Preprint Cambridge Centre for Computational Chemical Engineering ISSN 1473 – 4273

Reduced kinetic mechanism for turbulent combustion of ammonia/air mixtures

Jethro Akroyd[1,2,3], Yu-Chi Kao[4], Sunkyu Shin[4], Timo T. Pekkanen[4], Sebastian Mosbach[1,2,3], William H. Green[4], Markus Kraft[1,2,3,4]

released: July 3, 2026

[1] Department of Chemical Engineering and Biotechnology
University of Cambridge
Philippa Fawcett Drive
Cambridge, CB3 0AS
United Kingdom

[2] CARES
Cambridge Centre for Advanced Research and Education in Singapore
1 Create Way
CREATE Tower, #05-05
Singapore, 138602

[3] CMCL
No. 9, Journey Campus
Castle Park
Cambridge
CB3 0AX
United Kingdom

[4] MIT, Chemical Engineering
77 Massachusetts Avenue, Room E17-504
Cambridge, MA 02139 USA

Preprint No. 346

UNIVERSITY OF CAMBRIDGE

Keywords: Ammonia combustion, Reduced kinetic mechanism, Chemical mechanism reduction, Turbulence-chemistry interaction, Partially-stirred reactor, NO_x emissions

Edited by
Computational Modelling Group
Department of Chemical Engineering and Biotechnology
University of Cambridge
Philippa Fawcett Drive
Cambridge, CB3 0AS
United Kingdom

E-Mail: mk306@cam.ac.uk
World Wide Web: https://como.ceb.cam.ac.uk/

Abstract

This paper presents compact reduced kinetic mechanisms for ammonia/air combustion at atmospheric pressure for use in computational fluid dynamics (CFD) simulations. Perfectly-stirred reactor (PSR) and partially-stirred reactor (PaSR) simulations were used to sample hot-lean and hot-rich initial conditions, residence time, and turbulent-mixing time. At short residence times, combustion stability depended strongly on turbulent mixing and initial conditions, with faster mixing shifting selected cases from quenching to sustained combustion. Flux analysis showed that the dominant nitrogen-species-pair network was broadly conserved between fast- and slow-mixing cases, while the different outcomes were explained by different equivalence-ratio–temperature histories. Reduced mechanisms were generated using the directed relation graph with error propagation (DRGEP) method from both PaSR-derived and PSR-derived datasets. Both reduction routes gave similar performance for the cases studied. The reduced mechanisms reproduced the detailed-model behaviour across the tested mixing histories while substantially reducing CPU time, yielding CFD-ready mechanisms for atmospheric-pressure ammonia/air combustion. The mechanisms are available for download.

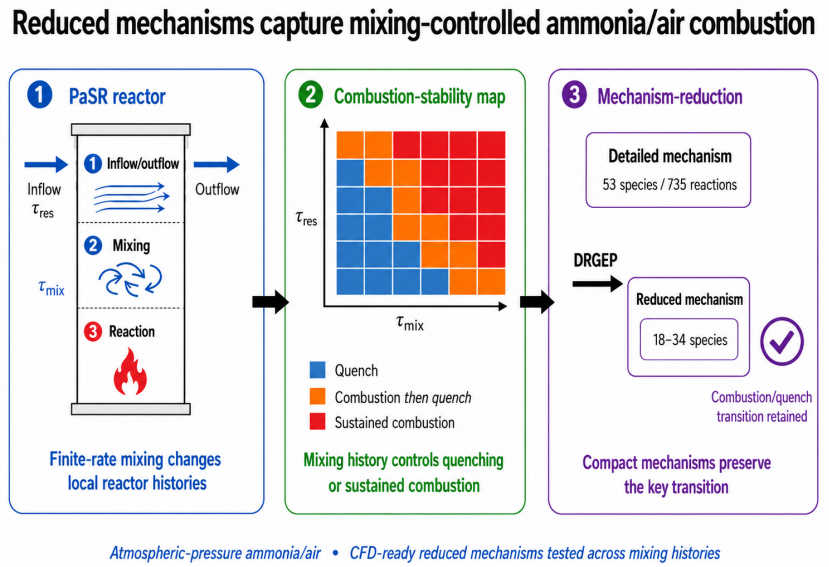

Highlights

- Compact reduced mechanisms developed for atmospheric-pressure ammonia/air CFD
- PaSR tests captured finite-mixing quenching and sustained combustion
- Hot-rich mixtures remained reactive over wider mixing-time ranges
- Nitrogen fluxes showed similar dominant pathways across mixing rates
- PSR-based reduction matched PaSR-derived performance for tested cases

Contents

1 Introduction

Ammonia has emerged as a promising carbon-free fuel for power generation and propulsion owing to its high hydrogen content and potential for carbon-neutral energy conversion [8, 21]. However, its use in practical combustion systems is challenging. Ignition and stable combustion are hindered by low laminar flame speed and high ignition temperature, resulting in incomplete oxidation and unstable flames. These effects often lead to excessive unburned ammonia (slip), while high temperatures promote nitrogen oxides (NO_x) formation [9, 21]. System-level analyses show that ammonia also faces challenges in energy density and cost [13], motivating efforts to develop combustion systems capable of stable, efficient, low-emission operation.

Fig. 1 illustrates a representative ammonia combustion device. Preheated ammonia and air enter primary and secondary zones, producing locally fuel-rich and -lean regions whose behaviour depends on local residence and turbulent-mixing times. Typical flame-zone residence times are of the order 10–100 ms and turbulent-mixing times are inferred to range from 0.1–1 ms, with feeds preheated to 400–800 K, equivalence ratios of 0.8–1.2, and 10–50 % v/v hydrogen enrichment in devices operating at pressures of 1–20 bar [1, 5, 14, 21]. Understanding and predicting ammonia oxidation under such conditions is essential for achieving stable operation while limiting emissions. Recent reviews indicate that enhanced mixing, high preheat, or hydrogen enrichment can shorten characteristic time scales to those of light hydrocarbon flames, permitting residence times in the single-millisecond range [5], whereas low-turbulence or distributed-reaction conditions can extend effective mixing times to several milliseconds [5, 21].

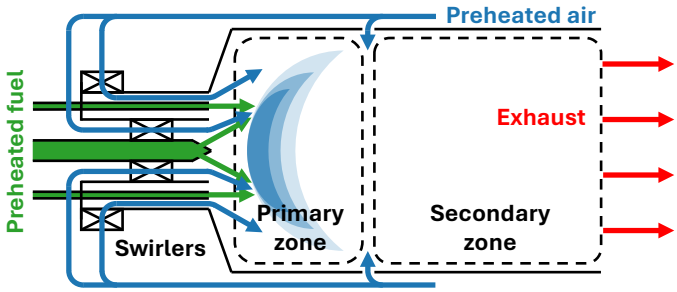

Figure 1: *Representative ammonia combustion device.*

Accurate computational fluid dynamics (CFD) simulations are essential for the design and optimisation of ammonia-fuelled combustion systems and for predicting ignition, flame stability, and emissions. Their accuracy depends critically on the kinetic mechanisms used to represent the underlying chemistry. Numerous detailed mechanisms for ammonia and ammonia–hydrogen combustion have been proposed in recent years (*e.g.*, [19, 24–26]), typically by extending hydrocarbon oxidation frameworks with nitrogen sub-mechanisms. Such composite mechanisms often suffer from inconsistencies in thermochemical data and empirically tuned rate parameters, which limit their predictive accuracy beyond the conditions for which they were optimised [10].

A recent mechanism for NH_3/H_2 combustion [10], generated with the Reaction Mechanism Generator (RMG) [4, 7, 15], addresses these issues by deriving all reactions and parameters from first-principles thermochemical and kinetic data, avoiding empirical fitting and ensuring internal consistency. The resulting model is expected to provide reliable predictions under conditions beyond those covered by existing experimental validations. However, as with other detailed mechanisms, the number of species and reactions is prohibitive for direct use in CFD. Mechanism reduction is therefore essential if detailed ammonia chemistry is to be used in practical simulations. Recent work has produced CFD-oriented reduced mechanisms for ammonia and ammonia blends, typically validated against laminar flame speeds, ignition delays and species measurements [17]. The same first-principles mechanism has also recently been reduced and tested against canonical kinetic, flame, reactor, extinction and CFD-fidelity cases for NH_3/H_2/air combustion [22]. The present work addresses a complementary question by assessing whether reduced mechanisms for atmospheric-pressure ammonia/air combustion reproduce detailed-model behaviour across finite-mixing partially stirred reactor (PaSR) histories as well as perfectly stirred reactor (PSR) histories.

The reduced mechanism must be compact, but compactness alone is not sufficient. In a turbulent combustor, material can experience different residence times, mixing times, and local mixture histories before it ignites, quenches, or forms emissions. A useful reduced mechanism must therefore retain the dominant ammonia-combustion chemistry and reproduce the chemical behaviour across these different thermochemical histories. This is the central issue addressed here: not only whether turbulent mixing changes combustion behaviour, but whether a compact mechanism can remain reliable when the mixing history changes.

The **purpose of this paper** is to develop compact, CFD-ready reduced kinetic mechanisms for ammonia/air combustion at atmospheric pressure and validate them across a range of residence and turbulent-mixing times. Flux analysis is used to interpret the effect of turbulent mixing. The reduction is based on the detailed RMG-derived mechanism of Kao *et al.* [10] and uses both perfectly-stirred and partially-stirred reactor models to sample the relevant states. The resulting reduced mechanisms are then tested against the detailed model to confirm that they preserve the shared chemistry needed for CFD-relevant ammonia-combustion simulations.

2 Method

A stochastic reactor model (SRM) was used to study turbulence–chemistry interaction by coupling detailed kinetics with a statistical description of subgrid inhomogeneity. Derived from the one-point joint composition probability density function (PDF) transport equation for homogeneous turbulence, the SRM represents inhomogeneous reactor contents as an ensemble of stochastic particles [11]. It unifies reactor types such as perfectly- and partially-stirred reactors (PSR and PaSR) and has been used to study the coupled effects of chemistry, turbulent mixing and soot formation (*e.g.*, [2, 6, 12, 16]).

2.1 PaSR and PSR simulations

Time-resolved adiabatic perfectly- and partially-stirred reactor (PSR and PaSR) models were used to generate data for reduction of the detailed mechanism of Kao *et al.* [10]. The PaSR model provides numerically robust solutions that reproduce the statistical behaviour predicted by the underlying PDF formulation [3]. PaSR models have recently been used to study ammonia–air combustion, showing that finite-rate mixing can affect temperature, NO_x formation and the distribution of thermochemical states [23].

Fig. 2 provides a schematic illustration of the PaSR concept. At each time step, inflow/outflow, mixing, and reaction occur sequentially within the reactor. The PaSR mimics the subgrid behaviour of a representative CFD cell, where inflows and outflows introduce material with different temperatures and compositions, creating local inhomogeneities that evolve through the combined effects of turbulent mixing and chemical reactions. The PaSR formulation provides a flexible framework that allows independent variation of the initial condition, feed composition, residence time (τ_{res}), and turbulent-mixing time (τ_{mix}), enabling systematic exploration of turbulence–chemistry interactions.

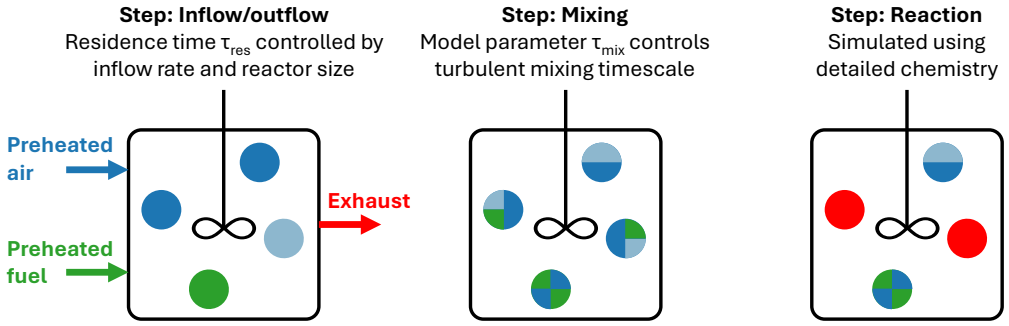

Figure 2: *Schematic illustration of PaSR concept.*

PSR simulations were performed to obtain steady-state compositions over a range of equivalence ratios for ammonia/air feeds at 1300 K. These provided hot initial conditions derived from lean ($\phi = 0.2$) and rich ($\phi = 4$) mixtures for PaSR simulations with air and ammonia preheated to 700 K and supplied in mass ratios giving an overall stoichiometric feed. The PaSR study covered residence times from 0.1 to 10 ms and turbulent-mixing times from 0.001 to 100 ms, with the shortest turbulent-mixing time approximating the perfectly-stirred limit. These values were chosen to sample short-residence and finite-rate mixing conditions motivated by practical ammonia-combustion devices.

Turbulent mixing in the PaSR was described using the Euclidean minimum spanning tree (EMST) model [20], in which stochastic particles mix locally in composition space. The stochastic inflow–outflow operator followed the PaSR formulation of Bhave and Kraft [3], with event-size parameter $\epsilon = 0.1$; further details are given in Appendix A.1. Separate PSR simulations were performed at the same residence times to provide a directly comparable perfectly-stirred reference dataset. All simulations were performed at 1 bar using ammonia fuel in ammonia/air mixtures.

2.2 Chemical mechanism reduction

The detailed mechanism of Kao et al. [10] was reduced using the directed relation graph with error propagation (DRGEP) method [18]. Species importance was quantified through overall interaction coefficients (OICs) that represent the normalised, time-integrated coupling of each species to the target species, NH_3, O_2, N_2, NO, NO_2, and N_2O. Species with OICs below a prescribed threshold ($\varepsilon_{\mathrm{OIC}}$) were removed, and the associated reactions pruned.

Two reduction datasets were considered. The first used the PaSR simulations described above, so that the OICs sampled residence time, turbulent-mixing time, and initial mixture history. The second used the corresponding perfectly-stirred data. The same target species, OIC definition, thresholds, and pruning procedure were applied to both datasets, giving paired PaSR- and PSR-reduced mechanisms at each threshold. The reduced mechanisms were then tested against the detailed mechanism across the validation cases.

3 Results

3.1 Influence of residence time and turbulent mixing

Fig. 3 summarises the PaSR residence and mixing times that yield combustion. Mixing strongly influences the system at residence times in the range $\tau_{\mathrm{res}} = 0.2$–1 ms: combustion quenches at long mixing times, with hot-rich mixtures remaining reactive at shorter residence times. At a residence time of 10 ms, stable combustion occurs for all conditions considered.

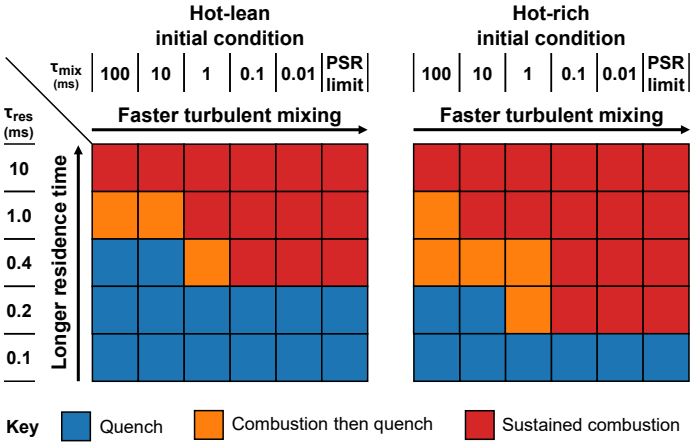

Figure 3: *Combustion stability as a function of residence time τ_{res} and turbulent-mixing time τ_{mix} at 1 bar for hot-lean and hot-rich initial conditions. Finite-rate mixing is decisive at short residence times, with hot-rich initial conditions sustaining combustion over a wider range than hot-lean initial conditions.*

Fig. 4 illustrates the critical role of mixing. It shows the temporal evolution of temperature, NH_3, OH, and NO in the hot-rich case across residence times and

mixing times. Ammonia depletion indicates conversion, OH marks ignition, and NO traces pollutant formation. As the turbulent-mixing time decreases, the cases with residence time $\tau_{\text{res}} \leq 1.0$ ms transition from quenching to sustained combustion, with faster mixing maintaining reactant supply and preventing extinction.

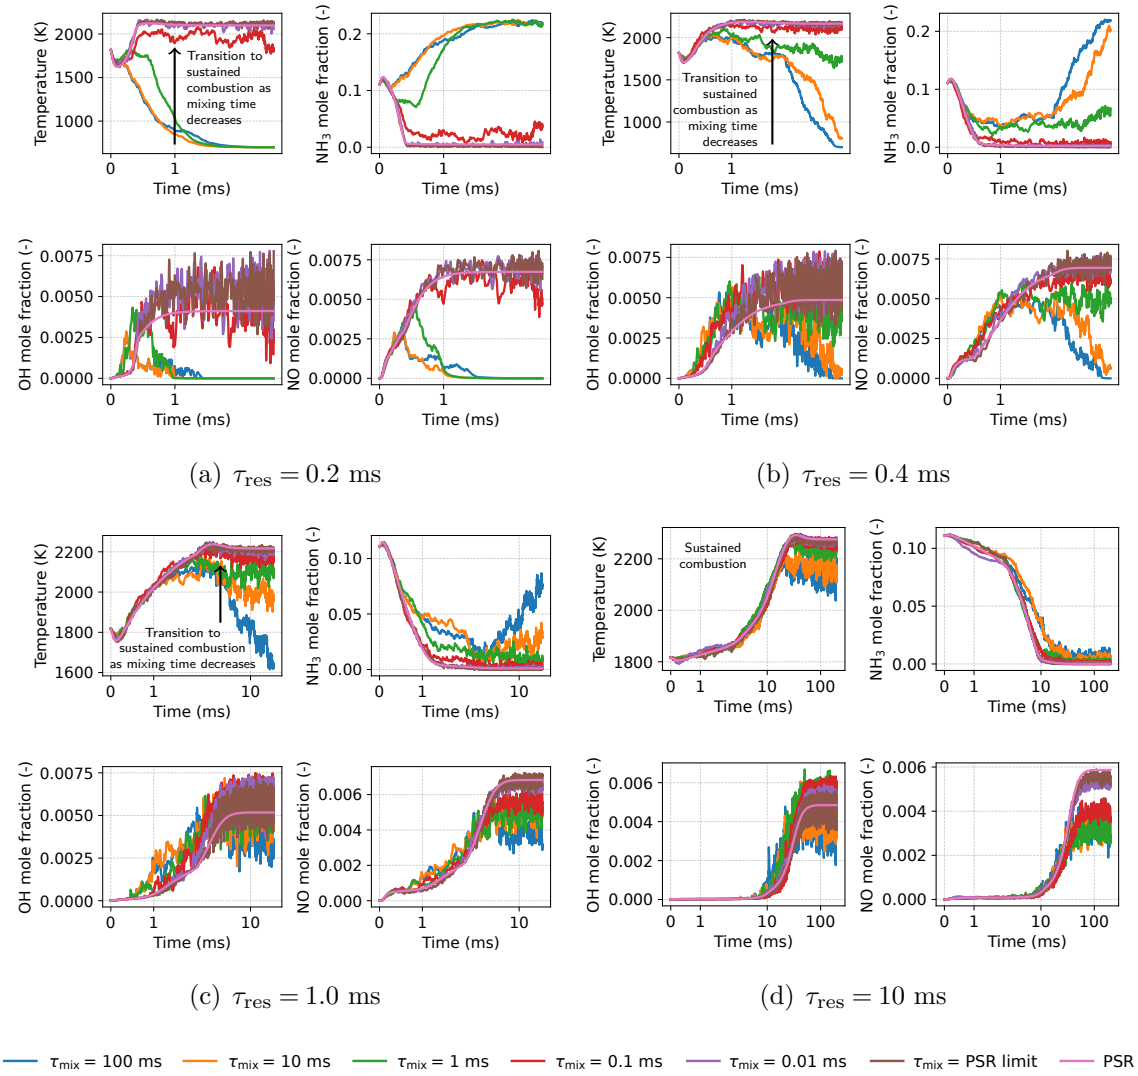

(a) $\tau_{\text{res}} = 0.2$ ms (b) $\tau_{\text{res}} = 0.4$ ms

(c) $\tau_{\text{res}} = 1.0$ ms (d) $\tau_{\text{res}} = 10$ ms

τ_{mix} = 100 ms τ_{mix} = 10 ms τ_{mix} = 1 ms τ_{mix} = 0.1 ms τ_{mix} = 0.01 ms τ_{mix} = PSR limit PSR

Figure 4: *Effect of turbulent mixing on transient evolution of temperature, NH_3, OH, and NO for hot-rich initial conditions. Decreasing the turbulent mixing time τ_{mix} moves the short-residence-time cases from extinction to sustained combustion.*

3.2 Mechanistic interpretation from flux analysis

The transient results show that turbulent mixing controls whether the hot-lean and hot-rich cases quench or achieve sustained combustion at residence times in the range $0.2 \leq \tau_{\text{res}} \leq 1.0$ ms. Elemental nitrogen flux analysis was used to examine whether this difference corresponds to a change in the underlying nitrogen chemistry.

Fig. 5 shows elemental-nitrogen species-pair fluxes for a fast-mixing case that achieves sustained combustion and a slow-mixing case that quenches after a period of initial combustion. The values of the fluxes for the leading species pairs and reactions are shown in Fig. A.1 in the Appendix. The fluxes show that the dominant species-pair network is broadly conserved across the fast- and slow-mixing cases. The largest fluxes involve the same core nitrogen pathways, indicating that fast and slow mixing do not activate fundamentally different nitrogen subnetworks.

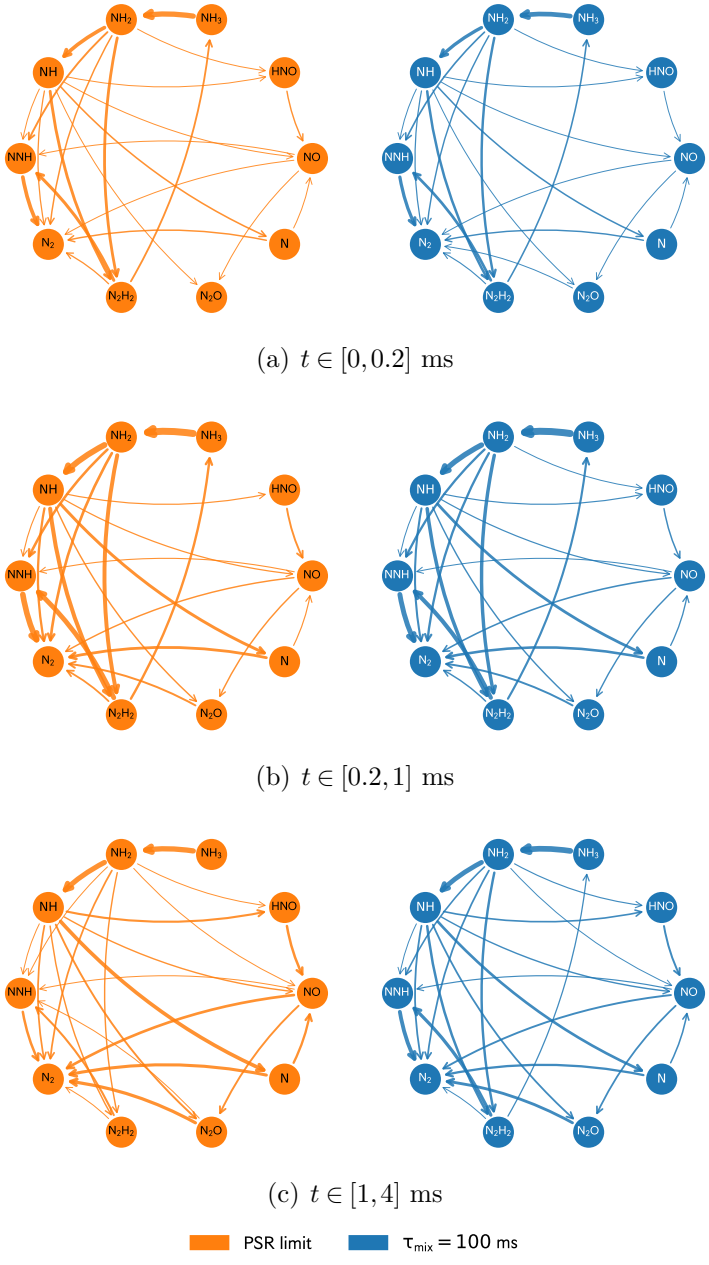

(a) $t \in [0, 0.2]$ ms

(b) $t \in [0.2, 1]$ ms

(c) $t \in [1, 4]$ ms

Figure 5: *Time-averaged elemental nitrogen species-pair fluxes for fast- and slow-mixing cases at $\tau_{res} = 1$ ms. The panels use the same species ordering and scale. The dominant nitrogen subnetwork is largely unchanged between the two cases.*

The difference between the fast- and slow-mixing cases is explained by their different equivalence-ratio–temperature histories. The histories for the fast- and slow-mixing cases at $\tau_{\mathrm{res}} = 1$ ms are shown in Figs. A.2 and A.3 in the Appendix. Fast mixing moves the reactor contents more rapidly through composition space and maintains a supply of reactants to the hot stochastic particles. Slow mixing leaves parts of the ensemble in regions where heat release and radical production are not sufficient to prevent extinction. The turbulent-mixing time changes the thermochemical trajectory and the timing and weighting of nitrogen-pathway activity, rather than activating a different nitrogen subnetwork.

These observations have a direct implication for the reduction step. A compact mechanism should retain the common nitrogen subnetwork, while remaining accurate for the different residence-time, mixing-time and initial-history combinations sampled by the PaSR calculations. The OIC analysis below provides this link.

3.3 Mechanism reduction

Fig. 6 shows the leading overall interaction coefficients (OICs) for species in the detailed mechanism, evaluated across all PaSR conditions. The leading species are consistent with the common nitrogen subnetwork identified by the flux analysis. However, the ranges show that species interaction coefficients change with residence time, turbulent-mixing time and initial mixture history. The reduction must therefore preserve the shared chemistry while remaining accurate across the different thermochemical histories sampled by the reactor calculations.

Reduced mechanisms were produced using the OICs evaluated across all PaSR conditions, as shown in Fig. 6, and using OICs evaluated across all PSR conditions (not shown). The values of the OICs and the species retained in the PaSR-derived and PSR-derived reduced mechanisms at different reduction thresholds are shown in Table A.1 in the Appendix.

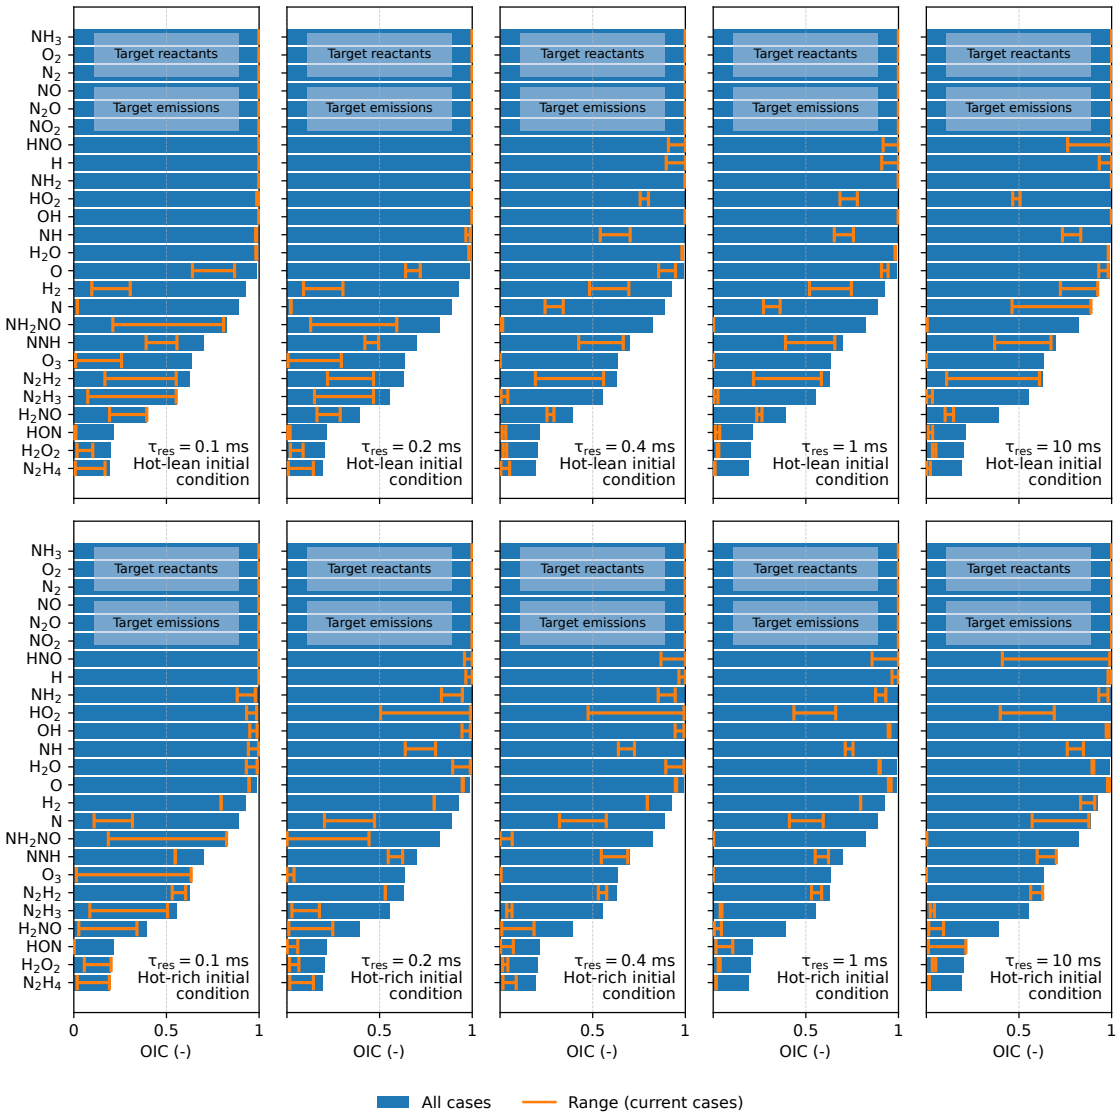

Figure 6: *Leading overall interaction coefficients (OICs). Bars: magnitude of the OIC across all PaSR cases. Lines: range of OICs across turbulent-mixing times τ_{mix} for different residence times τ_{res} (left-to-right) for hot-lean (top) and hot-rich (bottom) initial conditions. The spread in OICs shows that residence time, mixing time and initial conditions alter species importance; the mechanism reduction must preserve shared chemistry across these conditions.*

The reduced mechanisms were validated against the detailed mechanism for the full set of PaSR reactor calculations. Fig. 7 compares the temperature and key species profiles for the detailed and PaSR-derived reduced mechanisms at residence time $\tau_{\text{res}} = 0.2$ ms. This represents the most challenging case across all the conditions studied, because the correct outcome depends on the coupling between mixing, heat release and radical production. The reduced mechanisms with threshold $\varepsilon_{\text{OIC}} \leq 0.5$ accurately reproduce the mixing-controlled transition between quenching and sustained combustion (see Fig. 4), showing closer agreement with the detailed model as the reduction threshold ε_{OIC} decreases. The PSR-derived reduced mechanisms behaved similarly. This indicates that, for the present cases, the reduction route is less important than validation across the relevant mixing histories.

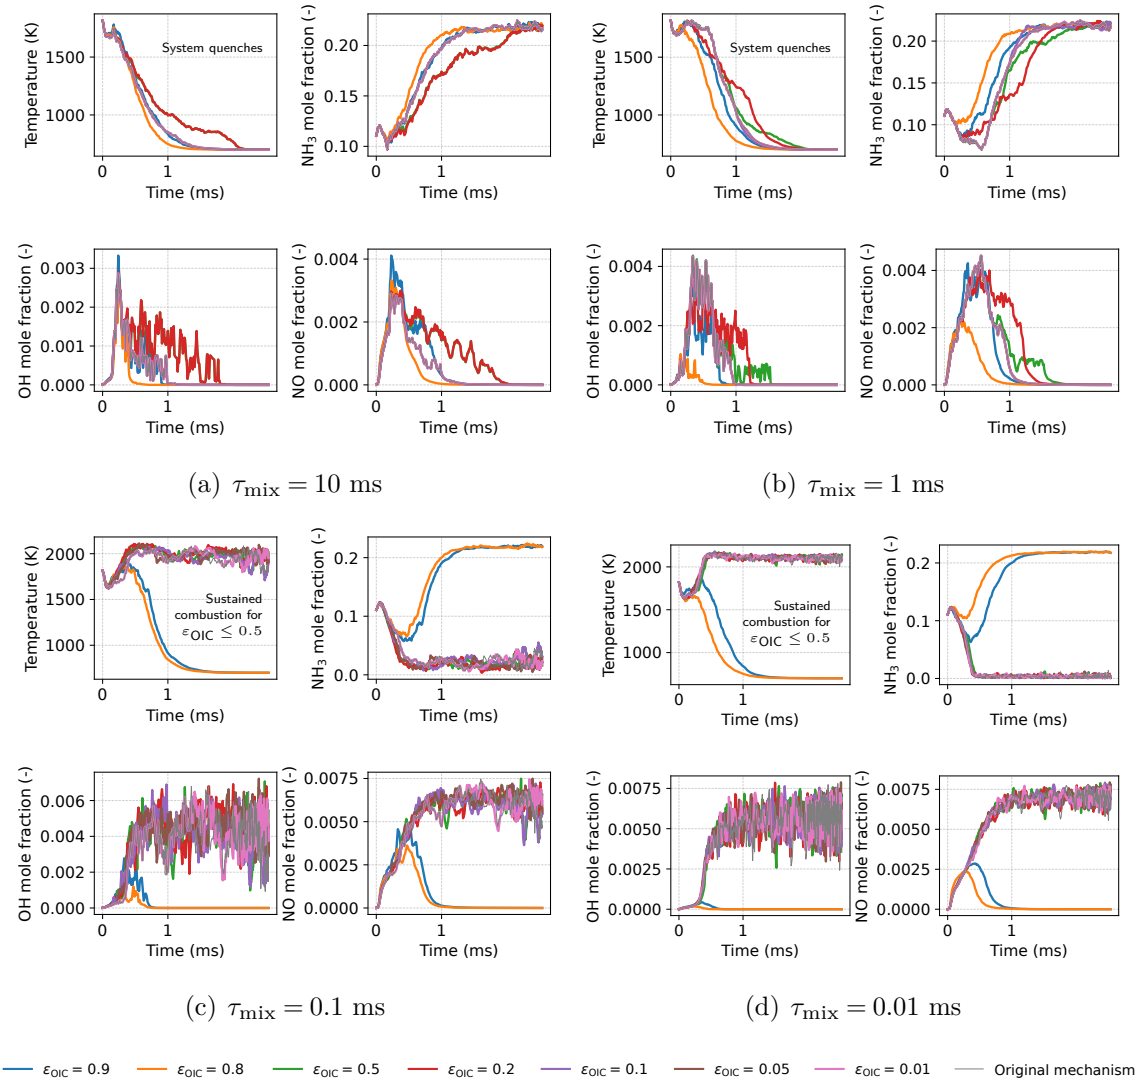

(a) $\tau_{\text{mix}} = 10$ ms (b) $\tau_{\text{mix}} = 1$ ms

(c) $\tau_{\text{mix}} = 0.1$ ms (d) $\tau_{\text{mix}} = 0.01$ ms

Figure 7: *Comparison of temperature, NH_3, OH and NO for the detailed and reduced mechanisms for hot-rich conditions at residence time $\tau_{res} = 0.2$ ms. The reduced mechanisms with thresholds $\varepsilon_{OIC} \leq 0.5$ retain the quench-to-combustion transition, while lower thresholds improve quantitative agreement with the detailed mechanism.*

Table 1 summarises the mechanism sizes, relative CPU times and mean absolute temperature errors for the PaSR-derived / PSR-derived reduced mechanisms. The temperature error is defined as the mean absolute difference between the detailed- and reduced-mechanism temperature profiles, averaged over all time points.

Table 1: *Mechanism reduction summary (PaSR-derived / PSR-derived). The PaSR- and PSR-derived mechanisms give substantial speed-up with modest temperature error; high thresholds over-reduce the mechanism.*

Threshold, $\varepsilon_{\mathrm{OIC}}$	# Species (-)	# Reactions (-)	Temperature error (K)	Relative CPU time (-)
Detailed mechanism	53	735	-	1
0.01	34 / 33	471 / 461	5 / 7	0.57 / 0.55
0.05	31 / 27	363 / 319	16 / 15	0.43 / 0.38
0.1	29 / 23	310 / 212	15 / 20	0.35 / 0.25
0.2	24 / 19	216 / 147	17 / 19	0.27 / 0.17
0.5	21 / 18	164 / 128	30 / 36	0.20 / 0.16
0.8	17 / 15	100 / 81	313 / 315	0.12 / 0.10
0.9	15 / 14	81 / 64	315 / 230	0.10 / 0.10

The reduced mechanisms show a clear trade-off between computational cost and accuracy. At $\varepsilon_{\mathrm{OIC}} = 0.01$, both the PaSR- and PSR-derived reduced mechanisms closely reproduce the behaviour of the detailed mechanism, with mean absolute temperature errors of 5–7 K, while reducing the CPU time to 55–57% of the original simulations. Increasing the OIC threshold gives larger speed-ups, with relative CPU times below 0.3 for $\varepsilon_{\mathrm{OIC}} \geq 0.2$, but at the cost of increased temperature error. This loss of accuracy is modest for intermediate thresholds, particularly between $\varepsilon_{\mathrm{OIC}} = 0.05$ and 0.2, but becomes substantial at very high thresholds. For $\varepsilon_{\mathrm{OIC}} = 0.8$ and 0.9, the mechanisms are too aggressively reduced to reproduce the original combustion behaviour.

4 Conclusions

This paper developed compact reduced mechanisms for ammonia/air combustion at atmospheric pressure for use in CFD simulations. The mechanisms were tested against partially-stirred reactor (PaSR) and perfectly-stirred reactor (PSR) calculations spanning residence time, turbulent-mixing time and initial mixture history, including cases close to the transition between quenching and sustained combustion. The results show that turbulent mixing has a strong effect on combustion stability at short residence times. Hot-rich mixtures remain reactive over a wider range of mixing times than hot-lean mixtures, and faster mixing can shift the system from quenching to sustained combustion.

Elemental nitrogen flux analysis showed that the dominant nitrogen-species-pair network is broadly conserved between the fast- and slow-mixing cases. The different combustion outcomes are explained mainly by the different equivalence-ratio–

temperature histories and by changes in the timing and weighting of nitrogen-pathway activity, rather than by the activation of fundamentally different nitrogen chemistry. This provides a useful interpretation of the reduction problem: the reduced mechanism must preserve the common controlling chemistry while remaining accurate across different thermochemical histories.

Reduced mechanisms were generated using the directed relation graph with error propagation (DRGEP) method using both PaSR-derived and PSR-derived simulations. Both approaches gave similar performance for the cases studied. For example, at an overall interaction coefficient threshold $\varepsilon_{\mathrm{OIC}} = 0.01$, the PaSR- and PSR-derived mechanisms reproduced the detailed-mechanism behaviour with mean absolute temperature errors of 5–7 K, while reducing the CPU time to 55–57% of the original simulations. More aggressive reductions continued to capture the mixing-controlled transition between quenching and sustained combustion up to $\varepsilon_{\mathrm{OIC}} = 0.5$, with a modest decrease in accuracy while reducing the CPU time to less than 30% of the original simulations. Reductions at higher thresholds lost this behaviour.

Overall, the results show that reduced mechanisms for ammonia/air combustion at atmospheric pressure should be tested across the relevant mixing histories. For the conditions considered here, PSR-based reduction was sufficient to reproduce the PaSR mixing-history effects, including the transition between quenching and sustained combustion. This supports the use of idealised reactor calculations for mechanism reduction, provided that the resulting mechanisms are validated against finite-rate mixing cases. The reduced mechanisms are available for download.

Future work will extend this approach to higher pressures and ammonia/hydrocarbon blends, and will test whether PSR-based reduction remains sufficient under those conditions.

Acknowledgements

This research was supported by the National Research Foundation, Prime Minister's Office, Singapore under its Campus for Research Excellence and Technological Enterprise (CREATE) programme. In particular, we benefited from HYCOMBS, which is part of the CREATE Thematic Programme in Decarbonisation. JA and MK gratefully acknowledge CMCL for software licences and technical support. YCK, SKS and WHG gratefully acknowledge support from ExxonMobil. YCK is supported by Taiwan NSTC (Graduate Students Study Abroad Program, 114-2917-I-002-025). TTP acknowledges the Finnish Foundation for Technology Promotion for fellowship funding and CSC IT Center for Science in Finland for computational resources.

Data and code availability

The PSR and PaSR simulations were performed using the Kinetics & SRM Engine Suite from CMCL (mailto:support@cmcl.io). The reduced mechanisms supporting the article are available via the University of Cambridge data repository (doi:10.17863/CAM.131370).

Declaration of Generative AI and AI-assisted technologies in the writing process

During the preparation of this work the authors used ChatGPT v5.2–v5.5 to enhance the readability and language of the manuscript. After using this tool, the authors reviewed and edited the content as needed and take full responsibility for the content of the publication.

A Appendix

A.1 Stochastic inflow–outflow event size

The stochastic inflow–outflow model follows the PaSR formulation of Bhave and Kraft [3]. Inflow–outflow is treated as a jump process in which each event exchanges a finite mass between the stochastic ensemble and the inlet streams. The total mass exchanged in one event is

$$W_{\mathrm{ex}} = \epsilon \frac{M}{N_{\mathrm{p}}}, \tag{A.1}$$

where M is the total mass of the stochastic ensemble, N_{p} is the number of stochastic particles, and ϵ is an event-size parameter. The waiting time between successive events is sampled from an exponential distribution with rate

$$\lambda = \frac{\sum_s \dot{m}_s}{W_{\mathrm{ex}}}, \tag{A.2}$$

where $\dot{m}_s$ is the mass-flow rate of inlet stream s. The expected exchanged mass per unit time is therefore

$$\lambda W_{\mathrm{ex}} = \sum_s \dot{m}_s, \tag{A.3}$$

so that the prescribed mean throughput is recovered independently of ϵ. For multiple inlet streams, W_{ex} is distributed among the streams in proportion to their mass-flow rates.

The parameter ϵ controls the granularity with which inflow–outflow introduces particle-to-particle composition differences into the stochastic ensemble. Smaller values split the same mean mass flux into smaller, more frequent exchange events, tending toward a more homogeneous particle ensemble. Larger values introduce inlet material more intermittently, producing stronger particle-to-particle composition differences that are subsequently reduced by turbulent mixing.

A.2 Flux analysis

Fig. A.1 shows time-resolved elemental nitrogen fluxes for the leading species pairs and reactions in the fast- and slow-mixing cases at $\tau_{\mathrm{res}} = 1$ ms. While the rates associated with the active pathways change with time, the dominant nitrogen-species-pair network remains similar.

The quantity $q_r(t)$ is the instantaneous rate of progress of reaction r. The contribution of reaction r to the directed species-pair flux $A \to B$ is

$$q_r^{A\to B}(t) = \nu_r^{A\to B} q_r(t), \tag{A.4}$$

where $\nu_r^{A\to B}$ is the stoichiometric coefficient associated with the directed transfer from species A to species B. If reaction r does not contribute to this directed pair, then $\nu_r^{A\to B} = 0$.

The total directed species-pair flux is the sum of the reaction-level contributions:

$$q_{A\to B}(t) = \sum_r q_r^{A\to B}(t). \tag{A.5}$$

Thus, $q_r(t)$ measures the instantaneous activity of one reaction, while $q_{A\to B}(t)$ measures the total instantaneous flux from species A to species B.

A.3 Equivalence-ratio–temperature histories

Figs. A.2 and A.3 compare the fast- and slow-mixing cases at $\tau_{\text{res}} = 1$ ms. Fig. A.2 shows the equivalence-ratio and temperature histories, while Fig. A.3 shows how the mean and standard deviation of the equivalence ratio change over time. The two cases follow different thermochemical trajectories.

A.4 Species retained in reduced mechanisms

Table A.1 shows the species retained in the PaSR-derived and PSR-derived reduced mechanisms at different reduction thresholds.

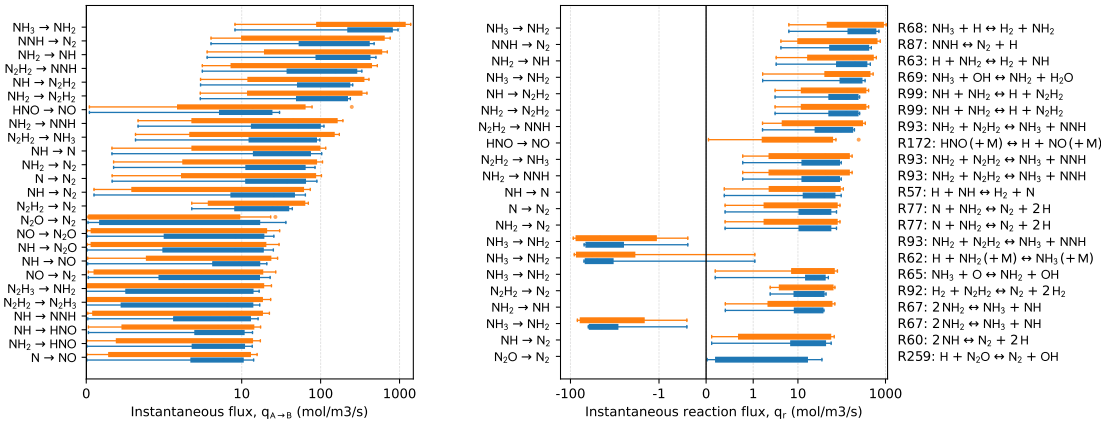

(a) $t \in [0, 0.2]$ ms

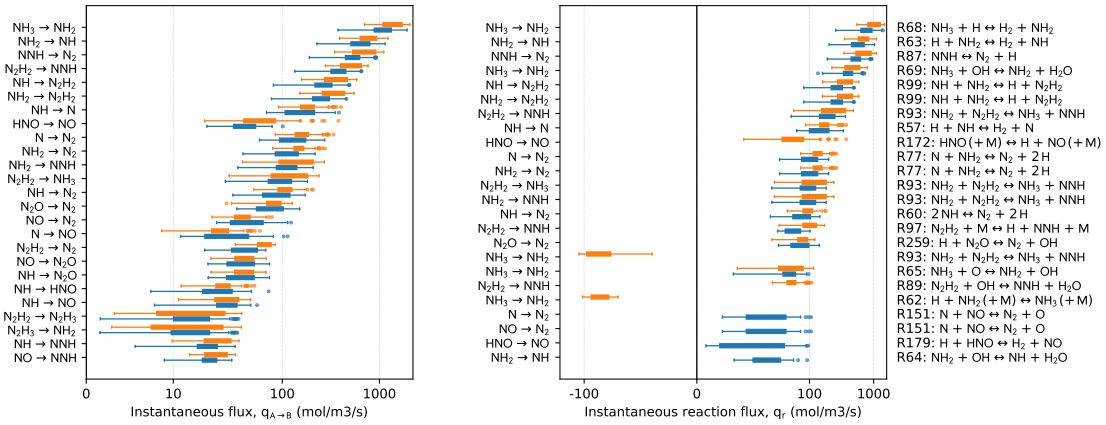

(b) $t \in [0.2, 1]$ ms

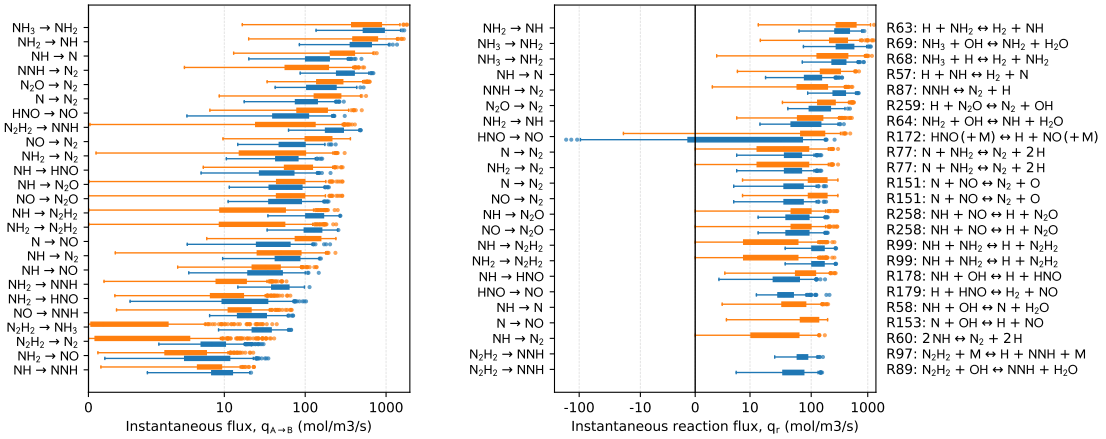

(c) $t \in [1, 4]$ ms

PSR limit $\tau_{mix} = 100$ ms

Figure A.1: *Time-averaged elemental nitrogen fluxes for the dominant species pairs and reaction-level contributions in the fast- and slow-mixing cases at $\tau_{res} = 1$ ms. The fast and slow mixing cases proceed via the same principal nitrogen-reaction network, while changing the timing and magnitude of the fluxes.*

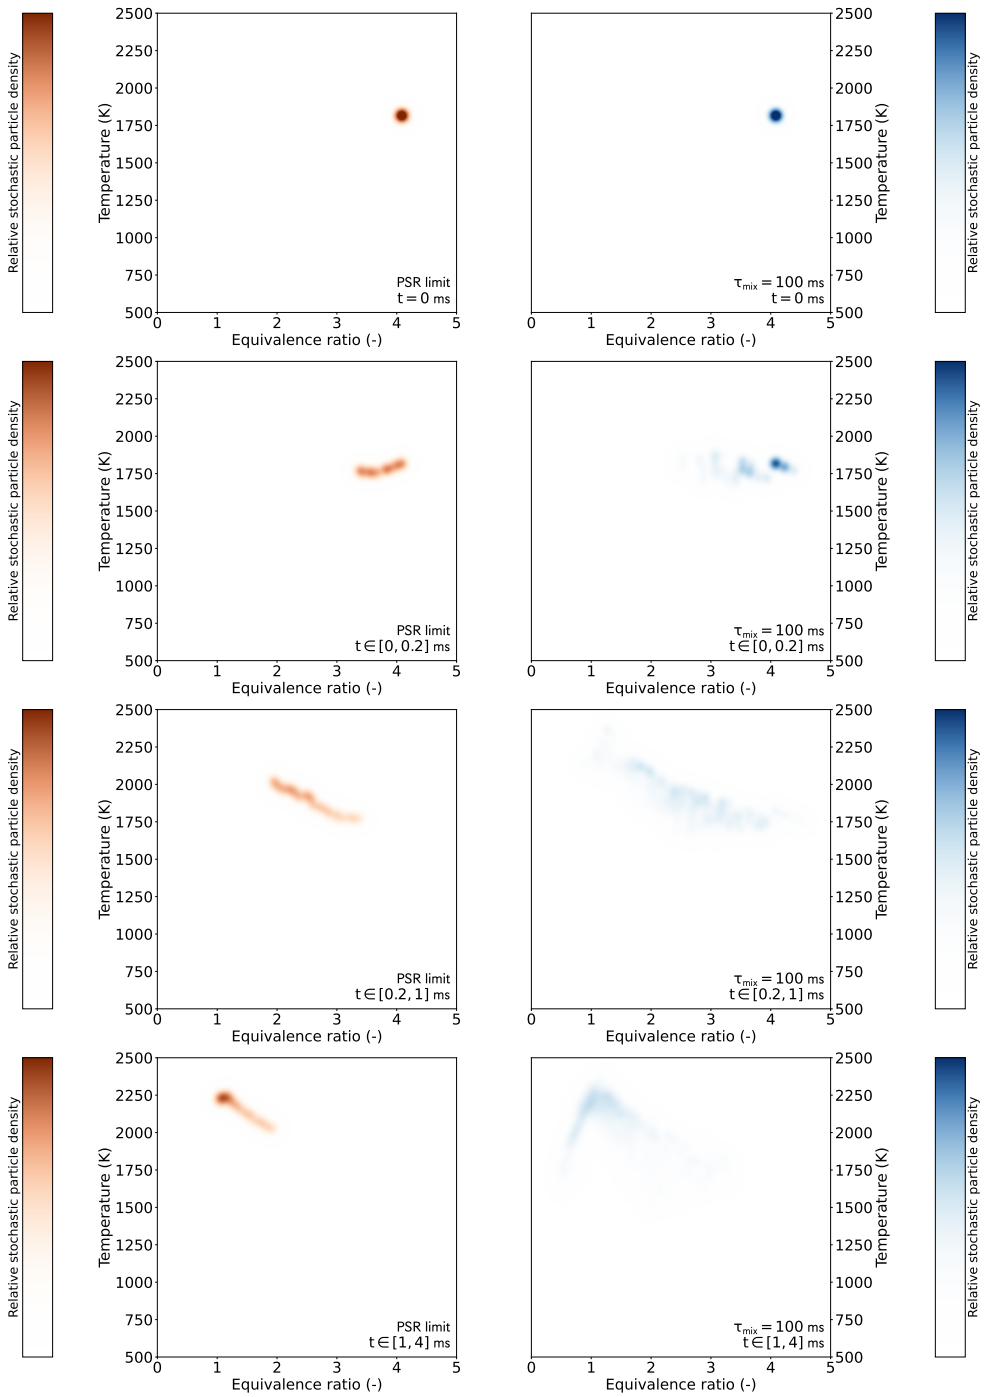

Figure A.2: *Equivalence-ratio–temperature histories for the fast- and slow-mixing cases at $\tau_{res} = 1$ ms. Markers indicate the relative stochastic particle density in each time window. The turbulent mixing changes the thermochemical states sampled by each case.*

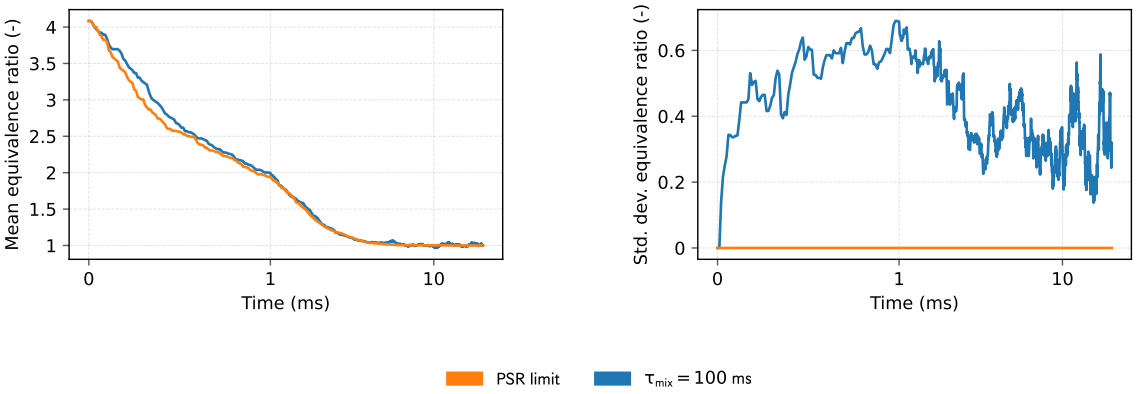

Figure A.3: *Time history of the mean and standard deviation of the equivalence ratio for the fast- and slow-mixing cases at $\tau_{res} = 1$ ms. Slow mixing results in significant inhomogeneities, whereas fast mixing drives the system toward a homogeneous state.*

Table A.1: *Species retained in PaSR-derived and PSR-derived reduced mechanisms. Both reductions retain the same core chemistry, explaining why the PaSR- and PSR-derived mechanisms give similar overall performance.*

Species	Overall interaction coefficient, OIC (-)		Species retained at $\varepsilon_{\mathrm{OIC}}$						
	PaSR	PSR	0.01	0.05	0.1	0.2	0.5	0.8	0.9
N_2	1.000	1.000	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
N_2O	1.000	1.000	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
NH_3	1.000	1.000	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
NO	1.000	1.000	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
NO_2	1.000	1.000	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
O_2	1.000	1.000	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
HNO	1.000	1.000	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
H	1.000	1.000	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
NH_2	1.000	1.000	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
HO_2	1.000	0.980	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
OH	0.999	0.999	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
NH	0.997	0.998	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
H_2O	0.993	0.996	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
O	0.988	0.986	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓
H_2	0.925	0.849	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓
N	0.888	0.589	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓	
NH_2NO	0.824	0.190	✓ ✓	✓ ✓	✓ ✓	✓	✓	✓	
NNH	0.701	0.646	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓		
O_3	0.633	1.211×10^{-2}	✓ ✓	✓	✓	✓	✓		
N_2H_2	0.627	0.693	✓ ✓	✓ ✓	✓ ✓	✓ ✓	✓ ✓		
N_2H_3	0.553	0.163	✓ ✓	✓ ✓	✓ ✓	✓	✓		
H_2NO	0.394	0.439	✓ ✓	✓ ✓	✓ ✓	✓ ✓			
HON	0.213	5.465×10^{-2}	✓ ✓	✓ ✓	✓	✓			
H_2O_2	0.201	0.153	✓ ✓	✓ ✓	✓ ✓	✓			
N_2H_4	0.192	0.180	✓ ✓	✓ ✓	✓ ✓				
HONO	0.148	4.197×10^{-2}	✓ ✓	✓	✓				
H_2NNO_2	0.145	4.846×10^{-2}	✓ ✓	✓	✓				
N_2H_3O	0.120	9.313×10^{-2}	✓ ✓	✓ ✓	✓				
N[N]O	0.118	9.315×10^{-2}	✓ ✓	✓ ✓	✓				
H_2NN	7.143×10^{-2}	8.241×10^{-2}	✓ ✓	✓ ✓					
NH_2OOH	5.654×10^{-2}	4.520×10^{-2}	✓ ✓	✓					
HNOH	4.753×10^{-2}	4.982×10^{-2}	✓ ✓						
$HONO_2$	3.743×10^{-2}	5.015×10^{-3}	✓						
NH_2OH	1.545×10^{-2}	1.583×10^{-2}	✓ ✓						
NH_2OO	4.196×10^{-3}	3.453×10^{-5}							
HNO_2	3.278×10^{-3}	3.223×10^{-3}							
$ONONO_2$	5.013×10^{-4}	2.089×10^{-5}							
O[N]O	4.253×10^{-4}	3.942×10^{-4}							
HONHO	4.247×10^{-4}	5.369×10^{-4}							
NO_3	3.943×10^{-4}	9.324×10^{-5}							
HONHOO	2.779×10^{-4}	3.375×10^{-5}							
H_2NONO	1.315×10^{-4}	1.438×10^{-4}							
N_2O_2	1.223×10^{-4}	1.586×10^{-7}							
NH_3O	9.980×10^{-5}	1.380×10^{-4}							
NNO[O]	7.021×10^{-5}	3.071×10^{-9}							
HN_2O	2.509×10^{-5}	1.004×10^{-6}							
$O_1O[N]_1$	1.554×10^{-6}	1.585×10^{-7}							
N_2O_3	6.344×10^{-7}	4.194×10^{-9}							
N_2O_4	1.455×10^{-8}	5.576×10^{-10}							
N_2O_5	1.990×10^{-12}	1.173×10^{-12}							
AR	0.000	0.000							
HE	0.000	0.000							
NE	0.000	0.000							

References

[1] M. Alnajideen, H. Shi, W. Northrop, D. Emberson, S. Kane, P. Czyzewski, M. Alnaeli, S. Mashruk, K. Rouwenhorst, C. Yu, S. Eckart, and A. Valera-Medina. Ammonia combustion and emissions in practical applications: a review. *Carbon Neutrality*, 3(1), 2024. doi:10.1007/s43979-024-00088-6.

[2] M. Balthasar, F. Mauss, A. Knobel, and M. Kraft. Detailed modeling of soot formation in a partially stirred plug flow reactor. *Combustion and Flame*, 128 (4):395–409, 2002. doi:10.1016/s0010-2180(01)00344-3.

[3] A. Bhave and M. Kraft. Partially stirred reactor model: Analytical solutions and numerical convergence study of a PDF/Monte Carlo method. *SIAM Journal on Scientific Computing*, 25(5):1798–1823, 2004. doi:10.1137/s1064827502411328.

[4] A. G. Dana, M. S. Johnson, J. W. Allen, S. Sharma, S. Raman, M. Liu, C. W. Gao, C. A. Grambow, M. J. Goldman, D. S. Ranasinghe, R. J. Gillis, A. M. Payne, Y. Li, X. Dong, K. A. Spiekermann, H. Wu, E. E. Dames, Z. J. Buras, N. M. Vandewiele, N. W. Yee, S. S. Merchant, B. Buesser, C. A. Class, F. Goldsmith, R. H. West, and W. H. Green. Automated reaction kinetics and network exploration (arkane): A statistical mechanics, thermodynamics, transition state theory, and master equation software. *International Journal of Chemical Kinetics*, 55(6):300–323, 2023. doi:10.1002/kin.21637.

[5] A. M. Elbaz, S. Wang, T. F. Guiberti, and W. L. Roberts. Review on the recent advances on ammonia combustion from the fundamentals to the applications. *Fuel Communications*, 10:100053, 2022. doi:10.1016/j.jfueco.2022.100053.

[6] J. Etheridge, S. Mosbach, M. Kraft, H. Wu, and N. Collings. Modelling soot formation in a DISI engine. *Proceedings of the Combustion Institute*, 33(2): 3159–3167, 2011. doi:10.1016/j.proci.2010.07.039.

[7] C. W. Gao, J. W. Allen, W. H. Green, and R. H. West. Reaction mechanism generator: Automatic construction of chemical kinetic mechanisms. *Computer Physics Communications*, 203:212–225, 2016. doi:10.1016/j.cpc.2016.02.013.

[8] A. Grinberg Dana, O. Elishav, A. Bardow, G. E. Shter, and G. S. Grader. Nitrogen-based fuels: A power-to-fuel-to-power analysis. *Angewandte Chemie International Edition*, 55(31):8798–8805, 2016. doi:10.1002/anie.201510618.

[9] O. Herbinet, P. Bartocci, and A. Grinberg Dana. On the use of ammonia as a fuel – a perspective. *Fuel Communications*, 11:100064, 2022. doi:10.1016/j.jfueco.2022.100064.

[10] Y.-C. Kao, A. C. Doner, T. T. Pekkanen, C. Cao, S. Shin, A. Grinberg Dana, Y.-P. Li, and W. H. Green. Detailed kinetic model for combustion of NH_3/H_2 blends. *Physical Chemistry Chemical Physics*, 28(10):6411–6424, 2026. doi:10.1039/d5cp04149j.

[11] M. Kraft. *Stochastic Modeling of Turbulent Reacting Flow in Chemical Engineering Applications*, volume 391 of *Fortschritt-Berichte VDI, Reihe 6: Energietechnik*. VDI Verlag, Düsseldorf, 1998. ISBN 3-18-339106-0.

[12] M. Kraft, P. Maigaard, F. Mauss, M. Christensen, and B. Johansson. Investigation of combustion emissions in a homogeneous charge compression injection engine: Measurements and a new computational model. *Proceedings of the Combustion Institute*, 28(1):1195–1201, 2000. doi:10.1016/s0082-0784(00)80330-6.

[13] L. C. Law, E. Mastorakos, and S. Evans. Estimates of the decarbonization potential of alternative fuels for shipping as a function of vessel type, cargo, and voyage. *Energies*, 15(20):7468, 2022. doi:10.3390/en15207468.

[14] J. Li, S. Lai, D. Chen, R. Wu, N. Kobayashi, L. Deng, and H. Huang. A review on combustion characteristics of ammonia as a carbon-free fuel. *Frontiers in Energy Research*, 9, 2021. doi:10.3389/fenrg.2021.760356.

[15] M. Liu, A. Grinberg Dana, M. S. Johnson, M. J. Goldman, A. Jocher, A. M. Payne, C. A. Grambow, K. Han, N. W. Yee, E. J. Mazeau, K. Blondal, R. H. West, C. F. Goldsmith, and W. H. Green. Reaction mechanism generator v3.0: Advances in automatic mechanism generation. *Journal of Chemical Information and Modeling*, 61(6):2686–2696, 2021. doi:10.1021/acs.jcim.0c01480.

[16] S. Mosbach, H. Su, M. Kraft, A. Bhave, F. Mauss, Z. Wang, and J.-X. Wang. Dual injection homogeneous charge compression ignition engine simulation using a stochastic reactor model. *International Journal of Engine Research*, 8(1): 41–50, 2007. doi:10.1243/14680874jer01806.

[17] K. Pan, S. Zhong, X. Huang, Z. Xu, X. Han, and L. Chen. A reduced chemical kinetic mechanism for cfd modeling of combustion fueled by ammonia and its blends with methane/hydrogen. *Chemical Engineering Journal*, 505:159253, 2025. doi:10.1016/j.cej.2025.159253.

[18] P. Pepiot-Desjardins and H. Pitsch. An efficient error-propagation-based reduction method for large chemical kinetic mechanisms. *Combustion and Flame*, 154 (1–2):67–81, 2008. doi:10.1016/j.combustflame.2007.10.020.

[19] A. Stagni, S. Arunthanayothin, M. Dehue, O. Herbinet, F. Battin-Leclerc, P. Bréquigny, C. Mounaïm-Rousselle, and T. Faravelli. Low- and intermediate-temperature ammonia/hydrogen oxidation in a flow reactor: Experiments and a wide-range kinetic modeling. *Chemical Engineering Journal*, 471:144577, 2023. doi:10.1016/j.cej.2023.144577.

[20] S. Subramaniam and S. Pope. A mixing model for turbulent reactive flows based on euclidean minimum spanning trees. *Combustion and Flame*, 115(4): 487–514, 1998. doi:10.1016/s0010-2180(98)00023-6.

[21] A. Valera-Medina, H. Xiao, M. Owen-Jones, W. David, and P. Bowen. Ammonia for power. *Progress in Energy and Combustion Science*, 69:63–102, 2018. doi:10.1016/j.pecs.2018.07.001.

[22] L. Wang, Y.-C. Kao, G. Krishnan, A. F. Ghoniem, W. H. Green, and S. Deng. Mechanism reduction and validation of ammonia/hydrogen/air combustion leveraging first-principles-derived kinetics. *ChemRxiv*, 2026. doi:10.26434/chemrxiv.15002824/v1. Preprint.

[23] C. Yu, L. Cai, and J.-Y. Chen. Stochastic modeling of partially stirred reactor (pasr) for the investigation of the turbulence-chemistry interaction for the ammonia-air combustion. *Flow, Turbulence and Combustion*, 112(2):509–536, 2023. doi:10.1007/s10494-023-00501-7.

[24] X. Zhang, S. P. Moosakutty, R. P. Rajan, M. Younes, and S. M. Sarathy. Combustion chemistry of ammonia/hydrogen mixtures: Jet-stirred reactor measurements and comprehensive kinetic modeling. *Combustion and Flame*, 234: 111653, 2021. doi:10.1016/j.combustflame.2021.111653.

[25] X. Zhang, K. K. Yalamanchi, and S. Mani Sarathy. Combustion chemistry of ammonia/C_1 fuels: A comprehensive kinetic modeling study. *Fuel*, 341:127676, 2023. doi:10.1016/j.fuel.2023.127676.

[26] Y. Zhu, H. J. Curran, S. Girhe, Y. Murakami, H. Pitsch, K. Senecal, L. Yang, and C.-W. Zhou. The combustion chemistry of ammonia and ammonia/hydrogen mixtures: A comprehensive chemical kinetic modeling study. *Combustion and Flame*, 260:113239, 2024. doi:10.1016/j.combustflame.2023.113239.