10 Common knowledge

A finite-state system is a system where each agent has only a finite number of local states.

Theorem 10.1. *Let $\mathcal{I}$ be an interpreted finite-state system, and let G be a nonempty set of agents. Then, there is a k such that for every formula φ,*

$$\mathcal{I} \models C_G\varphi \Leftrightarrow E_G^k\varphi.$$

Proof. Let L_i be the set of local states of the agent i. Define $l = \min\{|L_i| \mid i \in G\}$, and set $k = 2l - 1$. Assume $(\mathcal{I}, r, m) \not\models C_G\varphi$. Then, there is a point (r', m') that is G-reachable from (r, m) for which $(\mathcal{I}, r', m') \not\models \varphi$. Let

$$(r, m) = (r_0, m_0) \sim_{i_0} (r_1, m_1) \sim_{i_1} \cdots \sim_{i_{t-1}} (r_t, m_t) = (r', m'), \quad i_j \in G$$

be a path of minimal length that connects (r, m) with (r', m'). Then, the minimal length implies

$$(r_a, m_a) \not\sim_i (r_b, m_b)$$

for all $i \in G$ and for $0 < a + 1 < b \leq t$. Hence, $t \leq 2l - 1 = k$, and consequently $(\mathcal{I}, r, m) \not\models E_G^k\varphi$. We have proved $\mathcal{I} \models E_G^k\varphi \Rightarrow C_G\varphi$. As a result, $\mathcal{I} \models C_G\varphi \Leftrightarrow E_G^k\varphi$. □

10.1 Simultaneity

Let A and B be two agents in a system where message delivery is reliable, and messages are either delivered instantly or after a delay of ϵ. Assume that A and B share a global clock and assume that A sends a message that do not specify the time to B at some point in time. Then, A and B will never reach common knowledge of the fact that a message have been sent.

Formally, let the local states of A be of the form (t, s) and $(t, \neg s)$, where t is the time and and s means that A has sent a message and $\neg s$ means that A has not sent a message. Similarly, let the local states of B be of the form (t, r) and $(t, \neg r)$ representing whether B has received the message or not. Then,

$$\langle (t, s), (t, r) \rangle \sim_A \langle (t, s), (t, \neg r) \rangle \sim_B \langle (t, \neg s), (t, \neg r) \rangle$$

and $\langle (t, \neg s), (t, \neg r) \rangle$ is $\{A, B\}$-reachable from $\langle (t, s), (t, r) \rangle$. Hence, there can never be common knowledge among A and B of the fact that a message has been sent.

If A sends the time to B, then the situation changes. Let the local states of A be of the form $(t, s_{t'})$ and $(t, \neg s)$, where $s_{t'}$ expresses that A sent a message to B at time t'. Analogously, let the local states of B be of the form $(t, r_{t'})$ and $(t, \neg r)$. Then, $\langle (t+\epsilon, s_t), (t+\epsilon, r_{t+\epsilon})\rangle$ is the only state $\{A, B\}$-reachable from $\langle (t+\epsilon, s_t), (t+\epsilon, r_t)\rangle$. Thus, if the message is sent at time t it is common knowledge that the message has been sent at time $t+\epsilon$. Notice that A and B achieves common knowledge simultaneously at time $t+\epsilon$.

10.1.1 Simultaneous events

Fix an interpreted system $\mathcal{I} = (\mathcal{R}, \pi)$, and let S be the set of points of the system. Take an event in $\mathcal{R}$ to be a subset of S. We say that an event e holds at a point (r, m) if $(r, m) \in e$.

We are particularly interested in events whose occurrence depends only on the local states of an agent. For instance, the event $send(\mu, j, i)$ is an event local to i. Formally, an event e is local to i if there is a set L_i^e of the local states of i such that

$$(r, m) \in e \qquad \text{iff} \qquad r_i(m) \in L_i^e.$$

Certain events depends only on the global state. An event e is a state event if there is a set $\mathcal{G}^e$ such that

$$(r, m) \in e \qquad \text{iff} \qquad r(m) \in \mathcal{G}^e.$$

If e is a state event then we define a proposition ψ_e such that $\pi\big(r(m)\big)(\psi_e) = \mathit{true}$ if and only if $(r, m) \in e$. Similarly, for every formula φ, we define an event

$$ev_{\mathcal{I}}(\varphi) = \{(r, m) \mid (\mathcal{I}, r, m) \models \varphi\}.$$

Lemma 10.2. *An event e is local to an agent i in a system $\mathcal{I}$ if and only if*

$$\mathcal{I} \models \psi_e \Rightarrow K_i \psi_e.$$

Similarly,

$$\mathcal{I} \models \varphi \Rightarrow K_i \varphi$$

if and only if $ev_{\mathcal{I}}(\varphi)$ is local to i.

Proof. Assume that e is local to i. If $(\mathcal{I}, r, m) \models \psi_e$, then $(r, m) \in e$ and for all (r', m') such that $r'(m') = r(m)$, $(r', m') \in e$. Hence, $(\mathcal{I}, r, m) \models K_i \psi_e$. Conversely, if $(\mathcal{I}, r, m) \models \psi_e \Rightarrow K_i \psi_e$ and $(r, m) \in e$, then $(r', m') \in e$ for all (r', m') such that $r'(m') = r(m)$. Hence, e is local to i.

To prove the other part, substitute $ev_{\mathcal{I}}(\varphi)$ for e and φ for ψ_e. □

An ensemble for G is a mapping $e : i \mapsto e_i$ that assigns to every $i \in G$ an event e_i local to i. An ensemble e for G is perfectly coordinated if the local events hold simultaneously, i.e., if $e(i) = e(j)$ for all $i, j \in G$. For an ensemble e, let $\psi_e = \bigvee_{i \in G} \psi_{e(i)}$.

Proposition 10.3. *Let $\mathcal{I}$ be an interpreted system and let G be a set of agents. Then,*

1. *for every formula φ, the ensemble e for G defined by $e(i) = ev_{\mathcal{I}}(K_i C_G \varphi)$ is perfectly coordinated,*
2. *if e is a perfectly coordinated ensemble for G, then $\mathcal{I} \models \psi_{e(i)} \Rightarrow K_i C_G \psi_{e(i)}$ for all $i \in G$, and*
3. *if e is a perfectly coordinated ensemble for G, then $\mathcal{I} \models \psi_e \Rightarrow C_G \psi_e$.*

Proof. Pick $i, j \in G$. Since

$$\models C_G\varphi \Leftrightarrow K_i C_G \varphi \qquad \text{and} \qquad \models C_G\varphi \Leftrightarrow K_j C_G \varphi,$$

$$ev_{\mathcal{I}}(C_G\varphi) = ev_{\mathcal{I}}(K_i C_G\varphi) = ev_{\mathcal{I}}(K_j C_G \varphi)$$

and it follows that the ensemble $e : i \mapsto ev_{\mathcal{I}}(K_i C_G \varphi)$ is perfectly coordinated.

If e is a perfectly coordinated ensemble for G, then $e(i)$ is local to i, and $\mathcal{I} \models \psi_{e(i)} \Rightarrow K_i \psi_{e(i)}$ for all $i \in G$. Moreover, $\mathcal{I} \models \psi_{e(i)} \Leftrightarrow \psi_{e(j)}$ for $i, j \in G$. Hence, $\mathcal{I} \models \psi_{e(i)} \Rightarrow E_G \psi_{e(i)}$ and by the Induction Rule $\mathcal{I} \models \psi_{e(i)} \Rightarrow C_G \psi_{e(i)}$. From the Fixed Point Axiom it now follows $\mathcal{I} \models \psi_{e(i)} \Rightarrow K_i C_G \psi_{e(i)}$.

Assume $(\mathcal{I}, r, m) \models \psi_e$. Then, $(\mathcal{I}, r, m) \models \psi_{e(i)}$ for some $i \in G$, and $(\mathcal{I}, r, m) \models K_i C_G \psi_{e(i)}$ follows. Since $\mathcal{I} \models K_i C_G \psi_{e(i)} \Rightarrow C_G \psi_{e(i)}$ and since $\mathcal{I} \models \psi_{e(i)} \Rightarrow \psi_e$, we get $(\mathcal{I}, r, m) \models C_G \psi_e$. □

In other words, the local events that correspond to common knowledge are perfectly coordinated, and the local events in a perfectly coordinated ensemble are common knowledge when they hold.

10.2 Temporal imprecision

A system $\mathcal{R}$ is said to have temporal imprecision if for all runs $r \in \mathcal{R}$, times m, and sets of processes G with $|G| \geq 2$, there exist $i, j \in G$, a run r', and a time m' such that

$$r'_i(m') = r_i(m) \qquad \text{and} \qquad r'_j(m') = r_j(m+1).$$

That is, in a system with clocks process i is at some time uncertain about the clock reading of process j.

An ensemble e for G in $\mathcal{I}$ is nontrivial if there is a run $r \in \mathcal{I}$ and if there are times m and m' such that

$$(r, m) \in \bigcup_{i \in G} e(i) \qquad \text{and} \qquad (r, m') \notin \bigcup_{i \in G} e(i).$$

Proposition 10.4. *In a system with temporal imprecision there are no nontrivial perfectly coordinated ensembles for G when $|G| \geq 2$.*

Proof. Assume that e is a nontrivial perfectly coordinated ensemble for G in a system $\mathcal{I}$ with temporal imprecision. Since e is nontrivial there is a run $r \in \mathcal{I}$ and there are times m and m' such that

$$(r, m) \in \bigcup_{i \in G} e(i) \qquad \text{and} \qquad (r, m') \notin \bigcup_{i \in G} e(i).$$

Suppose $m' > m$. Then, there is a time k such that

$$(r, k) \in \bigcup_{i \in G} e(i) \qquad \text{and} \qquad (r, k+1) \notin \bigcup_{i \in G} e(i),$$

and it follows that there are agents $j, j' \in G$ and a point (r', k') such that

$$r'_j(k') = r_j(k) \qquad \text{and} \qquad r'_{j'}(k') = r_{j'}(k+1).$$

Since e is perfectly coordinated, $\bigcup_{i \in G} e(i) = e(j)$. Hence, $r_j(k) = r'_j(k')$ and $(r, k) \in \bigcup_{i \in G} e(i)$ implies

$$(r', k') \in e(j) = e(j')$$

as $e(j)$ is local to j. But then, $r'_{j'}(k') = r_{j'}(k+1)$ implies $(r, k+1) \in e(j')$, which contradicts $(r, k+1) \notin \bigcup_{i \in G} e(i)$. The case $m' < m$ is handled similarly. ◻

Corollary 10.5. *Let $\mathcal{I}$ be an interpreted system with temporal imprecision, let φ be a formula, and let $|G| \geq 2$. Then, for all r and m,*

$$(\mathcal{I}, r, m) \models C_G\varphi \qquad \textit{iff} \qquad (\mathcal{I}, r, 0) \models C_G\varphi.$$

Proof. By Proposition 10.3, the ensemble e for G defined by

$$e(i) = ev_{\mathcal{I}}(K_i C_G \varphi)$$

is perfectly coordinated. By Proposition 10.4, it is trivial. Hence,

$$(r, m) \in \bigcup_{i \in G} e(i) \qquad \text{iff} \qquad (r, 0) \in \bigcup_{i \in G} e(i)$$

and the corollary follows. ◻

10.3 Common knowledge as a fixed point

Common knowledge of φ can be seen as a situation in which every agent knows φ and every agent knows that they all are in this situation. Thus, $C_G\varphi$ can be seen as a solution to the equation $x = E_G(\varphi \wedge x)$, where x describes the situation. In other words, $C_G\varphi$ can be seen as a fixed point of the operator $E_G(\varphi \wedge x)$. We now make this precise.

We fix an interpreted system $\mathcal{I}$ and consider functions that take as arguments sets of points in $\mathcal{I}$.

Definition 10.6. Let A and B be sets of points. For a function f and a set A, let $f^0(A) = A$ and $f^{k+1}(A) = f\big(f^k(A)\big)$. Define

$$\begin{aligned}
f_\neg(A) &= \overline{A} \qquad \text{(the complement of } A\text{)}\\
f_\wedge(A,B) &= A \cap B\\
f_{K_i}(A) &= \{(r,m) \mid r_i(m) = r_i'(m') \text{ implies } (r',m') \in A\}\\
f_{E_G}(A) &= \bigcap_{i \in G} f_{K_i}(A)\\
f_{C_G}(A) &= \bigcap_{k>0} f_{E_G}^k(A).
\end{aligned}$$

Recall that the intension of a formula φ is

$$\varphi^{\mathcal{I}} = \{(r,m) \mid (\mathcal{I},r,m) \models \varphi\}.$$

Notice that

$$\begin{aligned}
(\neg\varphi)^{\mathcal{I}} &= f_\neg(\varphi^{\mathcal{I}})\\
(\varphi \wedge \psi)^{\mathcal{I}} &= f_\wedge(\varphi^{\mathcal{I}}, \psi^{\mathcal{I}})\\
(K_i\varphi)^{\mathcal{I}} &= f_{K_i}(\varphi^{\mathcal{I}})\\
(E_G\varphi)^{\mathcal{I}} &= f_{E_G}(\varphi^{\mathcal{I}})
\end{aligned}$$

and

$$(C_G\varphi)^{\mathcal{I}} = f_{C_G}(\varphi^{\mathcal{I}}).$$

Extend the language $\mathcal{L}_n^C$ with the variable x and define

$$f_x(A) = A.$$

Furthermore for any formula φ, define $f_\varphi(A)$ by induction on the structure of φ as follows:

$$\begin{aligned}
f_p(A) &= p^{\mathcal{I}} \qquad \text{for a primitive proposition } p\\
f_{\neg\varphi} &= f_{\neg}\big(f_\varphi(A)\big)\\
f_{\varphi\wedge\psi} &= f_{\wedge}\big(f_\varphi(A), f_\psi(A)\big)\\
f_{K_i\varphi} &= f_{K_i}\big(f_\varphi(A)\big)\\
f_{E_G\varphi} &= f_{E_G}\big(f_\varphi(A)\big)\\
f_{C_G\varphi} &= f_{C_G}\big(f_\varphi(A)\big).
\end{aligned}$$

Then,

$$f_\varphi(A) = \varphi^{\mathcal{I}}$$

for every formula $\varphi \in \mathcal{L}_n^C$ and any set A. Moreover, let $\varphi[x/\psi]$ be the result of replacing all occurrences of x in φ by ψ. Then,

$$f_\varphi(\psi^{\mathcal{I}}) = (\varphi[x/\psi])^{\mathcal{I}}.$$

By the Fixed Point Axiom,

$$\mathcal{I} \models C_G\varphi \Leftrightarrow E_G(\varphi \wedge C_G\varphi).$$

Thus,

$$(C_G\varphi)^{\mathcal{I}} = \big(E_G(\varphi \wedge C_G\varphi)\big)^{\mathcal{I}}$$

and therefore

$$f_{E_G(\varphi\wedge x)}\big((C_G\varphi)^{\mathcal{I}}\big) = (C_G\varphi)^{\mathcal{I}}.$$

Hence, $(C_G\varphi)^{\mathcal{I}}$ is a fixed point of the function $f_{E_G(\varphi\wedge x)}$.

We want to prove that common knowledge is the greatest fixed point of $f_{E_G(\varphi\wedge x)}$.

A function f is monotone if $A \subseteq B$ implies $f(A) \subseteq f(B)$. We have the following lemma.

Lemma 10.7. *Let f be a monotone function. If $A \subseteq f(A)$, then $A \subseteq$ gfp (f), where gfp (f) denotes the greatest fixed point of f.*

Proof. Define

$$S = \bigcup_{A\subseteq f(A)} A \qquad\qquad \big(\emptyset \subseteq f(\emptyset)\big).$$

Then, $A \subseteq f(A)$ implies $A \subseteq S$, which in turn implies $f(A) \subseteq f(S)$, by the monotonicity of f. Hence, $A \subseteq f(S)$, and consequently

$$S = \bigcup_{A \subseteq f(A)} A \subseteq \bigcup_{A \subseteq f(A)} f(S) = f(S).$$

Now, $S \subseteq f(S)$ implies $f(S) \subseteq f\big(f(S)\big)$, which by the definition of S implies $f(S) \subseteq S$. Thus, $S = f(S)$. Moreover, for any fixed point A,

$$A \subseteq f(A) \quad \text{implies} \quad A \subseteq S,$$

and hence $\mathit{gfp}\,(f) = S$ by definition. $\square$

In particular, every monotone function has a greatest fixed point.

We extend the language with a greatest fixed point operator νx, where $\nu x[\varphi]$ is read "the greatest fixed point of φ". If f_φ is monotone, then we define

$$f_{\nu x[\varphi]} = \mathit{gfp}\,(f_\varphi)\,.$$

We denote the extended language by $\mathcal{L}_n^{C_\nu}$, and define

$$(\mathcal{I}, r, m) \models \varphi \qquad \text{iff} \qquad (r, m) \in f_\varphi(\emptyset).$$

Notice that this definition of $\models$ agrees with the ordinary one for formulas in $\mathcal{L}_n^C$. Call a formula closed if every occurrence of the variable x is in the scope of a νx operator.

Theorem 10.8. *For every closed formula $\varphi \in \mathcal{L}_n^{C_\nu}$ and nonempty set G,*

$$\models C_G\varphi \Leftrightarrow \nu x[E_G(\varphi \wedge x)].$$

Proof. Let A be an arbitrary fixed point of $f_{E_G(\varphi \wedge x)}$. Then,

$$\begin{aligned} A = f_{E_G(\varphi\wedge x)}(A) = f_{E_G}\big(f_{\varphi\wedge x}(A)\big) = \\ f_{E_G}(\varphi^{\mathcal{I}} \cap A) = f_{E_G}(\varphi^{\mathcal{I}}) \cap f_{E_G}(A) \subseteq f_{E_G}(\varphi^{\mathcal{I}}) = f_{E_G}^1(\varphi^{\mathcal{I}}) \end{aligned}$$

Assume $A \subseteq f_{E_G}^k(\varphi^{\mathcal{I}})$. By the monotonicity of f_{E_G} and since $A \subseteq f_{E_G}(A)$,

$$A \subseteq f_{E_G}(A) \subseteq f_{E_G}\big(f_{E_G}^k(\varphi^{\mathcal{I}})\big) = f_{E_G}^{k+1}(\varphi^{\mathcal{I}})$$

Induction now gives,

$$A \subseteq \bigcap_{k>0} f_{E_G}^k(\varphi^{\mathcal{I}}) = f_{C_G}(\varphi^{\mathcal{I}}) = (C_G\varphi)^{\mathcal{I}}$$

and we are done. $\square$

10.3.1 Downward continuity and infinite conjunctions

A function f is downward continuous if

$$f\left(\bigcap_{i=0}^{\infty} A_i\right) = \bigcap_{i=0}^{\infty} f(A_i)$$

whenever $A_0 \supseteq A_1 \supseteq \cdots$. All downward continuous functions are monotone. Namely, if $A \subseteq B$, then $f(A) \cap f(B) = f(A \cap B) = f(A)$ from which follows $f(A) \subseteq f(B)$.

Lemma 10.9. *If $f : 2^S \to 2^S$ is downward continuous, then*

$$\mathit{gfp}\,(f) = \bigcap_{i=0}^{\infty} f^i(S).$$

Proof. Note that $S = f^0(S) \supseteq f^1(S)$. Assume $f^i(S) \supseteq f^{i+1}(S)$. Since f is monotone,

$$f^{i+1}(S) = f\big(f^i(S)\big) \supseteq f\big(f^{i+1}(S)\big) = f^{i+2}(S),$$

and induction shows that $f^i(S) \supseteq f^{i+1}(S)$ for all $i \geq 0$. Hence, as

$$f\left(\bigcap_{i=0}^{\infty} f^i(S)\right) = \bigcap_{i=0}^{\infty} f\big(f^i(S)\big) = \bigcap_{i=1}^{\infty} f^i(S) = \bigcap_{i=0}^{\infty} f^i(S),$$

$\bigcap_{i=0}^{\infty} f^i(S)$ is a fixed point of f.

Let $A = f(A)$. Then, $A \subseteq S = f^0(S)$. Suppose $A \subseteq f^i(S)$. By the monotonicity of f, $A = f(A) \subseteq f\big(f^i(S)\big) = f^{i+1}(S)$. Consequently,

$$A \subseteq \bigcap_{i=0}^{\infty} f^i(S),$$

and $\bigcap_{i=0}^{\infty} f^i(S)$ is the greatest fixed point of f. □

Since $f_{E_G(\varphi \wedge x)}$ is downward continuous and since $\top^{\mathcal{I}}$ consists of all points in the system $\mathcal{I}$,

$$\mathit{gfp}\,\big(f_{E_G(\varphi \wedge x)}\big) = \bigcap_{i=0}^{\infty} f^i_{E_G(\varphi \wedge x)}(\top^{\mathcal{I}}).$$

Define the sequence $\varphi_0, \varphi_1, \ldots$ inductively by taking $\varphi_0 = \top$ and $\varphi_{i+1} = E_G(\varphi \wedge x)[x/\varphi_i] = E_G(\varphi \wedge \varphi_i)$. Then,

$$f^i_{E_G(\varphi \wedge x)}(\top^{\mathcal{I}}) = \varphi_i^{\mathcal{I}}$$

and therefore

$$\big(\nu x[E_G(\varphi \wedge \varphi_i)]\big)^{\mathcal{I}} = gfp\ \big(f_{E_G(\varphi\wedge x)}\big) = \bigcap_{i=0}^{\infty} \varphi_i^{\mathcal{I}}.$$

This can also be written as

$$\mathcal{I} \models \nu x[E_G(\varphi \wedge \varphi_i) \Leftrightarrow \bigwedge_{i=0}^{\infty} \varphi_i.$$

Since $\varphi_i \Leftrightarrow E_G\varphi \wedge \cdots \wedge E_G^i\varphi$, we have again arrived at

$$\mathcal{I} \models \nu x[E_G(\varphi \wedge \varphi_i) \Leftrightarrow C_G\varphi.$$

10.4 Approximations of common knowledge

Let $\epsilon \geq 0$. We say that within an ϵ interval everyone in G knows φ, written $E_G^\epsilon\varphi$, if there is an interval of ϵ time units in which every process at some time knows φ. Formally, $(\mathcal{I}, r, m) \models E_G^\epsilon\varphi$ if there exists an interval $[m', m' + \epsilon]$ such that $m \in [m', m' + \epsilon]$ and for every $i \in G$ there is an $m_i \in [m', m' + \epsilon]$ for which $(\mathcal{I}, r, m_i) \models K_i\varphi$ holds. Define ϵ-common knowledge $C_G^\epsilon\varphi$ as

$$C_G^\epsilon\varphi = \nu x[E_G^\epsilon(\varphi \wedge x)].$$

An ensemble e for G is ϵ-coordinated if the local events in e never hold more than ϵ time units apart. That is, if $(r, m) \in e(i)$ for some $i \in G$, then there is an interval $[m', m' + \epsilon] \ni m$ such that for all $j \in G$ there is a time $m_j \in [m', m' + \epsilon]$ for which $(r, m_j) \in e(j)$.

Proposition 10.10. *Let $\mathcal{I}$ be an interpreted system and let G be a set of agents. Then,*

1. *for every formula φ, the ensemble e for G defined by $e(i) = ev_{\mathcal{I}}(K_iC_G^\epsilon\varphi)$ is ϵ-coordinated,*
2. *if e is an ϵ-coordinated ensemble for G, then $\mathcal{I} \models \psi_{e(i)} \Rightarrow K_iC_G^\epsilon\psi_{e(i)}$ for all $i \in G$, and*
3. *if e is an ϵ-coordinated ensemble for G, then $\mathcal{I} \models \psi_e \Rightarrow C_G^\epsilon\psi_e$.*

An even weaker form of common knowledge is eventual common knowledge. Define $E_G^\Diamond\varphi$ to hold at $(\mathcal{I}, r, m)$ if for each $i \in G$ there is some time m_i such that $(\mathcal{I}, r, m_i) \models K_i\varphi$. Thus, $E_G^\Diamond\varphi$ holds if every agent eventually knows φ. Eventual common knowledge is then defined as

$$C_G^\Diamond\varphi = \nu x[E_G^\Diamond(\varphi \wedge x)].$$