq U^i$, if $i \geq 1$, or so that $T^0 \leq T \leq U^0$, if $i = 0$.

Then $|c^i_p(T) - c^i_q(T)| \leq \beta + 2\rho(1 + \rho)(\beta + \delta + \varepsilon)$.

Proof: *Basis:* $i = 0$. Then $T^0 \leq T \leq U^0$.

$|c^0_p(T) - c^0_q(T)| \leq |(c^0_p(T) - c^0_q(T)) - (c^0_p(T^0) - c^0_q(T^0))| + |c^0_p(T^0) - c^0_q(T^0)|$

$\leq 2\rho(T - T^0) + \beta$, by Lemma 4-2 and assumption 3

$\leq \beta + 2\rho(1 + \rho)(\beta + \delta + \varepsilon)$.

Induction: $i \geq 0$. Choose T with $U^{i-1} \leq T \leq U^i$.

$$|c^i_p(T) - c^i_q(T)| \leq |(c^i_p(T) - c^i_q(T)) - (c^i_p(U^{i-1}) - c^i_q(U^{i-1}))| + |c^i_p(U^{i-1}) - c^i_q(U^{i-1})|$$

$$\leq 2\rho P + \beta/2 + 2\varepsilon + 2\rho(3\beta + 2\delta + 3\varepsilon) + 4\rho^2(\beta + \delta + \varepsilon), \text{ by Lemmas 4-2 and 4-11.}$$

The upper bound on P implies the result. ∎

Next, we prove a bound for a nonfaulty process' (i + 1)-st clock, in terms of nonfaulty processes' i-th clocks.

Lemma 4-15: Let p be nonfaulty, $i \geq 0$. Then there exist nonfaulty processes, q and r, such that for $u^i_p \leq t \leq \text{umax}^i$,

$$C^i_q(t) - \alpha \leq C^{i+1}_p(t) \leq C^i_r(t) + \alpha,$$

where $\alpha = \varepsilon + \rho(4\beta + \delta + 5\varepsilon) + 4\rho^2(\beta + \delta + \varepsilon) + 2\rho^3(\beta + \delta + \varepsilon)$.

Proof: $C^{i+1}_p(t) = C^i_p(t) + T^i + \delta - AV^i_p$. Therefore, by Lemma 4-5 there are nonfaulty processes, q and r, for which

$$C^i_p(t) + T^i + \delta - ARR^i_p(q) \leq C^{i+1}_p(t) \leq C^i_p(t) + T^i + \delta - ARR^i_p(r).$$

We show the right-hand inequality first. Let $a = c^i_p(ARR^i_p(r))$, the real time at which the message arrives at p from r. Thus, $C^i_p(a) = ARR^i_p(r)$. Note that $C^i_r(a) \geq T^i + (1 - \rho)(\delta - \varepsilon)$.

$$C^{i+1}_p(t) \leq C^i_p + T^i + \delta - ARR^i_p(r), \text{ from above}$$

$$\leq C^i_r(t) + C^i_p(a) - C^i_r(a) + T^i + \delta - ARR^i_p(r) + (C^i_p(t) - C^i_r(t)) - (C^i_p(a) - C^i_r(a))$$

$$\leq C^i_r(t) + C^i_p(a) - C^i_r(a) + T^i + \delta - ARR^i_p(r) + 2\rho(t - a), \text{ by Lemma 4-2 since } t > a$$

$$\leq C^i_r(t) + ARR^i_p(r) - T^i - (1 - \rho)(\delta - \varepsilon) + T^i + \delta - ARR^i_p(r) + 2\rho(t - a)$$

$$= C^i_r(t) + \varepsilon + \rho\delta - \rho\varepsilon + 2\rho(t - a).$$

It remains to bound t – a. The worst case occurs when $t = \text{umax}^i$. The longest possible elapsed real time between a particular nonfaulty process reaching T^i and U^i on the same clock is $(1 + \rho)^2(\beta + \delta + \varepsilon)$. Thus, $\text{umax}^i - \text{tmin}^i \leq \beta + (1 + \rho)^2(\beta + \delta + \varepsilon)$. But $a \geq \text{tmin}^i + \delta - \varepsilon$. Therefore, $t - a \leq \beta + (1 + \rho)^2(\beta + \delta + \varepsilon) - \delta + \varepsilon$

Thus, $C^{i+1}_p(t) \leq C^i_r(t) + \varepsilon + \rho\delta - \rho\varepsilon + 2\rho(\beta + (1 + \rho)^2(\beta + \delta + \varepsilon) - \delta + \varepsilon)$

$$= C^i_r(t) + \varepsilon + \rho(4\beta + \delta + 3\varepsilon) + 4\rho^2(\beta + \delta + \varepsilon) + 2\rho^3(\beta + \delta + \varepsilon)$$

$$< C^i_r(t) + \alpha.$$

For the left-hand inequality, we see that $C^i_q(t) - \varepsilon - \rho\delta - \rho\varepsilon - 2\rho(t - a) \leq C^{i+1}_p(t)$, where $a = c^i_p(ARR^i_p(q))$. The factor t – a is bounded exactly as before, so that we obtain:

$C^i_q(t) - \alpha \leq C^{i+1}_p(t)$. ∎

4.7 Agreement and Validity Conditions

We are now ready to show that the agreement and validity properties hold. The main effort is in restating bounds proved earlier concerning the closeness in real times when clocks reach the same value, in terms of the closeness of clock values at the same real time.

4.7.1 Agreement

The first lemma implies that the local times of two nonfaulty processes are close in those intervals where both use a clock with the same index.

Lemma 4-16: Let p, q be nonfaulty. Then

$$|C^i_p(t) - C^i_q(t)| \leq (1 + \rho)(\beta + 2\rho(1 + \rho)(\beta + \delta + \varepsilon))$$

for $\max\{u^{i-1}_p, u^{i-1}_q\} \leq t \leq \max\{u^i_p, u^i_q\}$, if $i \geq 1$,

and for $\min\{t^0_p, t^0_q\} \leq t \leq \max\{u^0_p, u^0_q\}$, if $i = 0$.

Proof: *Basis:* $i = 0$. Lemma 4-14 implies that

$$|c^i_p(T) - c^i_q(T)| \leq \beta + 2\rho(1 + \rho)(\beta + \delta + \varepsilon)$$

for all T, $U^{i-1} \leq T \leq U^i$ if $i \geq 1$ and for all T, $T^0 \leq T \leq U^0$ if $i = 0$. Then Lemma 4-3 immediately implies the needed result for $i = 0$.

Induction: $i \geq 1$. Lemma 4-3 implies the result for all t with

$$\min\{c^i_p(U^{i-1}), c^i_q(U^{i-1})\} \leq t \leq \max\{u^i_p, u^i_q\}.$$

It remains to show the bound for t with

$$\max\{u^{i-1}_p, u^{i-1}_q\} \leq t < \min\{c^i_p(U^{i-1}), c^i_q(U^{i-1})\}.$$

Without loss of generality, assume that $c^i_p(U^{i-1}) \leq c^i_q(U^{i-1})$, so that the minimum is equal to $c^i_p(U^{i-1})$.

$$|C^i_p(t) - C^i_q(t)| \leq |(C^i_p(t) - C^i_q(t)) - (C^i_p(c^i_p(U^{i-1})) - C^i_q(c^i_p(U^{i-1})))|$$
$$+ |C^i_p(c^i_p(U^{i-1})) - C^i_q(c^i_p(U^{i-1}))|$$

The first term, by Lemma 4-2, is at most $2\rho(c^i_p(U^{i-1}) - t)$. Since $t \geq \max\{u^{i-1}_p, u^{i-1}_q\} \geq u^{i-1}_p \geq c^{i-1}_p(U^{i-1})$, we have

$$2\rho(c^i_p(U^{i-1}) - t) \leq 2\rho(c^i_p(U^{i-1}) - c^{i-1}_p(U^{i-1})).$$

Since $c^{i-1}_p(U^{i-1}) = c^i_p(T)$ for some T with $|T - U^{i-1}| \leq |ADJ^i_p|$, this quantity is

$\leq 2\rho|c^i_p(U^{i-1}) - c^i_p(T)|$

$\leq 2\rho(1 + \rho)|U^{i-1} - T|$, by Lemma 4-1

$\leq 2\rho(1 + \rho)|ADJ^i_p|$

$\leq 2\rho(1 + \rho)((1 + \rho)(\beta + \varepsilon) + \rho\delta)$, by Lemma 4-6.

To bound the second term we note that Lemma 4-11 implies that

$|c^i_p(U^{i-1}) - c^i_q(U^{i-1})| \leq \beta/2 + 2\varepsilon + 2\rho(3\beta + 2\delta + 3\varepsilon) + 4\rho^2(\beta + \delta + \varepsilon) = \alpha$,

and so Lemma 4-3, with $T_1 = T_2 = U^{i-1}$, implies that

$|C^i_p(c^i_p(U^{i-1})) - C^i_q(c^i_p(U^{i-1}))| \leq (1 + \rho)\alpha$.

The assumed lower bound on β gives the result that

$2\rho(1 + \rho)((1 + \rho)(\beta + \varepsilon) + \rho\delta) + (1 + \rho)\alpha \leq (1 + \rho)(\beta + 2\rho(1 + \rho)(\beta + \delta + \varepsilon))$ ∎

Here is the main result, bounding the error in the synchronization at any time.

Theorem 4-17: The algorithm guarantees γ-agreement,

where $\gamma = \beta + \varepsilon + \rho(7\beta + 3\delta + 7\varepsilon) + 8\rho^2(\beta + \delta + \varepsilon) + 4\rho^3(\beta + \delta + \varepsilon)$.

Proof: The result for intervals in which the processes use clocks with the same indices has been covered in the preceding lemma. The expression in the statement of that lemma simplifies to

$\beta + \rho(3\beta + 2\delta + 2\varepsilon) + 4\rho^2(\beta + \delta + \varepsilon) + 2\rho^3(\beta + \delta + \varepsilon)$,

which is less than γ.

Next, we must consider the case where one of the processes has changed to a new clock, while the other still retains the old clock. Consider $|C^{i+1}_p(t) - C^i_q(t)|$ for some t with $u^i_p \leq t \leq u^i_q$. Lemma 4-15 implies that there exist nonfaulty processes r and s such that

$C^i_r(t) - \alpha \leq C^{i+1}_p(t) \leq C^i_s(t) + \alpha$,

where $\alpha = \varepsilon + \rho(4\beta + \delta + 5\varepsilon) + 4\rho^2(\beta + \delta + \varepsilon) + 2\rho^3(\beta + \delta + \varepsilon)$.

$|C^{i+1}_p(t) - C^i_q(t)| \leq \alpha + \max\{|C^i_r(t) - C^i_q(t)|, |C^i_s(t) - C^i_q(t)|\}$

$\leq \alpha + (1 + \rho)(\beta + 2\rho(1 + \rho)(\beta + \delta + \varepsilon))$, by the preceding lemma

$= \beta + \varepsilon + \rho(7\beta + 3\delta + 7\varepsilon) + 8\rho^2(\beta + \delta + \varepsilon) + 4\rho^3(\beta + \delta + \varepsilon)$, as needed. ∎

In some applications, it may never be the case that clocks with different indices are compared, perhaps because use of the clocks for processing ceases during the interval in which confusion is possible. In that case, the closeness of synchronization achieved by Algorithm 4-1 is given by Lemma 4-16, and is approximately $\beta + \rho(3\beta + 2\delta + 2\varepsilon)$. This value is more than ε less than the bound obtained when clocks with different indices must be compared.

Now we can sketch why it is reasonable for β to be approximately $4\varepsilon + 4\rho P$, as mentioned at the end of Section 4.5.1. Assume P is fixed. The i-th clocks reach T^i within β of each other. After the processes reset their clocks, the new clocks reach U^i within $\beta/2 + 2\varepsilon$ (ignoring ρ terms). By the end of the round, the clocks reach T^{i+1} within about $\beta/2 + 2\varepsilon + 2\rho P$ of each other, because of drift. This quantity must be at most β. The inequality $\beta/2 + 2\varepsilon + 2\rho P \leq \beta$ yields $\beta \geq 4\varepsilon + 4\rho P$.

Suppose we alter the algorithm so that during each round, the processes exchange clock values k times instead of just once. Then we get $\beta/2^k + (4 - 2^{2-k})\varepsilon + 2\rho P \leq \beta$, which simplifies to $\beta \geq 4\varepsilon + 2\rho P(2^k/(2^k-1))$. It appears that $\beta \geq 4\varepsilon + 2\rho P$ is approachable.

If the number of processes, n, increases while f, the number of faulty processes remained fixed, a greater closeness of synchronization can be achieved by modifying Algorithm 4-1 so that it computes the mean instead of the midpoint of the range of values.

As in [1], we show that the convergence rate of algorithms that use the mean instead of the midpoint is roughly f/(n–2f).

The result is based on the following lemma concerning multisets.

Lemma 4-18: Let U, V, and W be multisets such that $|U| = |V| = n \geq 3f + 1$ and $|W| = n - f$. If $d_x(W,U) = d_x(W,V) = 0$, then

$$|\text{mean}(\text{reduce}(U)) - \text{mean}(\text{reduce}(V))| \leq \text{diam}(W)f/(n-2f) + 2x.$$

The analysis of the modified Algorithm 4-1 parallels that just presented. However, the upper bound on P becomes

$$P \leq \beta(n-3f)/(n-2f)2\rho - \varepsilon/\rho - \rho(\beta + \delta + \varepsilon) - 2\beta - \delta - 2\varepsilon.$$

This bound implies $\beta \geq 2(n-2f)(\varepsilon + \rho P)/(n-3f)$, which approaches $\beta \geq 2\varepsilon + 2\rho P$ as n approaches infinity.

We now demonstrate that this bound is reasonable. After updating the clock and then waiting until the clocks reach the next T^i, the clocks must still be within β, giving $f\beta/(n-2f) + 2\varepsilon + 2\rho P \leq$

β, which implies $\beta \geq (2\varepsilon + 2\rho P)(n-2f)/(n-3f)$, which approaches $2\varepsilon + 2\rho P$ as n approaches infinity.

4.7.2 Validity

Next, we show the validity condition. The first lemma bounds the values of the zero-index clocks.

Lemma 4-19: $T^0 + (1-\rho)(t - t^0_p) \leq C^0_p(t) \leq T^0 + (1+\rho)(t - t^0_p)$ for $t \geq t^0_p$.

Proof: By Lemma 4-1. ∎

The next lemma is the main one.

Lemma 4-20: Let p be nonfaulty, $i \geq 0$. Then

$$(1-\rho)(t - tmax^0) + T^0 - i\varepsilon \leq C^i_p(t) \leq (1+\rho)(t - tmin^0) + T^0 + i\varepsilon$$

for all $t \geq u^{i-1}_p$ if $i \geq 1$, and for all $t \geq t^0_p$ if $i = 0$.

Proof: We proceed by induction on i. When proving the result for i + 1, we will assume the result for i, for all executions of the algorithm (rather than just the execution in question).

Basis: i = 0. This case follows immediately by Lemma 4-19.

Induction: Assume the result has been shown for i and show it for i + 1.

We argue the right-hand inequality first. The left-hand inequality is entirely analogous.

Assume in contradiction that we have a particular execution in which $C^{i+1}_p(t) > (1+\rho)(t - tmin^0) + T^0 + (i+1)\varepsilon$ for some $t \geq u^i_p$. Then by the limitations on rates of clocks, it is clear that $C^{i+1}_p(u^i_p) > (1+\rho)(u^i_p - tmin^0) + T^0 + (i+1)\varepsilon$.

Recall that p resets its clock at real time u^i_p, by adding $T^i + \delta - AV^i_p$. In this case, the inductive hypothesis implies that the adjustment must be an increment.

By Lemma 4-5, this increment is $\leq T^i + \delta - ARR^i_p(q)$ for some nonfaulty q. Therefore,

$$C^i_p(u^i_p) + T^i + \delta - ARR^i_p(q) > (1+\rho)(u^i_p - tmin^0) + T^0 + (i+1)\varepsilon.$$

Next, we claim that if p had done the adjustment just when the message arrived from q rather than waiting till real time u^i_p, the bound would still have been exceeded. That is, $ARR^i_p(q) + T^i + \delta - ARR^i_p(q) > (1+\rho)(t' - tmin^0) + T^0 + (i+1)\varepsilon$, where $t' = c^i_p(ARR^i_p(q))$. (This again follows by the limits on the rates of clocks.) Thus,

$$T^i + \delta > (1+\rho)(t' - tmin^0) + T^0 + (i+1)\varepsilon.$$

Now consider an alternative execution of the algorithm in which everything is exactly like the one we have been describing, except that immediately after q sends out clock reading T^i, q's clock C^i_q begins to move at rate 1. This change cannot affect p's (i+1)-st clock because q doesn't send any more messages until t^{i+1}_q, and these